



HAL
open science

Finite unions of balls with inner and outer margins

Tuong Nguyen

► **To cite this version:**

Tuong Nguyen. Finite unions of balls with inner and outer margins. Signal and Image Processing. Université Grenoble Alpes, 2018. English. NNT : 2018GREAT022 . tel-01871579

HAL Id: tel-01871579

<https://theses.hal.science/tel-01871579>

Submitted on 11 Sep 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ DE GRENOBLE ALPES

Spécialité : **SIGNAL IMAGE PAROLE TELECOMS**

Arrêté ministériel : 25 mai 2016

Présentée par

Tuong NGUYEN

Thèse dirigée par **Dominique ATTALI** et
codirigée par **Isabelle SIVIGNON**

préparée au sein du

**Laboratoire Grenoble Images Parole Signal Automatique
(GIPSA)**

dans l'École Doctorale **Electronique, Electrotechnique,
Automatique, Traitement du Signal (EEATS)**

Unions finies de boules avec marges interne et externe

Finite unions of balls with inner and outer margins

Thèse soutenue publiquement le **27 mars 2018**,
devant le jury composé de:

Monsieur David COEURJOLLY

Directeur de Recherche, CNRS, Président du jury

Monsieur Olivier DEVILLERS

Directeur de Recherche, INRIA, Rapporteur

Monsieur Xavier GOAOC

Professeur des Universités, Université Paris-Est Marne-la-Vallée,
Rapporteur

Monsieur Boris THIBERT

Maître de Conférences, Université Grenoble Alpes, Examineur

Monsieur André LIEUTIER

Ingénieur, Dassault Systèmes, Invité

Madame Dominique ATTALI

Directeur de Recherche, CNRS, Directeur de thèse

Madame Isabelle SIVIGNON

Chargée de Recherche, CNRS, Co-encadrante de thèse



UNIVERSITÉ DE GRENOBLE ALPES
ÉCOLE DOCTORALE EEATS
Électronique, Électrotechnique, Automatique, Traitement du Signal

THÈSE

pour obtenir le titre de

docteur en sciences

de l'Université de Grenoble Alpes

Mention : SIPT

Présentée et soutenue par

Tuong NGUYEN

Unions finies de boules avec marges interne et externe

Thèse dirigée par Dominique ATTALI

préparée au laboratoire Grenoble Images Parole Signal Automatique
(GIPSA)

soutenue le 27 mars 2018

Jury :

<i>Président :</i>	David COEURJOLLY	-	CNRS
<i>Rapporteurs :</i>	Olivier DEVILLERS	-	INRIA
	Xavier GOAOC	-	Université Paris-Est Marne-la-Vallée
<i>Examineurs :</i>	Boris THIBERT	-	Université Grenoble Alpes
<i>Invité :</i>	André LIEUTIER	-	Dassault Systèmes
<i>Directeur :</i>	Dominique ATTALI	-	CNRS
<i>Encadrant :</i>	Isabelle SIVIGNON	-	CNRS

Acknowledgements

As a fairly indecisive individual, I was hesitant and apprehensive of what was to come when I engaged in this PhD. In hindsight, this may be the greatest venture I have undertaken, and I owe a great many people for supporting me these past years.

The most valuable contributors to this work are undeniably Dominique Attali and Isabelle Sivignon. I cannot thank them enough for the guidance they have provided me with, nor the patience they demonstrated as they soldiered on through many of my convoluted ideas. Their expertise was key to turning my trademark *plats de nouilles* into the manuscript you are now reading.

I am very grateful to Olivier Devillers and Xavier Goaoc who have accepted to review this thesis. My meeting with them has greatly influenced the shape this thesis has taken. Many thanks also to David Coeurjolly, Boris Thibert, and André Lieutier for being part of the defense committee. I really appreciate their remarks and the interest they have shown for this work.

My stay at Gipsa-Lab was very fruitful and enriching. I had several interesting and entertaining conversations with Florent and Jocelyn. I also had the opportunity to work in collaboration with Steeve, Li, Quentin and Nadege, on their project.

Finally, on a more personal note, I would like to thank Tuan, Milio, Marine, and Jonathan for producing many animated debates. I am also grateful to Rafaël and Sara for their ability to always cheer me up.

Contents

Acknowledgements	i
Contents	iii
List of Figures	vii
0 Résumé en français	1
Introduction	1
0.1 Section 1.5 : Un résultat sur l'inclusion d'ensembles	4
0.2 Chapitre 2: Union finies de boules	4
0.3 Chapitre 3: Approximation par boules à (δ, ε) -près	5
0.4 Chapitre 4: Algorithme glouton optimal	8
0.5 Chapitre 5: Généralisations à d'autres formes	10
Introduction	11
1 Technical tools	15
1.1 Preliminary notions and notations	15
1.2 Balls and spheres	16
1.3 Medial axis and skeleton	17
1.4 Distance diagrams	18
1.5 A result about set inclusion	19
2 Finite union of balls	21
2.1 Pencil of balls	22
2.2 Convex subsets of pencils	36
2.3 Local inclusion in parabolic pencils	54

2.4	Boundary of a finite union of balls	58
2.5	Medial axis, skeleton, boundary and pencils	64
2.6	Conclusion	81
3	(δ, ε)-ball approximation	85
3.1	Problem statement	85
3.2	Boundary of the erosion	89
3.3	The rational (δ, ε) -ball approximation is in NP	92
3.4	Rational (δ, ε) -ball approximation is NP-hard	95
3.5	Conclusion	100
4	Greedy optimal algorithm for $(0, \varepsilon)$-ball approximation	103
4.1	Preliminary notions	104
4.2	Projection on the medial axis	107
4.3	Testing for set inclusion in finite unions of balls	109
4.4	Critical balls	115
4.5	Critical balls and T -maximal candidates	127
4.6	Computing critical balls	135
4.7	Algorithm	142
4.8	Conclusion	147
5	Generalization to other shapes	149
5.1	Generalized problem and algorithm	150
5.2	Generalities on computing critical balls	153
5.3	Classic point-point pencil	155
5.4	Line-line pencil	164
5.5	Point-line pencil	170
5.6	Conclusion	180

Conclusion	187
Conclusion	187
(O, I) -ball approximations in higher dimensions	188
A Additional technical results	193
A.1 Boundary of union and intersection	193
A.2 Interior of a cone	194
A.3 An alternative way to test the inclusion $(\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B}$	195
Bibliography	200
Abstract	202

List of Figures

1	Generation of a geometric model from a real shape.	11
2	Generation of a ball approximation.	12
2.1	The three types of k -pencils in \mathbb{R}^d . Illustrated for $k = 1$ and $d = 1$	23
2.2	Example of elliptic 1-pencil in \mathbb{R}^2	23
2.3	Example of elliptic 2-pencil in \mathbb{R}^3	24
2.4	Schematic representation of the hyperplanes $H(c_0)$ and $H(c_0, r_0^2)$	32
2.5	Influence of V on $\mathcal{P}(V, c_0, r_0^2)$	33
2.6	Local but non-global inclusion in parabolic pencils.	37
2.7	Boon of \mathcal{B}_0 over \mathcal{B}	38
2.8	Power cell and contribution to the boundary.	40
2.9	Normal cone to a polytope.	41
2.10	Normal cones and power cells.	42
2.11	Image of m_b in a parabolic case.	45
2.12	Inclusion normal cone.	55
2.13	Illustration for the proof of Lemma 2.41.	56
2.14	Lack of “local representation” for ellipses.	59
2.15	Ambiguity for some cardinal-minimum local representations.	62
2.16	Vertex configurations in \mathbb{R}^2	63
2.17	Medial axis in a degenerate configuration.	66
2.18	Star and link in a convex polytope.	67
2.19	Star and link in the boundary of a convex polytope.	69
2.20	Joins for various boundary points.	70
2.21	Illustration for the proof of Lemma 2.49.	72

2.22	Illustration for the proof of Lemma 2.50.	73
2.23	Illustration for the construction of v in the proof of Lemma 2.52.	77
2.24	Illustration for case (a) in the proof of Lemma 2.52.	78
2.25	Illustration for case (b) in the proof of Lemma 2.52.	79
2.26	Non-collinearity of \mathbf{n} and \mathbf{e}' in the proof of Lemma 2.52.	80
2.27	Arc of points x_λ in the proof of Lemma 2.52.	83
2.28	Computing the join of a simple vertex.	84
3.1	(δ, ε) -ball approximation of a shape.	87
3.2	Wavefront, erosion and arrangement.	90
3.3	Wavefront from the inclusion normal cone.	91
3.4	Notations for the proof of Lemma 3.3.	94
3.5	Construction of $\mathcal{S}(G, \delta, \varepsilon)$ from an orthogonal grid drawing.	97
3.6	Canonical covering of an edge gadget.	98
3.7	Grid division into blocks.	100
3.8	Block conversions for vertices.	100
3.9	Block conversions for edges.	101
4.1	T -large and T -small components of a medial ball.	106
4.2	Face join decomposition.	108
4.3	The mapping π	109
4.4	Illustration for Property 4.2.	110
4.5	Illustration for the proof of Lemma 4.6.	114
4.6	$C(S, b, -)$ and $C(S, b, +)$ are not always disjoint.	116
4.7	Projection of a path of balls by π	117
4.8	Critical ball for a point.	119
4.9	Preimage of pcrit	122

4.10	The collection of balls $F_{\text{cover}}(S, A)$.	123
4.11	The collection of balls $F_{\text{crit}}(S, A)$.	123
4.12	Critical balls for a set.	124
4.13	Importance of the assumption $\bar{A} \subseteq \mathring{S}$ for Property 4.16.	127
4.14	The subset $A(S, \varepsilon, \mathcal{B})$.	129
4.15	Illustration for the proof of Lemma 4.19.	131
4.17	Illustration for the proof of Lemma 4.20.	132
4.18	Schematic illustration for the proof of Lemma 4.23.	135
4.19	$T^{\mathcal{P}}$ order on a pencil \mathcal{P} .	136
4.20	Local restriction of $A(S, \varepsilon, \mathcal{B})$.	140
4.21	Illustration of the different steps of the algorithm.	144
5.1	Exemple of general shape	151
5.2	Arc constraint in a line-line pencil.	154
5.3	Point-point pencil.	155
5.4	Point constraint in pp-pencils.	157
5.5	Ball constraint in pp-pencils.	158
5.6	Arc constraint in pp-pencils.	163
5.7	Line-line pencil.	164
5.8	Point constraint in ll-pencils.	166
5.9	Ball constraint in ll-pencils.	168
5.10	Point-line pencil.	170
5.11	Point constraint in pl-pencils.	173
5.12	Ball constraint in pl-pencils: above L .	182
5.13	Ball constraint in pl-pencils: out of bounds.	183
5.14	Ball constraint in pl-pencils: tangential to bounds.	183
5.15	Ball constraint in pl-pencils: last case.	184

5.16 Arc constraint in pl-pencils: support ball above and tangent to L at F	184
5.17 Arc constraint in pl-pencils: support ball below and tangent to L at F	185
5.18 Arc constraint in pl-pencils: support ball through F . (1/2)	185
5.19 Arc constraint in pl-pencils: support ball through F . (2/2)	186

Résumé en français

Sommaire

Introduction	1
0.1 Section 1.5 : Un résultat sur l'inclusion d'ensembles	4
0.2 Chapitre 2 : Union finies de boules	4
0.3 Chapitre 3 : Approximation par boules à (δ, ε)-près	5
0.4 Chapitre 4 : Algorithme glouton optimal	8
0.5 Chapitre 5 : Généralisations à d'autres formes	10

Introduction

Représenter une forme à l'aide d'un ordinateur est une tâche avec de nombreuses applications pratiques. Comparé à l'usage historique de schémas et plans, l'utilisation d'ordinateurs permet en effet de grandement simplifier la façon de concevoir, visualiser, tester et échanger des modèles d'objets physiques [FHK02]. Bien que l'émergence des technologies de CAD/CAM ait principalement été motivée par l'industrie automobile, en particulier pour la conception et production de pièces de carrosserie, les techniques de modélisation géométrique ainsi développées se sont depuis répandues dans d'autres domaines, tels que l'animation par ordinateur et les effets spéciaux dans l'industrie cinématographique [Hab01].

Il existe deux façons principales de créer un modèle numérique pour une nouvelle forme. La première consiste à saisir le modèle à la main. Bien que plusieurs méthodes de saisie et outils dédiés aient été conçus à cet effet, créer un modèle avec une précision suffisante reste une tâche difficile et du ressort des experts. La seconde méthode consiste à scanner un objet réel, ce qui permet en général d'obtenir un nuage de points comme données brutes. Les nuages de points sont un échantillonnage de la surface de la forme qu'ils représentent, et par nature ce sont des représentations de formes lacunaires. En effet, les nuages de points ne fournissent aucune information sur l'objet entre les points échantillonnés. La surface exacte de l'objet est par conséquent indisponible, et il est impossible de distinguer avec certitude les points intérieur et extérieur. Avec si peu d'information sur la géométrie de la forme, il est difficile de manipuler ces représentations de façon contrôlée et significative. Il est donc nécessaire d'appliquer un traitement supplémentaire sur ces nuages de points afin d'obtenir des représentations de forme

plus adéquates. Ce processus est appelé la reconstruction de surface : étant donné un nuage de points, en déduire un modèle numérique qui représente au mieux la forme d'origine. De façon plus générale, étant donné de l'information incomplète sur une forme, comment obtenir la meilleure représentation pour cette forme? Voir Figure 1.




FIGURE 1 – Construction d'un modèle numérique à partir d'une forme réelle.

Il existe de nombreuses méthodes de reconstruction différentes. Cette diversité est directement liée au choix de l'application cible. En effet, des applications différentes vont nécessiter des traitements de formes différents. Il est ainsi courant d'adapter les modèles géométriques utilisés afin d'améliorer les performances d'une application. Nous ne mentionnerons que les schémas de représentation les plus connus, avec quelques exemples de représentations. En raison de ses origines, la représentation des formes dites "solides" est très étudiée dans la discipline [Req77] ; [RR92]. En particulier, toute forme "solide" peut être déterminée de façon unique par son bord, ce qui donne lieu aux représentations de bord (B-rep). Les très populaires B-splines rationnelles non uniformes (NURBS) en sont un exemple. Comme son nom l'indique, le schéma B-rep conserve seulement de l'information sur le bord de la forme. Les applications qui ne se soucient pas des points intérieurs, telles que le rendu d'images par ordinateur, optent souvent pour ce type de représentation. Comme le bord définit la forme de façon unique, il est toujours possible de tester si un point est contenu à l'intérieur de la forme à l'aide de méthodes telles que le lancer de rayons (*ray tracing*). Il est cependant plus efficace d'adopter un schéma différent s'il est nécessaire de traiter en majorité des points intérieurs. On peut alors considérer l'utilisation d'une représentation par maillage. Si les maillages triangulaires ne reflètent que le bord d'une forme 3D, les maillages tétraédriques en reflètent également le volume. Un autre schéma de représentation populaire est la géométrie de construction de solides (CSG), dans laquelle on utilise des formes primitives que l'on combine par des opérations booléennes. En particulier, les unions finies de boules appartiennent à ce schéma.

Dans cette thèse, nous étudions les unions finies de boules. Les boules font partie des formes géométriques les plus simples, et leurs unions héritent partiellement de cette simplicité, contribuant ainsi à la popularité de ce mode de représentation. Un autre facteur de popularité est l'aisance théorique avec laquelle on peut obtenir une approximation par unions finies de boules, et ce de n'importe quelle forme. Une boule b est appelée (inclusion-)maximale pour une forme, si toute boule qui contient b et qui est contenue dans la forme, est égale à b . Si l'on suppose les boules fermées et la forme compacte, alors l'union des boules maximales décrit exactement la forme. La collection des boules maximales est en général de taille infinie, néanmoins un échantillonnage fini de cette collection produit naturellement une approximation de la forme (Figure 2). Au delà de la simplicité théorique, il existe plusieurs algorithmes de conversion prouvés pour obtenir des unions finies de boules à partir d'autres types de

représentation, comme par exemple les nuages de points (reconstruction de surface) [AK00], les maillages [CKM99]; [Hub96] ou les formes numériques [CM07]. Étant donnée l'existence de certains prédicats sur la forme, il est aussi possible d'utiliser une méthode de conversion générique [MGP10].




FIGURE 2 – Construction d'une approximation par boules.

Les unions de boules sont également utiles dans plusieurs applications. Elles sont par exemple utilisées en biochimie pour représenter des molécules [EK05]; [Caz+14], ou en infographie pour la détection de collisions [BO04]; [WZ09]. Pour les tâches gourmandes en calcul, telles que la simulation de processus physiques [Fei+15], l'interpolation ou l'appariement de formes [RF96]; [Cab+09], il est souhaitable d'avoir les représentations de forme les plus compactes qui soient, tout en restant aussi fidèle que possible à la forme d'origine. Avec les unions finies de boules, la taille de la représentation est directement proportionnelle au nombre de boules utilisées. Quant à la qualité de l'approximation, elle est souvent mesurée par diverses quantités, comme la distance de Hausdorff, ou encore la différence de volume entre la forme et sa représentation. Dans certaines applications comme la détection de collision, il peut être important d'imposer des *contraintes géométriques* que l'union de boules doit respecter. Par exemple, on peut contraindre les boules à recouvrir toute la forme, ou au minimum un ensemble de points prédéterminés.

Dans cette thèse, nous allons étudier une nouvelle méthode d'approximation de forme par union finie de boules. Afin de contrôler la précision de l'approximation, nous contraignons l'union de boules à contenir un sous-ensemble interne de la forme, tout en étant contenu dans un sur-ensemble externe. Dans notre approche, nous paramétrons ces ensembles de contraintes par deux nombres réels positifs δ et ε , et appelons par conséquent ces unions des approximations par boules à (δ, ε) -près. Nous définissons formellement ces approximations dans le Chapitre 3, et montrons que calculer des approximations par boules à (δ, ε) -près de cardinal minimum est un problème NP-complet. Nous présentons au Chapitre 4 un algorithme glouton qui est optimal pour des formes simples dans \mathbb{R}^2 , avant de l'étendre à une plus large variété de contraintes internes et externes dans le Chapitre 5. Ces algorithmes sont en grande partie fondés sur les propriétés des boules et de leurs unions. Le Chapitre 2 fournit les fondements théoriques nécessaires à la démonstration de ces propriétés. Nous y décrivons ainsi la structure des unions finies de boules, sans hypothèse de position générale. Nous rappelons dans le Chapitre 1 plusieurs définitions et structures qui sont utilisées dans cette thèse. Ce premier chapitre présente également un résultat générique sur les inclusions d'ensembles, et dont la spécialisation aux unions de boules justifie la faisabilité de notre algorithme.

0.1 Section 1.5 : Un résultat sur l'inclusion d'ensembles

Nous démontrons à présent un résultat général qui caractérise l'inclusion d'un ensemble dans un autre. Nous n'utilisons que peu d'hypothèses sur ces deux ensembles. Plus précisément, ils doivent être des sous-ensembles d'un espace topologique connexe. Dans ce cas, il suffit de vérifier des conditions d'inclusion "locale" au voisinage du bord des ensembles pour obtenir l'inclusion "globale".

Lemme 0.1. *Soient A et $X \neq \emptyset$ deux sous-ensembles d'un espace topologique connexe. Si*

$$(i) \quad \partial A \subseteq \overline{X}$$

$$(ii) \quad \forall x \in \partial X, \exists N_x \text{ un voisinage de } x \text{ tel que } N_x \cap A \subseteq X$$

alors $A \subseteq X$.

Démonstration. Soit $H = A \cap X^c$, la preuve se déroule en deux étapes :

(a) Dans un premier temps, on montre que $\partial H \subseteq \partial X$. Comme $\partial(A \cap B) \subseteq (\partial A \cap \overline{B}) \cup (\overline{A} \cap \partial B)$ (voir Lemme A.1), on peut injecter $B = X^c$ pour obtenir $\partial H \subseteq (\partial A \cap \overline{X^c}) \cup (\overline{A} \cap \partial X^c)$. Puisque $\partial A \subseteq \overline{X}$ et $\partial X^c = \partial X$, on en déduit alors $\partial H \subseteq \partial X$.

(b) Nous montrons à présent que $H = \emptyset$. Supposons par contradiction que $H \neq \emptyset$. Comme X est non vide, X^c ne peut pas être l'espace topologique entier, ce qui implique que H , à son tour, ne peut pas être l'espace entier. Par conséquent, $\partial H \neq \emptyset$ et il existe $x \in \partial H$. Comme $\partial H \subseteq \partial X$, nous pouvons utiliser l'hypothèse (ii) sur le point x , et il existe un voisinage N_x de x tel que $N_x \cap A \subseteq X$. Puisque $x \in \partial H$, nécessairement $N_x \cap H \neq \emptyset$. Cependant, $N_x \cap H = N_x \cap A \cap X^c \subseteq X \cap X^c = \emptyset$, et l'on aboutit à une contradiction. \square

Il est à noter que l'inclusion $A \subseteq X$ implique à la fois (i) et (ii), on obtient donc une équivalence. Si l'on travaille avec une classe de formes spécifique, il peut être avantageux de spécialiser ce lemme à cette classe. En particulier, nous verrons dans la Section 4.3 une spécialisation du Lemme 0.1 dans laquelle l'ensemble extérieur X est une union finie de boules.

0.2 Chapitre 2 : Union finies de boules

§ Introduction Les unions finies de boules sont au centre des considérations et discussions qui seront menées tout au long de cette thèse. Nous dédions par conséquent ce chapitre à l'étude de leur structure, ainsi que quelques unes de leurs propriétés. À cet effet, nous allons baser notre approche sur les faisceaux de boules. Les faisceaux sont des familles affines de boules avec de nombreuses propriétés, et nous montrons que toute union finie de boules peut être décomposée en un nombre fini de sous-ensembles de faisceaux. En étudiant soigneusement

les propriétés des faisceaux, nous aboutissons à une description du bord des unions finies de boules qui reste valide sans hypothèse de position générale. Cette description du bord donne ensuite lieu à une description de l'axe médian en termes de faisceaux, elle aussi valide sans hypothèse de position générale. Au final, l'usage des faisceaux nous permet d'utiliser une terminologie et des outils capables de traiter tous les cas dégénérés qui peuvent survenir au sein d'une union finie de boules.

§ Conclusion Au cour de ce chapitre, nous nous sommes intéressés aux faisceaux de boules et à leurs nombreuses propriétés. En tirant parti de l'existence de représentations locales dans une union finie de boules S , nous sommes en mesure de décrire le voisinage de tout point du bord de S à l'aide de faisceaux. Par ce procédé, les configurations dégénérées de boules de \mathbb{R}^d produisent des faisceaux paraboliques. Gérer les dégénérescences des unions finies de boules peut donc se ramener à l'étude des propriétés et des configurations de faisceaux paraboliques. Grâce aux nombreuses propriétés des faisceaux de boules, et leur simple définition comme sous-espaces affines de \mathbb{R}^{d+1} , cette façon de décrire les unions finies de boules permet de simplifier l'étude des cas dégénérés. En particulier, nous avons obtenu des descriptions robustes du bord et de l'axe médian des unions finies de boules. Ces descriptions montrent qu'il existe une relation duale entre le bord et l'axe médian.

Dans la continuité de ce chapitre, il serait intéressant d'explorer plus en détail cette relation duale. Nous présumons qu'une telle étude puisse aboutir à un algorithme efficace pour le calcul de l'axe médian d'une union finie de boules. Suite aux résultats présentés dans ce chapitre, il est en effet possible de calculer l'axe médian, y compris dans les configurations dégénérées. Cependant, nous ne nous sommes pas encore intéressés aux meilleures façons de réaliser ce calcul.

0.3 Chapitre 3 : Approximation par boules à (δ, ε) -près

§ Introduction Maintenant que nous avons introduit nos principaux outils sur les unions de boules, nous abordons le problème que nous allons étudier dans le reste de cette thèse, c'est-à-dire les approximations par boules à (δ, ε) -près. Étant donné un ensemble S dans \mathbb{R}^d , nous appelons approximation par boules à (δ, ε) -près de S , une collection de boules qui contient l'érosion morphologique de S (par une boule de rayon ε), et qui est contenu dans le dilaté morphologique de S (par une boule de rayon δ). Nous étudions le problème de calculer de telles approximations pour certaines classes de formes (en particulier des formes qui sont elles-mêmes des unions finies de boules). Ce problème est proche des problèmes de couverture par ensembles géométriques (*geometric set cover*), même s'il reste de nature différente. Nous verrons que ce problème offre un nouveau cadre pour simplifier une collection de boules, tout en contrôlant à la fois la distance interne et externe à la collection d'origine. Dans ce chapitre, nous évoquons quelques propriétés générales de ces approximations, et montrons en particulier que calculer une approximation par boules à (δ, ε) -près de cardinal minimum est déjà NP-complet pour $d = 2$. Dans les chapitres suivants, nous étudierons un algorithme polynomial qui calcule de

telles approximations dans certains cas particuliers.

§ Approximations par boules Dans ces travaux, nous introduisons une nouvelle façon de contraindre une collection de boules de sorte que son union soit proche de la forme. L'approche consiste à forcer les boules à couvrir un sous-ensemble de la forme, tout en restant à l'intérieur d'un sur-ensemble de la forme, permettant ainsi aux boules de ne couvrir que partiellement la forme, mais aussi à en dépasser les bords. Étant donné $S \subseteq \mathbb{R}^d$ et un nombre réel $r \geq 0$, la *dilatation* de S par une boule de rayon r est $S^{\oplus r} = \cup_{x \in S} b(x, r)$ et l'*érosion* de S par une boule de rayon r est $S^{\ominus r} = \{x \mid b(x, r) \subseteq S\}$. Lorsque S est fermé, alors $S^{\ominus r}$ et $S^{\oplus r}$ sont également fermés. Dorénavant, nous désignerons également $S^{\ominus r}$ et $S^{\oplus r}$ par les appellations de r -érode et r -dilaté.

Définition 14. Soit $S \subseteq \mathbb{R}^d$, $\varepsilon \geq 0$ et $\delta \geq 0$. Une collection de boules \mathcal{B} est une *approximation par boules à (δ, ε) -près* de S si $S^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B} \subseteq S^{\oplus \delta}$.




FIGURE 3 – Une forme (bleue), son ε -érode (orange), son δ -dilaté (délimité par la courbe en tirets bleus) et une approximation par boules à (δ, ε) -près (les trois boules vertes).

Faisons quelques remarques. Considérons \mathcal{B} une approximation par boules à (δ, ε) -près de S , voir Figure 3 pour un exemple. Soit $B = \bigcup \mathcal{B}$, nous avons alors $B \subseteq S^{\oplus \delta}$ et $(S^{\ominus \varepsilon})^{\oplus \varepsilon} \subseteq B^{\oplus \varepsilon}$. On en déduit que si $S = (S^{\ominus \varepsilon})^{\oplus \varepsilon}$, on a alors $B \subseteq S^{\oplus \delta}$ et $S \subseteq B^{\oplus \varepsilon}$, ou de façon équivalente la distance de Hausdorff entre B et S est inférieure ou égale au maximum entre δ et ε , *i.e.* $d_H(S, B) \leq \max\{\delta, \varepsilon\}$. La condition $S = (S^{\ominus \varepsilon})^{\oplus \varepsilon}$ peut être vue comme une hypothèse de régularité¹ sur S . En effet, l'égalité $S = (S^{\ominus \varepsilon})^{\oplus \varepsilon}$ est satisfaite si tout point de S peut être couvert par une boule de rayon ε contenue dans S . En particulier, cette condition est vérifiée

¹Un ensemble S est dit *r-régulier* si $S = (S^{\ominus r})^{\oplus r} = (S^{\oplus r})^{\ominus r}$ [Ser82]. Autrement dit, une forme S est *r-régulière* si elle reste inchangée après une ouverture ou une fermeture morphologique par une boule de rayon r .

si toutes les boules médiales ont un rayon supérieur ou égal à ε , condition elle-même vérifiée dès que le *reach* de S est supérieur ou égal à ε . Pour rappel, le *reach* d'un ensemble S a été introduit par Federer [Fed59] et est l'infimum des distances entre les points de l'axe médian de S et les points du complémentaire de S . On le dénote par $\text{reach}(S)$. Nous résumons nos observations ci-dessus dans la remarque suivante :

Remarque 44. Soit \mathcal{B} une approximation par boules à (δ, ε) -près de S . Si $\text{reach}(S) \geq \max\{\delta, \varepsilon\}$ alors $d_H(\bigcup \mathcal{B}, S) \leq \max\{\delta, \varepsilon\}$.

Remarque 45. Si S est compact et si $\delta + \varepsilon > 0$, alors S admet des approximations par boules à (δ, ε) -près de cardinal fini.

Supposons que S soit compact et $\delta + \varepsilon > 0$. Pour construire une approximation par boules à (δ, ε) -près de S , considérons la collection des boules ouvertes de rayon $\delta + \varepsilon$ et dont le centre appartient à $S^{\ominus \varepsilon}$. Par construction, cette collection recouvre $S^{\ominus \varepsilon}$ qui est compact, par compacité de S , ce qui implique que la collection de boules admet un sous-recouvrement fini. En prenant l'adhérence des boules d'un tel sous-recouvrement, on obtient une collection de boules fermées, toutes contenues dans $S^{\oplus \delta}$, et qui constitue donc une approximation à (δ, ε) -près de S de cardinal fini.

§ **Résultat principal** Dans cette section, nous nous intéressons à une variante discrète du problème suivant :

Problème 1 (Approximation par boules à (δ, ε) -près). Étant donné un sous-ensemble $S \subseteq \mathbb{R}^d$, deux nombres réels positifs δ et ε , calculer une approximation par boules à (δ, ε) -près de S de cardinal minimum.

Notre résultat principal est qu'une version discrète du problème de décision associé à ce problème est NP-complet. Afin d'énoncer ce problème, nous qualifions de *rationnelle* toute collection de boules dont les centres ont des coordonnées rationnelles et un rayon rationnel.

Problème 2 (Problème de décision des boules rationnelles à (δ, ε) -près). Étant donné un sous-ensemble $S \subseteq \mathbb{R}^d$ qui est l'union d'une collection de boules finie et rationnelle, deux nombres réels positifs ε et δ , et un entier $k > 0$, répondre à la question suivante : est-ce que S admet une approximation par boules à (δ, ε) -près qui soit rationnelle et de cardinalité inférieure ou égale à k ?

Nous montrons que ce problème est déjà difficile en dimension 2 :

Théorème 5. Le problème de décision des boules rationnelles à (δ, ε) -près dans \mathbb{R}^2 est NP-complet.

§ **Conclusion** Dans notre preuve de dureté, nous avons soigneusement choisi l'ensemble des disques qui constituent la forme, de telle sorte que son érosion ne contienne qu'un ensemble fini de points (les rotules). Il s'ensuit que la réduction depuis le problème de couverture par

sommets induit un problème de couverture par ensembles, pour un ensemble fini de points (les rotules), avec une famille infinie de boules (l'axe médian du dilaté).

On observe cependant qu'il est possible de réaliser une construction similaire, dans laquelle l'érodé de la forme est également un ensemble infini de points. De plus, une construction où la forme, son érosion et sa dilatation ont tous le même type d'homotopie est aussi possible. Pour obtenir ces variantes, il suffit d'ajuster les rayons des disques qui composent les gadgets, et les distances entre disques consécutifs, de telle sorte que les huit propriétés (i)-(viii) soient vérifiées. Ces changements n'ont pas d'impact sur le résultat final.

Étant donné que l'approximation par boules rationnelles à (δ, ε) -près est NP-complet, il y a peu de chance de trouver un algorithme polynomial pour résoudre ce problème de façon générale. Néanmoins, pour des formes suffisamment simples, il est possible de le résoudre avec des complexités en temps et en espace polynomiales. Dans le Chapitre 4, nous présentons un algorithme qui opère sur les unions finies de boules sans cycle, et capable de calculer des approximations par boules à $(0, \varepsilon)$ -près. En nous reposant sur les notions introduites par ce premier algorithme, nous l'étendons ensuite à une classe de formes plus diverse dans le Chapitre 5.

0.4 Chapitre 4 : Algorithme glouton optimal

§ **Introduction** Dans ce chapitre, nous présentons un algorithme qui permet de calculer une solution de cardinal minimum pour certaines instances du problème d'approximation par boules à (δ, ε) -près dans \mathbb{R}^2 . Comme le problème est NP-complet, nous opérons avec certaines hypothèses simplificatrices. Ainsi, nous supposons $\delta = 0$ et $\varepsilon > 0$. Nous restreignons également la classe de formes pour S . Plus précisément, nous nous intéressons seulement aux unions finies de disques qui ne contiennent pas de cycle dans leur axe médian. Il est déjà bien connu que les sous-ensembles ouverts et bornés de \mathbb{R}^d ont le même type d'homotopie que leur axe médian [Lie04], et par conséquent l'intérieur de ces formes ne contient pas de cycle non plus. Cependant, nous considérons ici des unions de disques fermés, pour lesquelles la configuration illustrée dans la Figure 4.6 peut survenir. Il s'ensuit que la forme peut contenir des cycles, malgré le fait que son axe médian n'en contienne pas.

Notre résultat principal est le suivant.

Théorème 7. Sous le modèle Real-RAM, il existe un algorithme en temps polynomial pour calculer des approximations par boules à (δ, ε) -près de cardinal minimum pour les unions finies de disques dont l'axe médian est sans cycle.

Remarque 46. La plupart des algorithmes et structures de données qui sont utilisés en géométrie algorithmique sont décrits sous le modèle Real-RAM [T04]; [BY95]. Ce modèle suppose que tout nombre réel peut être représenté par une seule cellule mémoire, et que les opérations arithmétiques usuelles peuvent être réalisées en temps constant.

Tout au long de ce chapitre, S désignera une union finie de disques dont l'axe médian est

sans cycle. Dans la Section 4.1, nous présentons l'intuition sous-jacente à notre algorithme. Les définitions qui découlent de cette intuition ne sont toutefois pas adaptées à une implémentation algorithmique. Nous introduisons ainsi plusieurs outils et propriétés dans les Sections 4.2 et 4.3, que nous utilisons ensuite pour définir et étudier les propriétés des boules critiques dans les Sections 4.4 à 4.6. Enfin, nous décrivons notre algorithme dans la Section 4.7.

§ Conclusion Pour arriver à l'Algorithme 1, nous sommes partis de la notion intuitive de boules candidates et sommes passés par la notion plus générale de boules critiques. À partir des diverses propriétés de ces boules, nous sommes parvenus à montrer que notre algorithme glouton calcule des solutions de cardinal minimum pour une certaine classe de formes. Nous commentons à présent quelques hypothèses qui sont faites dans notre approche, en rapport avec cette classe de formes.

Pour commencer, abordons le sujet de l'ordre partiel T défini sur l'axe médian. D'après la description fournie dans la Section 4.1, tout point de l'axe médian peut être choisi comme racine de T , et donc définir une orientation et un ordre partiel. Cette variabilité est sans conséquence sur notre méthode, et le choix de la racine est entièrement libre. Indépendamment de ce choix, l'orientation et l'ordre partiel qui en résultent auront toutes les propriétés requises pour le bon fonctionnement de notre algorithme. En pratique, changer la position de la racine peut avoir une influence sur la position des boules qui seront sélectionnées, mais le cardinal de la collection finale reste invariant.

Si l'on veut s'affranchir de l'hypothèse d'acyclicité de $\text{MA}(S)$, on peut facilement étendre l'Algorithme 1 afin de gérer les cycles. Toutefois, il est à noter que cette extension ne préserve pas la garantie d'optimalité de la solution obtenue. Pour cela, nous rappelons que $\text{MA}(S)$ peut s'interpréter comme le dessin d'un graphe. En théorie des graphes, un *coupe-cycles de sommets* d'un graphe $G = (V, E)$ est un sous-ensemble F des sommets qui contient au moins un sommet dans chaque cycle de G . Il est connu que calculer un coupe-cycles de sommets de cardinalité minimum est un problème NP-complet [Kar72]. Étant donné un tel coupe-cycles de sommets F , nous avons alors la possibilité d'initialiser notre algorithme à partir de la collection \mathcal{B}_F qui contient toutes les boules médiales dont les centres sont un sommet dans F . Ceci a pour conséquence de supprimer tous les cycles de l'axe médian, et l'on peut alors définir une orientation sur $\text{MA}(S)$ en choisissant une racine dans $\text{MA}(S) \setminus F$. À partir de là, on est en mesure d'appliquer le reste de notre algorithme après avoir calculé l'ensemble de contrainte initiale $A(S, \varepsilon, \mathcal{B}_F)$. Soit $\mathcal{B}_{\text{algo}}$ la collection de boules obtenues par notre algorithme. Les sections précédentes garantissent que $\mathcal{B}_{\text{algo}} \setminus \mathcal{B}_F$ est de cardinal minimum parmi les collections de boules dans S qui couvrent $S^{\ominus \varepsilon} \setminus \bigcup \mathcal{B}_F$. Il s'ensuit que $\mathcal{B}_{\text{algo}}$ dépasse le cardinal des approximations optimales à $(0, \varepsilon)$ -près de S d'au plus $|\mathcal{B}_F| = |F|$ boules.

Bien que nous ayons supposé $\delta = 0$, on remarque que cette hypothèse n'a que peu d'importance dans notre approche, hormis la simplification des notations. En effet, tant que l'axe médian de $S^{\oplus \delta}$ est sans cycle, on peut alors substituer S par $S^{\oplus \delta}$ dans la chapitre ci-dessus. Comme l'hypothèse d'acyclicité doit être satisfaite par $\text{MA}(S^{\oplus \delta})$ au lieu de $\text{MA}(S)$, utiliser une valeur $\delta > 0$ restreint seulement les formes S pour lesquelles on obtient une solution

optimale.

Compte tenu des observations précédentes, l’Algorithme 1 n’a en réalité besoin de boucler que sur deux listes : une liste de faisceaux de la forme contenant en ordre topologique, et une liste d’éléments du bord de la forme contenue. Tant que les prédicats sont disponibles, notre méthode peut ainsi s’adapter à d’autres classes de formes. Ceci est précisément le sujet du prochain Chapitre 5.

0.5 Chapitre 5 : Généralisations à d’autres formes

§ Introduction Dans ce chapitre, nous montrons enfin comment implémenter les prédicats nécessaires au calcul de boules critiques. En fait, nous verrons que ces prédicats peuvent être implémentés pour plusieurs classes de familles de boules, et ne sont pas restreints aux 1-faisceaux elliptiques qui sont requis pour les unions finies de boules. Par conséquent, nous étendons de fait l’algorithme présenté précédemment à une plus large variété de formes, et qui englobe en particulier les formes polygonales.

Nous commençons par quelques définitions dans la Section 5.1, afin d’étendre notre schéma d’approximation et les notions en rapport aux boules critiques à ces nouvelles formes. Dans la Section 5.2 nous exposons la stratégie générale pour établir les prédicats nécessaires à notre algorithme, que nous appliquons ensuite aux diverses familles de boules médiales dans les Sections 5.3 à 5.5.

§ Conclusion Étant donné que le calcul de boules critiques ne repose que sur quelques prédicats, il est possible de généraliser l’Algorithme 1, que nous avons utilisé pour calculer des approximations par boules à (δ, ε) -près pour des unions finies de boules, à d’autres classes de formes. Il suffit en effet de montrer que ces prédicats peuvent être implémentés pour ces nouvelles classes de formes. Afin de faciliter cette généralisation, il est avantageux de découpler au maximum les ensembles de contrainte interne et externe, menant ainsi au concept d’approximation par boules à (O, I) -près.

On peut montrer que l’axe médian d’un polygone peut être décomposé en portions de faisceaux point-point, ligne-ligne, et point-ligne. Notre étude des prédicats pour ces familles de boules permet donc d’étendre notre approche aux polygones. Il est à noter que nous n’avons pas mené d’étude détaillée sur ces nouveaux faisceaux, et qu’il reste donc une incertitude sur l’optimalité et la convergence de l’algorithme étendu. Nous laissons ainsi ces deux propriétés en tant que conjectures.

Une extension naturelle de notre algorithme dans \mathbb{R}^2 serait d’étudier l’existence de prédicats pour les faisceaux “arc-point”, “arc-ligne”, et “arc-arc”, permettant ainsi aux ensembles de contrainte interne et externe de couvrir la même classe de formes.

Introduction

The task of representing a shape with the help of a computer has several important applications. Indeed, compared to the historical use of drawn blueprints, the use of computers has greatly improved and simplified the way to create, visualize, test and exchange models of physical objects [FHK02]. Although the advent of CAD/CAM technology was greatly driven by the car industry, to help with the design and manufacturing process of car bodies, the techniques developed in geometric modeling since extended to other domains, such as computer animation and special effects for the film industry [Hah01].

There are two main ways to create a digital model for a new shape. One is direct, manual input. Though special tools and input methods have been devised, obtaining a model with the desired accuracy can be a daunting task, or require a high degree of skill from the operator. Another method consists in scanning a real object, which usually yields a point cloud as raw data. Because point clouds are samples of the surface of the shape they represent, they are by nature incomplete. Indeed, this representation contains no information about the object in between the sample points. The exact surface is thus unavailable, and it is impossible to distinguish interior and exterior points with certainty. Because there is little actual information on the geometry of the shape, manipulating this representation in a controlled and meaningful way can be difficult. Point clouds thus require additional processing in order to derive an appropriate shape representation. This is precisely the problem of surface reconstruction: given a point cloud, infer a digital model that best describes the original shape. More generally, given some incomplete information about the shape, derive the best representation for that shape (Figure 1).




Figure 1 – Generation of a geometric model from a real shape.

A number of different reconstruction methods exists. One impetus for this variety is the target application of the geometric model. Different applications will indeed require different types of shape processing, and it is common to tailor the model to the application to improve performance. We shall only mention the most well-known schemes, along with a few examples of shape representations. Due to its origins, developments in this field have a keen interest in representing “solid” shapes [Req77]; [RR92]. In particular, any such shape is uniquely determined by its boundary, thus leading to boundary representations (B-rep). The widely popular non-uniform rational basis splines (NURBS) fall into this category. As the name implies, B-rep schemes only store information regarding the boundary of the shape. Applications that do not care about interior points, such as rendering for an animation, often

favour such representations. Of course, because the boundary uniquely defines the shape, one can still test for interior points with methods such as ray casting. It is however more efficient to adopt a different scheme if a lot of processing on interior points is required. One possibility is to use polygonal meshes. Although triangle meshes only capture the surface of a 3D shape, tetrahedral meshes also capture the volume. Another popular representation scheme is constructive solid geometry (CSG). It uses primitive shapes and combines them together through boolean operations. In particular, finite unions of balls belong to this scheme.

In this thesis, we focus on finite unions of balls. Balls are among the simplest geometric shapes, and their unions partly retain that simplicity. This greatly contributed to the popularity of finite unions of balls. Another factor to that popularity is the theoretical ease with which one can derive an approximation by finite union of balls from any shape. A ball b is said to be (inclusion-wise) *maximal* for a given shape if every ball that contains b and is contained in the shape is equal to b . Assuming that balls are closed and the shape is compact, the union of maximal balls provides an exact description of the shape, whose size is in general infinite. Nevertheless, finite samplings of the collection of maximal balls are naturally an approximation of the shape (Figure 2). Beyond simple theoretical simplicity, there exists many provably good conversion algorithms that output finite unions of balls from other representations, including point clouds (surface reconstruction) [AK00], polygonal meshes [CKM99]; [Hub96] or digital shapes [CM07]. Given the existence of some predicates on the shape, one can also use a general conversion method [MGP10].




Figure 2 – Generation of a ball approximation.

Finite unions of balls are also useful for a number of applications. For instance they are used in biochemistry to model molecules [EK05]; [Caz+14] or in computer graphics to detect collisions between objects [BO04]; [WZ09]. For computationally demanding tasks, such as the simulation of physical processes [Fei+15], shape interpolation or shape matching [RF96]; [Cab+09], it is desirable that shape representations have a small size, while still providing an accurate approximation of the shape. With finite unions of balls, size of the representation directly translates to the cardinal of the collection of balls. The quality of the approximation is usually measured by various quantities such as the Hausdorff distance or the difference in volume between the shape and its representation. In other situations such as collision detection, it may be important that the union of balls satisfies some *geometric constraints*. For instance, we may want the balls to cover the shape or at least a prescribed set of points.

In this thesis, we will investigate a novel approach to perform approximation by finite union of balls. To control the fidelity of the approximation to the original shape, we force the union of balls to contain an inner subset of the shape, while still being contained within an

outer superset. As we parameterize these constraint sets by two non-negative real numbers δ and ε , we call these unions of balls (δ, ε) -ball approximations. We formally introduce them in Chapter 3, and show that finding such an approximation with minimum cardinality is an NP-complete problem. We provide in Chapter 4 a greedy optimal algorithm that outputs these cardinal minimum approximations for simple shapes in \mathbb{R}^2 , and then extend this algorithm to a wider class of inner and outer constraint sets in Chapter 5. These algorithms will rely heavily on the properties of balls and their unions. Chapter 2 provides the theoretical foundation required to demonstrate these properties; we describe there the structure of finite unions of balls without relying on the general position assumption. We recall in Chapter 1 some definitions and structures that we use throughout this thesis. This first chapter in particular provides a generic result regarding set inclusion, and whose specialization to unions of balls is the underlying reason for the feasibility of our algorithm.

Technical tools

Contents

1.1 Preliminary notions and notations	15
1.2 Balls and spheres	16
1.3 Medial axis and skeleton	17
1.4 Distance diagrams	18
1.5 A result about set inclusion	19

1.1 Preliminary notions and notations

We recall here several basic notions and definitions on which we will build upon.

§ **Euclidean structure** \mathbb{R}^d is endowed with the Euclidean structure. We denote the inner product of vectors \mathbf{u} and \mathbf{v} by $\langle \mathbf{u}, \mathbf{v} \rangle$, and the norm of \mathbf{u} by $\|\mathbf{u}\|$. For any two points $x, y \in \mathbb{R}^d$, their Euclidean distance is $\|x - y\|$. More generally for two sets $X, Y \subseteq \mathbb{R}^d$, we also denote the distance from X to Y by $d(X, Y)$, and simplify the notation to $d(x, Y)$ if X is reduced to the singleton $\{x\}$.

§ **Set notations** For any subset $X \subseteq \mathbb{R}^d$, we respectively denote its closure, interior, complement, and boundary by \overline{X} , $\overset{\circ}{X}$, X^c and ∂X .

§ **Lines, half-lines and segments** Let $x, y \in \mathbb{R}^d$ be two distinct points. We denote by (xy) the line that goes through x and y , and $[xy)$ the half-line originating from x and that goes through y . The closed line segment joining x and y is $[xy]$, the open line segment is $]xy[$, and $[xy[$ denotes the line segment closed at x , open at y .

§ **Affine subspaces** Let $X \subseteq \mathbb{R}^d$ be an affine subspace of \mathbb{R}^d . When X has dimension k , $0 \leq k \leq d$, we call X a k -flat for short. Additionally, we denote by \vec{X} the vector subspace

associated to X . \vec{X} is also sometimes called the direction of X .

Given any subset Y of \mathbb{R}^d , we denote its affine hull by $\text{aff}(Y)$, and the associated vector subspace of that affine hull by $\vec{\text{aff}}(Y)$.

1.2 Balls and spheres

§ Definition A d -ball in \mathbb{R}^d with center $c \in \mathbb{R}^d$ and radius $r \geq 0$ is the collection of points $\{x \in \mathbb{R}^d \mid \|x - c\| \leq r\}$. We denote it by $b(c, r)$. Given an integer k , $0 \leq k \leq d - 1$, a point $c \in \mathbb{R}^d$, a k -flat V passing through c , and a real number $r \geq 0$, the k -ball centered at c with radius r and support space V is the intersection between the d -ball $b(c, r)$ and V . We denote it by $b(c, r, V) = b(c, r) \cap V$. For k , $-1 \leq k \leq d - 1$, a k -sphere is the relative boundary of a $(k + 1)$ -ball. The radius, center and support space of a sphere is inherited from the corresponding ball. For instance, a 0-ball is a singleton, a 1-ball is a line segment, a 2-ball is a disk and a 3-ball is a regular ball in three dimensions. It follows that a 0-sphere is a pair of points, a 1-sphere is a circle and a 2-sphere is the usual sphere in \mathbb{R}^3 . Note that in our definition, we consider (-1) -spheres to be well defined: they are equal to the singleton defined by their corresponding 0-ball. Also, any ball (sphere) of any dimension with null radius only contains its center. In these particular cases, we say that the ball (sphere) is *degenerate*.

§ Power of a ball Consider a d -ball $b = b(c, r)$ and a point $x \in \mathbb{R}^d$. Then the *power* of x with respect to b is

$$\text{pow}(x, b) = \|x - c\|^2 - r^2$$

It is straightforward to see that the power of a point is negative when it belongs to the interior of the ball, null when it is on its boundary, and positive when in its complement.

In particular it is possible to fully characterize a ball b by its center c and the power of the origin to b , $\text{pow}(O, b) = \|c\|^2 - r^2$. We can thus define a space of balls using the identity

$$b(c, r) \iff \left(c, \|c\|^2 - r^2 \right) \quad (1.1)$$

We will thus use the same notation b to designate both the d -ball of \mathbb{R}^d with center c and radius r , and the point of \mathbb{R}^{d+1} $(c, \|c\|^2 - r^2)$. Given a ball $b \in \mathbb{R}^{d+1}$, we denote by $c(b)$ its center.

Consider $(x_1, \dots, x_d, x_{d+1}) \in \mathbb{R}^{d+1}$, and let \mathcal{P} be the paraboloid of equation $x_{d+1} = \sum_{i=1}^d x_i^2$. Any point lying below \mathcal{P} corresponds to a d -ball with positive radius, the points of \mathcal{P} are balls with null radius, whereas points above \mathcal{P} are *imaginary balls*. Using the identity (1.1), we indeed obtain $r^2 < 0$, implying that r is a pure imaginary number. The interpretation of imaginary balls is beyond the scope of this work, and we will mostly ignore them thereafter. As such, we consider imaginary balls to be empty.

A natural operation to perform on subsets of \mathbb{R}^{d+1} is to retrieve the collection of the centers of the balls. To do so it suffices to project the subset onto the hyperplane of equation $x_{d+1} = 0$. Later on, we let $c(\mathcal{B})$ denote the collection of centers of balls in \mathcal{B} , $c(\mathcal{B}) = \{c(b) \mid b \in \mathcal{B}\}$.

§ Bisector of two balls Consider two distinct balls $b_1 = b(c_1, r_1)$ and $b_2 = b(c_2, r_2)$, with respective centers c_1, c_2 , and radii r_1 and r_2 . It is known that the locus of points with equal power from b_1 and b_2 is a hyperplane called the *bisector* of b_1 and b_2 , and denoted $H(b_1, b_2) = \{x \in \mathbb{R}^d \mid \text{pow}(x, b_1) = \text{pow}(x, b_2)\}$. Indeed we have

$$\begin{aligned} \text{pow}(x, b_1) &= \text{pow}(x, b_2) \\ \iff \|x\|^2 - 2\langle x, c_1 \rangle + \|c_1\|^2 - r_1^2 &= \|x\|^2 - 2\langle x, c_2 \rangle + \|c_2\|^2 - r_2^2 \\ \iff 2\langle x, c_2 - c_1 \rangle + \|c_1\|^2 - \|c_2\|^2 - r_1^2 + r_2^2 &= 0 \end{aligned}$$

It follows that $H(b_1, b_2)$ is indeed a hyperplane and admits $c_2 - c_1$ as a normal vector. For reasons that will become clear hereafter, we shall also call this hyperplane the *radical hyperplane* of b_1 and b_2 .

Remark 1. When the two balls b_1 and b_2 have a non-empty intersection, the hyperplane $H(b_1, b_2)$ contains $\partial b_1 \cap \partial b_2 = H(b_1, b_2) \cap \partial b_1$. To see this, notice that the sphere bounding a ball corresponds to points which have null power to that ball. Hence, for all $x \in \partial b_1 \cap \partial b_2$, $\text{pow}(x, b_1) = 0 = \text{pow}(x, b_2)$ and $x \in H(b_1, b_2)$. Conversely, if $x \in H(b_1, b_2) \cap \partial b_1$, $\text{pow}(x, b_1) = \text{pow}(x, b_2) = 0$ and $x \in \partial b_1 \cap \partial b_2$.

1.3 Medial axis and skeleton

Let S be a subset of \mathbb{R}^d . The notion of medial axis and/or skeleton of S has been widely studied ([CL05]; [ABE09]; [CCT11]; [DZ04]...) but unfortunately their precise definitions often vary from an author to another. We thus dedicate this section to specifying the naming convention that we will adopt.

Consider a ball $b \subseteq S$. We say that b is *medial* in S if $|\partial b \cap \partial S| \geq 2$. In other words, b is a medial ball for S if its boundary intersects the boundary of S at least twice. The *medial axis* of S is then the collection of centers of the medial balls of S ,

$$\text{MA}(S) = \{c(b) \mid b \text{ medial in } S\}$$

On the other hand, we say that b is (inclusion-wise) *maximal* in S if any ball that contains b and is contained in S , is equal to b . We call *skeleton* of S the collection of centers of maximal balls of S ,

$$\text{Sk}(S) = \{c(b) \mid b \text{ inclusion-wise maximal in } S\}$$

These two notions of medial and maximal balls for a shape are similar, but strictly distinct. Indeed, any medial ball is also inclusion-wise maximal, however the converse does not hold.

A ball can effectively be inclusion-wise maximal, but only intersect the boundary of S at one unique point. In general we have the inclusions [Mat88]:

$$\text{MA}(S) \subseteq \text{Sk}(S) \subseteq \overline{\text{MA}(S)}$$

As was mentioned previously, any compact shape S can exactly be described as the union of its inclusion-maximal balls, $S = \bigcup_{c(b) \in \text{Sk}(S)} b$. Due to the existence of inclusion-maximal balls that are not medial, the union of medial balls may not perfectly reconstruct S . The relationship between maximal and medial balls however ensures that the closure of the family of medial balls will reconstruct S (when it is compact). Because the medial axis is generally simpler to compute directly than the skeleton, some authors opt to study the object $\overline{\text{MA}(S)}$.

The popularity of medial representations can be attributed to two properties. Those collection of balls (nearly) reconstruct the shape, and any of their ball is intuitively a “largest” ball of the shape. Any application that requires a small number of balls will thus naturally benefit from using medial balls.

1.4 Distance diagrams

§ **Voronoi diagram** Consider a collection of points \mathcal{V} . The Voronoi diagram of \mathcal{V} is a partition of the space based on the distance to points in \mathcal{V} . Indeed, it encodes for each $u \in \mathcal{V}$ which points in space are closest to u than other points in $\mathcal{V} \setminus \{u\}$.

Definition 1 (Voronoi cell). Let $\mathcal{V} \subseteq \mathbb{R}^d$ be a collection of points and $u \in \mathcal{V}$. We call *Voronoi cell* of u to \mathcal{V} and denote $\text{vor}(u, \mathcal{V})$ the subset of points that admit u as a closest point in \mathcal{V} ,

$$\text{vor}(u, \mathcal{V}) = \left\{ x \in \mathbb{R}^d \mid \|x - u\| \leq \|x - v\|, \forall v \in \mathcal{V} \right\}$$

The *Voronoi diagram* of \mathcal{V} is the subdivision of \mathbb{R}^d induced by the Voronoi cells of points in \mathcal{V} .

Voronoi diagrams have a wealth of properties that are beyond the scope of this thesis. They are in particular closely related to Delaunay triangulations, that we will not address here. Note that Voronoi cells can be interpreted as the intersection of closed half-spaces, and it follows that they are closed and convex. Finally, note that Voronoi diagrams can be computed in polynomial time for any finite collection \mathcal{V} . For a more in-depth discussion on Voronoi diagrams, see for instance [BY95].

§ **Power diagram** Similarly to Voronoi diagrams, other diagrams can be built using different notions of distance. In particular, we can consider collections of balls instead of points. An interesting replacement for the notion of distance between two points is then the power of a point to a ball.

Definition 2 (Power cell). Let \mathcal{B} be a collection of balls. Consider a ball $b \in \mathcal{B}$. We call *power cell* of b to \mathcal{B} and denote $\text{pcell}(b, \mathcal{B})$ the collection of points

$$\text{pcell}(b, \mathcal{B}) = \left\{ x \in \mathbb{R}^d \mid \forall b' \in \mathcal{B}, \text{pow}(x, b) \leq \text{pow}(x, b') \right\}$$

The *power diagram* of \mathcal{B} is the subdivision of \mathbb{R}^d induced by the power cells of balls in \mathcal{B} .

Remark 2. Similarly to Voronoi cells, power cells are closed and convex.

1.5 A result about set inclusion

We prove here a very general result to characterize the inclusion of one set into another one. We assume very little on the two sets, only that they are subsets of a connected topological space. We show that it suffices to test “local” inclusions at the boundary points of the two sets.

Lemma 1.1. *Let A and $X \neq \emptyset$ be two subsets of a connected topological space. If*

$$(i) \quad \partial A \subseteq \overline{X}$$

$$(ii) \quad \forall x \in \partial X, \exists N_x \text{ an open neighbourhood of } x \text{ such that } N_x \cap A \subseteq X$$

then $A \subseteq X$.

Proof. Letting $H = A \cap X^c$, we prove the lemma in two stages:

(a) First, we prove that $\partial H \subseteq \partial X$. Using $\partial(A \cap B) \subseteq (\partial A \cap \overline{B}) \cup (\overline{A} \cap \partial B)$ (see Lemma A.1) and plugging $B = X^c$, we get $\partial H \subseteq (\partial A \cap \overline{X^c}) \cup (\overline{A} \cap \partial X^c)$. Since $\partial A \subseteq \overline{X}$ and $\partial X^c = \partial X$, we deduce $\partial H \subseteq \partial X$.

(b) Second, we prove that $H = \emptyset$. Suppose for a contradiction that $H \neq \emptyset$. Since X is not empty, X^c is not the entire topological space, so that H cannot be the entire topological space either. Thus, $\partial H \neq \emptyset$ and we can pick a point $x \in \partial H$. Because $\partial H \subseteq \partial X$, we can apply (ii) on point x , showing that there is N_x a neighbourhood of x such that $N_x \cap A \subseteq X$. Since $x \in \partial H$, necessarily $N_x \cap H \neq \emptyset$. However, $N_x \cap H = N_x \cap A \cap X^c \subseteq X \cap X^c = \emptyset$, yielding a contradiction. \square

Note that the inclusion $A \subseteq X$ implies both (i) and (ii), thus this is an equivalence. When working with a specific class of shapes, it can be interesting to look at a specialization of this lemma. In particular, we will see in the later Section 4.3 a specialization of Lemma 1.1 where the containing set X is a finite union of balls.

Finite union of balls

Contents

2.1 Pencil of balls	22
2.2 Convex subsets of pencils	36
2.2.1 Normal cones and power diagram	39
2.2.2 Structure of the boundary	44
2.2.3 Proofs of results	53
2.3 Local inclusion in parabolic pencils	54
2.3.1 Technical lemmas	55
2.3.2 Proving the property of local inclusion cones	57
2.4 Boundary of a finite union of balls	58
2.5 Medial axis, skeleton, boundary and pencils	64
2.5.1 State of the art when assuming general position	65
2.5.2 Link of a polytope	67
2.5.3 From joins to Voronoi cells.	69
2.5.4 . . . to polytopes.	73
2.5.5 . . . to the medial axis	77
2.5.6 Proof of Theorem 3	80
2.6 Conclusion	81

Finite unions of balls are central to the discussion that we conduct throughout this thesis. We thus dedicate this chapter to the study of their structure, and some of their properties. To that effect, we rely on pencils of balls. Pencils are affine families of balls with many properties, and we show that any finite union of balls can be decomposed into finitely many subsets of pencils. By carefully investigating the properties of pencils, we eventually obtain a description of the boundary of finite unions of balls that holds without the general position assumption. This in turn leads to a description of the medial axis in terms of pencils, which also does not require the balls to be in general position. All in all, the use of pencils allows us to rely on a terminology and a toolset that can handle any degeneracy within finite unions of balls.

2.1 Pencil of balls

§ Introduction

Definition 3 (Pencil of balls). Let \mathcal{P} be a k -flat in \mathbb{R}^{d+1} , $0 \leq k \leq d$. \mathcal{P} corresponds to balls in \mathbb{R}^d . Consider $c(\mathcal{P})$ the collection of centers of balls in \mathcal{P} . If $c(\mathcal{P})$ also has dimension k , that is if the direction $\vec{\mathcal{P}}$ of \mathcal{P} does not contain the vertical $(d+1)$ -th direction, we say that \mathcal{P} is a *pencil of balls* of dimension k , or k -pencil, in \mathbb{R}^d . We call $c(\mathcal{P})$ the *centers' space* of \mathcal{P} . A pencil \mathcal{P}' is a *subpencil* of \mathcal{P} if we have $\mathcal{P}' \subseteq \mathcal{P}$.

Remark 3. If we allowed the affine subspace \mathcal{P} to contain the vertical $(d+1)$ -th direction, we would allow collections of concentric balls. Indeed consider any vertical line L of \mathbb{R}^{d+1} , intersecting \mathbb{R}^d in c . L is in fact the collection of every balls of \mathbb{R}^d centered at c .

Remark 4. The definition ensures that for all $c \in c(\mathcal{P})$, there exists a unique $b \in \mathcal{P}$ such that $c(b) = c$. Due to this one-to-one mapping between a ball $b \in \mathcal{P}$ of the pencil, and its center $c \in c(\mathcal{P})$ in the centers' space, it is equivalent to consider a ball or its center, as long as the pencil is properly identified.

Remark 5. The centers' space $c(\mathcal{P})$ is a k -flat in \mathbb{R}^d . We sometimes consider its associated vector subspace, $\vec{c}(\mathcal{P})$.

There are many different ways to define pencils of balls. Although the final objects resulting from these various definitions express the same basic concept, there are many variations on which families of balls may be called a pencil or not. The above definition is similar to Devillers' [DMT92] and Boissonnat and Yvinec's [BY95], with the differences that we do not restrict ourselves to lines in \mathbb{R}^{d+1} , and forbid "concentric pencils" that span the orthogonal direction to the \mathbb{R}^d hyperplane. Accordingly, depending on how pencils are defined, the various terms used to describe and classify them may or may not make sense. We chose to partially adopt the terminology from [Sch79]. As such we refer to pencils that contain imaginary balls as *hyperbolic*, pencils that contain a unique singleton ball as *parabolic*, and *elliptic* otherwise. Recall that the paraboloid separates imaginary balls from balls with positive radius. Hence hyperbolic pencils intersect the convex region delimited by the paraboloid, parabolic pencils are tangent to the paraboloid, and elliptic pencils are strictly below the paraboloid. See Figures 2.1, 2.2 and 2.3 for examples of pencils. In Boissonnat and Yvinec's terms, hyperbolic, parabolic and elliptic pencils are respectively pencils with two limit points, tangent pencils, and pencils with a "supporting sphere". This latter naming convention reflects more closely the geometric configuration of the pencil and the paraboloid in \mathbb{R}^{d+1} . Indeed in the context of [BY95], a pencil is a line of \mathbb{R}^{d+1} , and it intersects the paraboloid \mathcal{P} in 0, 1, or 2 points. When the (line) pencil intersects \mathcal{P} in 2 points, it is then called in [BY95] a "two limit points" pencil. As for the "supporting sphere" label, it originates from the property that an elliptic pencil (of dimension at most $d-1$) defines a sphere with positive radius. We shall see how to precisely define this sphere in a broader context. This particular sphere will be called the *radical sphere* of the pencil. Finally, we will later define and focus on pencils that satisfy the property of having a "real radical sphere". In particular we prove that these pencils are either elliptic with dimension $k \leq d-1$, or parabolic (of any dimension). We now address various properties of pencils.




Figure 2.1 – The three types of k -pencils in \mathbb{R}^d . Illustrated for $k = 1$ and $d = 1$.

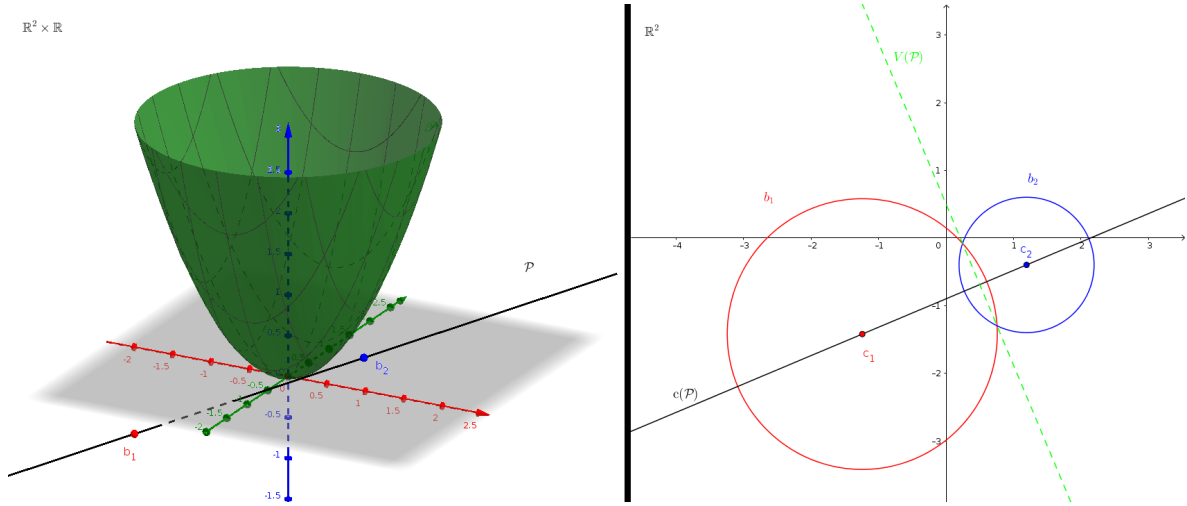


Figure 2.2 – Example of elliptic 1-pencil in \mathbb{R}^2 .

§ **Affine combinations of balls** Consider n balls $b_1, \dots, b_n \in \mathbb{R}^{d+1}$, and n coefficients $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ such that $\sum_{i=1}^n \lambda_i = 1$. We denote by $\sum_{i=1}^n \lambda_i b_i$ the affine combination of the b_i 's with coefficients λ_i 's, that is the ball which corresponds to the $(d + 1)$ -dimensional point

$$\left(\sum_{i=1}^n \lambda_i c_i, \sum_{i=1}^n \lambda_i \text{pow}(O, b_i) \right)$$

which is the affine combination of points $b_i = (c_i, \text{pow}(O, b_i))$. We say that a collection of balls \mathcal{B} spans a pencil \mathcal{P} if $\text{aff}(\mathcal{B}) = \mathcal{P}$.


Property 2.1. $\forall x \in \mathbb{R}^d, \text{pow}(x, \sum_{i=1}^n \lambda_i b_i) = \sum_{i=1}^n \lambda_i \text{pow}(x, b_i)$

Proof. Consider $x \in \mathbb{R}^d$ and let $b = b(c, r)$ be any d -ball. Then we have

$$\begin{aligned} \text{pow}(x, b) &= \|x - c\|^2 - r^2 = \|x\|^2 - 2\langle x, c \rangle + \|c\|^2 - r^2 \\ &= \|x\|^2 - 2\langle x, c \rangle + \text{pow}(O, b) \end{aligned} \tag{2.1}$$

In particular, $\text{pow}(x, b)$ is linear in $c(b)$ and $\text{pow}(O, b)$. If $b = \sum_{i=1}^n \lambda_i b_i$, then $c(b) = \sum_{i=1}^n \lambda_i c_i$ and $\text{pow}(O, b) = \sum_{i=1}^n \lambda_i \text{pow}(O, b_i)$. Thus,

$$\text{pow}\left(x, \sum_{i=1}^n \lambda_i b_i\right) = \sum_{i=1}^n \lambda_i \left(\|x\|^2 - 2\langle x, c_i \rangle + \text{pow}(O, b_i) \right) = \sum_{i=1}^n \lambda_i \text{pow}(x, b_i) \quad \square$$

Figure 2.3 – Example of elliptic 2-pencil in \mathbb{R}^3 .

Given a (possibly infinite) collection of balls \mathcal{B} , its affine hull in \mathbb{R}^{d+1} , $\text{aff}(\mathcal{B})$, may not necessarily define a pencil, since it may span the $(d + 1)$ -th direction. Note however that if $\text{aff}(\mathcal{B})$ is indeed a pencil, then it is the smallest pencil that contains \mathcal{B} . In fact, \mathcal{B} is the subset of a pencil if its affine hull is a pencil. To see this, assume there exists a pencil that contains \mathcal{B} , then it must contain $\text{aff}(\mathcal{B})$. Thus the affine hull $\text{aff}(\mathcal{B})$ is also a pencil. For easy reference, we summarize these observations in the remark below.

Remark 6. There is a pencil that contains \mathcal{B} if and only if $\text{aff}(\mathcal{B})$ is a pencil. Then, $\text{aff}(\mathcal{B})$ is the smallest pencil that contains \mathcal{B} .

We now give a sufficient condition on \mathcal{B} which ensures that $\text{aff}(\mathcal{B})$ is a pencil.

Property 2.2. Consider a (possibly infinite) collection of balls \mathcal{B} which satisfies $\bigcap_{b \in \mathcal{B}} \partial b \neq \emptyset$. Then $\text{aff}(\mathcal{B})$ is a pencil of balls.

Proof. Without loss of generality, we can assume that $\mathcal{B} = \{b_1, \dots, b_n\}$ is a finite collection of balls. Indeed, it suffices to select a maximum family of affinely independent balls in \mathcal{B} . Such a family still has non-empty intersection of the bounding $(d - 1)$ -spheres, and its affine hull coincide with that of the complete collection \mathcal{B} . Denote by $\mathcal{P} = \text{aff}(\mathcal{B})$. By definition \mathcal{P} is an affine space in \mathbb{R}^{d+1} , thus it is a pencil if and only if \mathcal{P} does not span the vertical, $(d + 1)$ -th direction. Suppose for a contradiction that \mathcal{P} spans the vertical direction. Then for all $b \in \mathcal{P}$, there is $b' \in \mathcal{P} \setminus \{b\}$ such that $c(b) = c(b')$. Because $b \neq b'$, we have $\text{pow}(O, b) \neq \text{pow}(O, b')$. In particular, let $b \in \mathcal{P}$ such that $c(b) = c(b_1)$ and $\text{pow}(O, b) \neq \text{pow}(O, b_1)$. Because $b \in \mathcal{P}$, there are $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ such that $b = \sum_{i=1}^n \lambda_i b_i$ and $\sum_{i=1}^n \lambda_i = 1$. Let $x \in \bigcap_{i=1}^n \partial b_i$. By Property 2.1, we deduce $\text{pow}(x, b) = \sum_{i=1}^n \lambda_i \text{pow}(x, b_i) = 0$. However by Equation (2.1), we also have $\text{pow}(x, b) = \|x\|^2 - 2\langle x, c(b) \rangle + \text{pow}(O, b)$. Because $\text{pow}(x, b) = 0 = \text{pow}(x, b_1)$, and $c(b) = c(b_1)$, we thus deduce that $\text{pow}(O, b) = \text{pow}(O, b_1)$. This contradicts the assumption $\text{pow}(O, b) \neq \text{pow}(O, b_1)$ and is impossible. Hence \mathcal{P} cannot span the vertical direction, and is thus a pencil. \square

In other words, as long as the bounding $(d - 1)$ -spheres of balls of \mathcal{B} have non-empty

intersection, \mathcal{B} defines a pencil through $\text{aff}(\mathcal{B})$.

§ **Defining radical objects** We introduce here a few objects which will turn useful to characterize pencils.

Definition 4 (Radical space). Given a (possibly infinite) collection of balls \mathcal{B} , we call *radical space* of \mathcal{B} and denote $V(\mathcal{B})$ the collection of points

$$V(\mathcal{B}) = \left\{ x \in \mathbb{R}^d \mid \forall b, b' \in \mathcal{B}, \text{pow}(x, b) = \text{pow}(x, b') \right\}$$

We establish hereafter that the radical space of a pencil \mathcal{P} is an affine space that relates to its centers' space $c(\mathcal{P})$.

Theorem 1. If \mathcal{B} is a finite collection of balls with $V(\mathcal{B}) \neq \emptyset$, then $V(\mathcal{B})$ is an affine subspace of \mathbb{R}^d , and

$$\mathbb{R}^d = \overrightarrow{\text{aff}}(c(\mathcal{B})) \oplus \overrightarrow{V}(\mathcal{B})$$

Remark 7. We observed in Remark 5 that $c(\mathcal{P})$ is a k -flat. Note however that because we assume \mathcal{B} to be finite, $c(\mathcal{B})$ is not a k -flat, and thus “ $\overrightarrow{c}(\mathcal{B})$ ” is ill-defined.

Proof. Let $\mathcal{B} = \{b_0, b_1, \dots, b_n\}$, with $b_i = b(c_i, r_i)$. Notice that by definition of $V(\mathcal{B})$, we have $V(\mathcal{B}) = \bigcap_{0 \leq i, j \leq n} H(b_i, b_j)$. That is, $V(\mathcal{B})$ is the intersection of the bisector of every pair of balls in \mathcal{B} . Thus $V(\mathcal{B})$ is a finite intersection of hyperplanes, it is either empty or an affine subspace.

We now prove by double inclusion that if $\overrightarrow{V}(\mathcal{B}) \neq \emptyset$, then $\overrightarrow{V}(\mathcal{B}) = \left(\overrightarrow{\text{aff}}(c(\mathcal{B})) \right)^\perp$. Consider $\mathbf{e} \in \overrightarrow{V}(\mathcal{B})$ and $c \in \text{aff}(c(\mathcal{B}))$. By definition of $\text{aff}(c(\mathcal{B}))$, there exist $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ such that $c - c_0 = \sum_{i=1}^n \lambda_i (c_i - c_0)$. Recall that $V(\mathcal{B}) = \bigcap_{i,j} H(b_i, b_j)$. In particular, for all i , $1 \leq i \leq n$, we have $\langle \mathbf{e}, c_i - c_0 \rangle = 0$. By linearity $\langle \mathbf{e}, c - c_0 \rangle = 0$ and we deduce $\mathbf{e} \in \left(\overrightarrow{\text{aff}}(c(\mathcal{B})) \right)^\perp$, thus $\overrightarrow{V}(\mathcal{B}) \subseteq \left(\overrightarrow{\text{aff}}(c(\mathcal{B})) \right)^\perp$.

Conversely consider $\mathbf{e} \in \left(\overrightarrow{\text{aff}}(c(\mathcal{B})) \right)^\perp$ and $v \in V(\mathcal{B})$. For all $i \neq j$ we have $\langle \mathbf{e}, c_i - c_j \rangle = 0$, thus $v + \mathbf{e} \in H(b_i, b_j)$. We have $v + \mathbf{e} \in V(\mathcal{B})$, hence $\mathbf{e} \in \overrightarrow{V}(\mathcal{B})$, and $\left(\overrightarrow{\text{aff}}(c(\mathcal{B})) \right)^\perp \subseteq \overrightarrow{V}(\mathcal{B})$. It follows that we indeed have $\overrightarrow{\text{aff}}(c(\mathcal{B})) \oplus \overrightarrow{V}(\mathcal{B}) = \mathbb{R}^d$. \square

Property 2.3. If \mathcal{P} is a pencil with dimension k , then $V(\mathcal{P})$ is a $(d - k)$ -flat.

Proof. Pencil \mathcal{P} has dimension k , thus there exists $k + 1$ balls $\mathcal{B} = \{b_0, b_1, \dots, b_k\}$ such that $\mathcal{P} = \text{aff}(\mathcal{B})$. Because $\mathcal{B} \subseteq \mathcal{P}$ we immediately deduce $V(\mathcal{P}) \subseteq V(\mathcal{B})$. We argue that the reverse inclusion also holds, which implies $V(\mathcal{P}) = V(\mathcal{B})$. Because $\text{aff}(c(\mathcal{B})) = \text{aff}(\mathcal{P})$ is a k -flat, Theorem 1 then yields the result.

Let $x \in V(\mathcal{B})$ and $b \in \mathcal{P}$. There are $\lambda_0, \dots, \lambda_k$ such that $b = \sum_{i=0}^k \lambda_i b_i$ and $\sum_{i=0}^k \lambda_i = 1$. By definition of $V(\mathcal{B})$, for all i we have $\text{pow}(x, b_i) = \text{pow}(x, b_0)$. Using Property 2.1 we obtain

$$\text{pow}(x, b) = \sum_{i=0}^k \lambda_i \text{pow}(x, b_i) = \left(\sum_{i=0}^k \lambda_i \right) \text{pow}(x, b_0) = \text{pow}(x, b_0)$$

Because this holds for every $b \in \mathcal{P}$, we deduce that $x \in V(\mathcal{P})$. Therefore, $V(\mathcal{P}) = V(\mathcal{B})$ which concludes the proof. \square

In the proof above, we in fact show that Theorem 1 extends to pencils of balls. Recall that as per Remark 5, the centers' space $c(\mathcal{P})$ of pencil \mathcal{P} is an affine space, thus $\vec{c}(\mathcal{P})$ is well defined. In addition, Property 2.3 implies that $V(\mathcal{P})$ is always non-empty. Theorem 1 restricted to pencils thus simplifies to the following:

Corollary 2.4. *If \mathcal{P} is a pencil, then $\mathbb{R}^d = \vec{c}(\mathcal{P}) \oplus \vec{V}(\mathcal{P})$.*

We also have the corollary below, that will be useful later on.

Corollary 2.5. *Let \mathcal{P} be a pencil and \mathcal{B} a finite collection of balls. If $\mathcal{P} = \text{aff}(\mathcal{B})$, then $V(\mathcal{P}) = V(\mathcal{B})$.*

Definition 5 (Radical ball). Let \mathcal{P} be a pencil. We call *radical center* of \mathcal{P} , and denote by $c_0(\mathcal{P})$, the unique intersection point between its radical space $V(\mathcal{P})$ and its centers' space $c(\mathcal{P})$. We call *radical ball* of \mathcal{P} , and denote by $b_0(\mathcal{P})$, the unique ball of \mathcal{P} centered at $c_0(\mathcal{P})$. Additionally, we let $-r_0^2(\mathcal{P}) = \text{pow}(c_0(\mathcal{P}), b_0(\mathcal{P}))$. The value $r_0(\mathcal{P})$ is the (potentially imaginary) radius of $b_0(\mathcal{P})$.

Remark 8. Because $c_0(\mathcal{P}) \in V(\mathcal{P})$, we have $-r_0^2(\mathcal{P}) = \text{pow}(c_0(\mathcal{P}), b)$, for all $b \in \mathcal{P}$.

Remark 9. We always have $r_0^2(\mathcal{P}) \in \mathbb{R}$. However $r_0(\mathcal{P})$ is either non-negative, or purely imaginary.

A key property of the radical ball is as follows.

Property 2.6. *Let $b \in \mathcal{P}$ be centered at c with radius r . Then*

$$r^2 = r_0^2(\mathcal{P}) + \|c - c_0(\mathcal{P})\|^2$$

Proof. This is a direct consequence of Remark 8. \square

Remark 10. Notice that $r^2 \geq r_0^2(\mathcal{P})$, with equality if and only if $c = c_0(\mathcal{P})$. Because r^2 is the signed vertical distance between the $(d+1)$ -dimensional point b and the (interior of the) paraboloid, the radical ball $b_0(\mathcal{P})$ is the ball with smallest signed vertical distance.

Remark 11. For non-hyperbolic pencils, $b_0(\mathcal{P})$ is the ball with smallest radius.

Corollary 2.7. *Let \mathcal{P} be a pencil. We have the three equivalences:*

- \mathcal{P} is hyperbolic if and only if $r_0^2(\mathcal{P}) < 0$
- \mathcal{P} is parabolic if and only if $r_0^2(\mathcal{P}) = 0$
- \mathcal{P} is elliptic if and only if $r_0^2(\mathcal{P}) > 0$

Proof. Direct consequence of Remark 10. □

Thanks to Property 2.6, the power of a point to a ball of the pencil can be expressed in terms of the radical ball.

Lemma 2.8. *Let \mathcal{P} be a pencil. Denote by $b_0 = b_0(\mathcal{P})$ and $c_0 = c_0(\mathcal{P})$. Consider $b \in \mathcal{P}$ centered at c . For all $x \in \mathbb{R}^d$, we have*

$$\text{pow}(x, b) = \text{pow}(x, b_0) - 2\langle x - c_0, c - c_0 \rangle$$

Proof. By Property 2.6 $r^2 = r_0^2 + \|c - c_0\|^2$. Thus

$$\begin{aligned} \text{pow}(x, b) &= \|x - c\|^2 - r^2 \\ &= \|x - c_0\|^2 - 2\langle x - c_0, c - c_0 \rangle + \|c - c_0\|^2 - r_0^2 - \|c - c_0\|^2 \\ &= \text{pow}(x, b_0) - 2\langle x - c_0, c - c_0 \rangle \end{aligned} \quad \square$$

Definition 6 (Radical sphere). Let \mathcal{P} be a pencil. We call *radical sphere* of \mathcal{P} and denote by $s_0(\mathcal{P})$ the subset $s_0(\mathcal{P}) = V(\mathcal{P}) \cap \partial b_0(\mathcal{P})$.

The radical sphere will frequently turn up in future discussions, especially when it is non-empty. As such, we explicit now the various forms the radical sphere can take.

Property 2.9. *Let \mathcal{P} be a pencil of \mathbb{R}^d .*

- *If \mathcal{P} is elliptic with dimension d , or hyperbolic, then $s_0(\mathcal{P}) = \emptyset$.*
- *If \mathcal{P} is parabolic, then $s_0(\mathcal{P}) = \{c_0(\mathcal{P})\}$ is a singleton.*
- *Otherwise \mathcal{P} is elliptic with dimension at most $d - 1$, and $s_0(\mathcal{P})$ is a proper $(d - 1 - k)$ -sphere, where $k = \dim \mathcal{P}$. By definition, the center, radius, and support space of the radical sphere are $c_0(\mathcal{P})$, $r_0(\mathcal{P})$, and $V(\mathcal{P})$.*

Proof. We proceed in order.

Recall that by convention, we consider imaginary balls to be empty. Thus if $b_0(\mathcal{P})$ is imaginary, or equivalently if \mathcal{P} is hyperbolic, then $s_0(\mathcal{P})$ is empty. Additionally, when $\dim \mathcal{P} = d$, we know by Corollary 2.4 that $\dim V(\mathcal{P}) = 0$, thus $V(\mathcal{P}) = \{c_0(\mathcal{P})\}$. As per Corollary 2.7,

pencil \mathcal{P} is elliptic when $r_0(\mathcal{P}) > 0$, which entails $c_0(\mathcal{P}) \notin \partial b_0(\mathcal{P})$. It follows that for elliptic d -pencils, $s_0(\mathcal{P})$ is also empty.

Assume now that \mathcal{P} is parabolic. According to Corollary 2.7 we have $r_0(\mathcal{P}) = 0$ and thus $b_0(\mathcal{P}) = \{c_0(\mathcal{P})\}$. By definition we always have $c_0(\mathcal{P}) \in V(\mathcal{P})$, thus $s_0(\mathcal{P})$ indeed reduces to the singleton radical center.

In the remaining cases, \mathcal{P} is an elliptic k -pencil, $k \leq d - 1$. As observed previously, $b_0(\mathcal{P})$ has positive radius. In addition, Corollary 2.4 yields $\dim V(\mathcal{P}) = d - k$. Therefore $s_0(\mathcal{P})$ is a $(d - 1 - k)$ -sphere with the specified parameters. \square

From the properties of the radical objects, we also easily obtain the property below.

Property 2.10. *Let \mathcal{P} be a pencil, then*

$$\bigcap_{b \in \mathcal{P}} \partial b = s_0(\mathcal{P})$$

Proof.

$$\begin{aligned} x \in s_0(\mathcal{P}) &\iff x \in V(\mathcal{P}) \cap \partial b_0(\mathcal{P}) \\ &\iff \forall b \in \mathcal{P}, \text{pow}(x, b) = \text{pow}(x, b_0(\mathcal{P})) = 0 \\ &\iff \forall b \in \mathcal{P}, x \in \partial b \\ &\iff x \in \bigcap_{b \in \mathcal{P}} \partial b \end{aligned} \quad \square$$

Although we put emphasis on the radical ball $b_0(\mathcal{P})$ while defining the radical sphere $s_0(\mathcal{P})$, the properties of the radical space $V(\mathcal{P})$ guarantee that the radical sphere would remain unchanged by choosing any other sphere of the pencil \mathcal{P} . In fact, a minor adaptation of the proof above immediately yields the corollary below.

Corollary 2.11. $\forall b \in \mathcal{P}, s_0(\mathcal{P}) = V(\mathcal{P}) \cap \partial b$.

As a side effect, we can now state a pencil-focused version of a well-known result: the intersection of finitely many $(d - 1)$ -spheres is also a sphere.

Theorem 2 (Finite intersection of $(d - 1)$ -spheres). Let \mathcal{B} be a finite collection of balls with $\bigcap_{b \in \mathcal{B}} \partial b \neq \emptyset$. Then

$$\bigcap_{b \in \mathcal{B}} \partial b = s_0(\mathcal{P})$$

where \mathcal{P} is the smallest pencil that contains \mathcal{B} .

Remark 12. Note that by Property 2.2 and Remark 6, we have $\mathcal{P} = \text{aff}(\mathcal{B})$.

Proof. By Remark 12, we have $\mathcal{P} = \text{aff}(\mathcal{B})$. We can thus apply Corollary 2.5 to obtain $V(\mathcal{B}) = V(\mathcal{P})$. By definition, points of $V(\mathcal{B})$ have equal power from all balls of \mathcal{B} . We can thus easily generalize Remark 1 to finitely many balls, and deduce that for all $b \in \mathcal{B}$, we can express the intersection of bounding spheres of \mathcal{B} as $\bigcap_{b' \in \mathcal{B}} \partial b' = V(\mathcal{B}) \cap \partial b = V(\mathcal{P}) \cap \partial b$. Because $b \in \mathcal{P}$, we can apply Corollary 2.11 and thus finally obtain $\bigcap_{b \in \mathcal{B}} \partial b = s_0(\mathcal{P})$. \square

Remark 13. Because the radical sphere is either empty, or a (possibly degenerate) sphere, this theorem indeed implies that the intersection of $(d - 1)$ -spheres is also a sphere.

Remark 14. With Property 2.9, we can easily obtain the well-known general position version of Theorem 2: given k balls in general position, $k \leq d$, the intersection of their bounding $(d - 1)$ -spheres is either empty, or a $(d - k)$ -sphere. Indeed, assuming general position, \mathcal{P} is then a non-parabolic pencil of dimension $k - 1 \leq d - 1$. If \mathcal{P} is hyperbolic, $s_0(\mathcal{P})$ is empty. Otherwise, \mathcal{P} is elliptic and $s_0(\mathcal{P})$ is a $(d - k)$ -sphere.

§ Intersection and union over a pencil We present here some results on pencils that give some insight on the way pencils cover the space.

Lemma 2.12. *Consider a pencil \mathcal{P} . Denote its radical space by V and its radical ball by b_0 . Then*

$$\forall b \in \mathcal{P}, b \cap V = b_0 \cap V \quad (2.2)$$

$$\bigcap_{b \in \mathcal{P}} b \setminus V = \emptyset \quad (2.3)$$

$$\bigcup_{b \in \mathcal{P}} \overset{\circ}{b} \setminus V = \mathbb{R}^d \setminus V \quad (2.4)$$

$$\bigcup_{b \in \mathcal{P}} \partial b \setminus V = \mathbb{R}^d \setminus V \quad (2.5)$$

Proof. (2.2) stems from the definition of the radical space V . Indeed, the radical space is the subset of points of \mathbb{R}^d whose power is equal from all balls of the pencil \mathcal{P} , $V = V(\mathcal{P}) = \{x \in \mathbb{R}^d \mid \forall b, b' \in \mathcal{P}, \text{pow}(x, b) = \text{pow}(x, b')\}$. Because the sign of the power of x to a ball fully characterizes the membership of x to that ball, it follows that within the radical space V , all the balls of the pencil share the same points.

For (2.3), (2.4), and (2.5), consider $x \in \mathbb{R}^d \setminus V$. We prove that there is always some ball in \mathcal{P} that contains x in its interior, some in its boundary, and some that does not contain x . Recall by Corollary 2.4 that $\vec{\mathcal{C}}(\mathcal{P}) \oplus \vec{V} = \mathbb{R}^d$. Let $\mathbf{e}_c \in \vec{\mathcal{C}}(\mathcal{P})$ and $\mathbf{e}_V \in \vec{V}$ such that $x - c_0 = \mathbf{e}_c + \mathbf{e}_V$, where c_0 denotes the radical center of \mathcal{P} . For $\lambda \in \mathbb{R}$ let $c_\lambda = c_0 + \lambda \mathbf{e}_c$ and denote by $b_\lambda \in \mathcal{P}$ the ball centered at c_λ . By Lemma 2.8 we have

$$\begin{aligned} \text{pow}(x, b_\lambda) &= \text{pow}(x, b_0) - 2\langle x - c_0, c_\lambda - c_0 \rangle \\ &= \text{pow}(x, b_0) - 2\langle \mathbf{e}_c, \lambda \mathbf{e}_c \rangle \\ &= \text{pow}(x, b_0) - 2\lambda \|\mathbf{e}_c\|^2 \\ &= A\lambda + B \end{aligned}$$

where $A, B \in \mathbb{R}$ are constant with respect to λ . Because $x \notin V$ we have $\|\mathbf{e}_c\| \neq 0$ and therefore $A \neq 0$. It follows that the power of point x with respect to balls of the pencil spans \mathbb{R} , which suffices to conclude. \square

From the proof, it is quite clear that for any point $x \in \mathbb{R}^d \setminus V$, there is an infinity of balls in \mathcal{P} that contain x , and also an infinity that do not contain x . Also if the pencil has dimension $k \geq 2$, consider $\mathbf{e} \in \vec{\mathcal{C}}(\mathcal{P})$ such that $\langle \mathbf{e}_c, \mathbf{e} \rangle = 0$. Let $c_{\lambda, \mu} = c_0 + \lambda \mathbf{e}_c + \mu \mathbf{e}$ and $b_{\lambda, \mu}$ be the corresponding ball of the pencil. For all $x \in \mathbb{R}^d \setminus V$ and $\mu \in \mathbb{R}$ we have $\text{pow}(x, b_{\lambda, \mu}) = \text{pow}(x, b_{\lambda=0, \mu})$. Hence there is also an infinity of balls that contain x in their boundary. In fact, the collection of balls $\{b_{\lambda=0, \mu} \mid \mu \in \mathbb{R}\}$ forms a subpencil of dimension 1 of \mathcal{P} .

§ Characterizing a pencil through its radical elements We have seen that we can associate to a pencil \mathcal{P} an affine subspace $V(\mathcal{P})$ and a radical ball $b_0(\mathcal{P})$. We show here that the converse is true. Given a d -ball $b_0 \subseteq \mathbb{R}^d$ and a k -flat V passing through the center of b_0 , there exists a unique pencil \mathcal{P} , such that $V(\mathcal{P}) = V$ and $b_0(\mathcal{P}) = b_0$.

For reasons that will be explained later on, we prefer to state our result using triplets (V, c_0, r_0^2) , rather than directly using (V, b_0) , where $b_0 = b(c_0, r_0)$. Recall from (1.1) that a ball $b = b(c, r)$ in \mathbb{R}^d can be represented by a point of coordinates $(c, \|c\|^2 - r^2)$ in \mathbb{R}^{d+1} .

Property 2.13. *Let $V \subseteq \mathbb{R}^d$ be a k -flat, $c_0 \in V$ a point, and $r_0^2 \in \mathbb{R}$ a real number. There exists a unique pencil \mathcal{P} such that $V(\mathcal{P}) = V$ and $b_0(\mathcal{P}) = (c_0, \|c_0\|^2 - r_0^2) \in \mathbb{R}^{d+1}$.*

Because radical balls can in general be imaginary, we choose here to express balls as points in \mathbb{R}^{d+1} rather than subsets of \mathbb{R}^d . Indeed in the latter case, all imaginary balls would equate to the empty set, and they would thus be undistinguishable from one another. As points of \mathbb{R}^{d+1} however, two balls are identical if and only if they have the same center and the same power to the origin. Equivalently, two balls are equal if they have the same center and the same (possibly imaginary) radius. To highlight this framework, we thus characterize pencils with the triplet (V, c_0, r_0^2) rather than (V, b_0) , where $b_0 = b(c_0, r_0)$.

We provide a constructive proof of Property 2.13. For any triplet (V, c_0, r_0^2) satisfying the assumptions of this property, we let $\mathcal{P}(V, c_0, r_0^2)$ be the collection of balls

$$\mathcal{P}(V, c_0, r_0^2) = \left\{ (c, \|c\|^2 - r^2) \in \mathbb{R}^{d+1} \mid c \in \left(c_0 + \vec{V}^\perp \right), r^2 = r_0^2 + \|c - c_0\|^2 \right\}$$

We show that $\mathcal{P}(V, c_0, r_0^2)$ is the unique pencil with radical space V and radical ball $b_0 = b(c_0, r_0)$.

For convenience, denote by $\mathcal{C} = c_0 + \vec{V}^\perp$. The collection of centers of $\mathcal{P}(V, c_0, r_0^2)$ is precisely $\mathcal{C} = c(\mathcal{P}(V, c_0, r_0^2))$. By definition of \mathcal{C} , $\vec{\mathcal{C}} = \vec{V}^\perp$ and we have $\vec{\mathcal{C}} \oplus \vec{V} = \mathbb{R}^d$. Thus, if $\mathcal{P}(V, c_0, r_0^2)$ is a pencil, its radical space is indeed V . Likewise, its radical center would be $c_0 \in \mathcal{C} \cap V$, and its radical ball would be $b_0 = (c_0, \|c_0\|^2 - r_0^2) \in \mathbb{R}^{d+1}$. Therefore, it suffices to prove that $\mathcal{P}(V, c_0, r_0^2)$ is a pencil to justify existence in Property 2.13.

Lemma 2.14. $\mathcal{P}(V, c_0, r_0^2)$ is a pencil.

Proof. First, V is a k -flat, thus \mathcal{C} is a $(d - k)$ -flat. Let $c_1, \dots, c_{d-k+1} \in \mathcal{C}$ be $d - k + 1$ affinely independent points, and denote by b_i the (unique) ball of $\mathcal{P}(V, c_0, r_0^2)$ centered at c_i , $1 \leq i \leq d - k + 1$. We argue that $\mathcal{P}(V, c_0, r_0^2) = \text{aff}(b_i \mid 1 \leq i \leq d - k + 1)$, and hence that $\mathcal{P}(V, c_0, r_0^2)$ is an affine subspace of \mathbb{R}^{d+1} . Next, by definition, $\mathcal{P}(V, c_0, r_0^2)$ admits one unique ball per point in \mathcal{C} . Therefore, it cannot span the vertical $(d + 1)$ -th direction of \mathbb{R}^{d+1} and is thus a pencil. Remains to prove that the b_i 's indeed generate $\mathcal{P}(V, c_0, r_0^2)$. We proceed by double inclusion.

Because $b_i \in \mathcal{P}(V, c_0, r_0^2)$ for all i , we trivially have $\text{aff}(b_i \mid 1 \leq i \leq d - k + 1) \subseteq \mathcal{P}(V, c_0, r_0^2)$. Conversely, consider $b \in \mathcal{P}(V, c_0, r_0^2)$ and denote its center by $c = c(b)$. Because $c \in \mathcal{C} = \text{aff}(c_i \mid 1 \leq i \leq d - k + 1)$, let $\lambda_1, \dots, \lambda_{d-k+1} \in \mathbb{R}$ such that $c = \sum_{i=1}^{d-k+1} \lambda_i c_i$ and $\sum_{i=1}^{d-k+1} \lambda_i = 1$. We show that $b = \sum_{i=1}^{d-k+1} \lambda_i b_i$. Because c is the affine combination of the c_i 's by definition, it suffices to show that $\text{pow}(O, b) = \sum_{i=1}^{d-k+1} \lambda_i \text{pow}(O, b_i)$. By definition of $\mathcal{P}(V, c_0, r_0^2)$, we have

$$\begin{aligned} \text{pow}(O, b) &= \|c\|^2 - r_0^2 - \|c - c_0\|^2 = -r_0^2 - \|c_0\|^2 + 2\langle c, c_0 \rangle \\ &= \sum_{i=1}^{d-k+1} \lambda_i \left(-r_0^2 - \|c_0\|^2 + 2\langle c_i, c_0 \rangle \right) \\ &= \sum_{i=1}^{d-k+1} \lambda_i \text{pow}(O, b_i) \end{aligned}$$

Thus, $b = \sum_{i=1}^{d-k+1} \lambda_i b_i$ and $\mathcal{P}(V, c_0, r_0^2) \subseteq \text{aff}(b_i \mid 1 \leq i \leq d - k + 1)$, which concludes the proof. \square

$\mathcal{P}(V, c_0, r_0^2)$ is a pencil with the adequate radical space and radical sphere. Remains to show that it is the unique such pencil.

Lemma 2.15. $\mathcal{P}(V, c_0, r_0^2)$ is the unique pencil with radical space V and radical ball $(c_0, \|c_0\|^2 - r_0^2)$.

Proof. Let \mathcal{P} be a pencil such that $V(\mathcal{P}) = V$ and $b_0(\mathcal{P}) = (c_0, \|c_0\|^2 - r_0^2) \in \mathbb{R}^{d+1}$. We argue that \mathcal{P} must be a subpencil of $\mathcal{P}(V, c_0, r_0^2)$. Assuming this is true, we can conclude. Indeed, by Corollary 2.4 we know that $\dim \mathcal{P} = d - \dim V = \dim \mathcal{P}(V, c_0, r_0^2)$. Thus we deduce that $\mathcal{P} = \mathcal{P}(V, c_0, r_0^2)$.

Remains to show that $\mathcal{P} \subseteq \mathcal{P}(V, c_0, r_0^2)$. Let $b = b(c, r) \in \mathcal{P}$. By Corollary 2.4, we have $c(\mathcal{P}) = c_0 + \vec{V}^\perp$. Therefore $c \in c_0 + \vec{V}^\perp$. In addition, we know by Property 2.6 that $r^2 = r_0^2 + \|c - c_0\|^2$. It follows that $b \in \mathcal{P}(V, c_0, r_0^2)$, and \mathcal{P} is a subpencil of $\mathcal{P}(V, c_0, r_0^2)$. \square

Property 2.13 thus follows as an immediate corollary to Lemma 2.15.

§ **Influence of (V, c_0, r_0^2) over $\mathcal{P}(V, c_0, r_0^2)$** Although we have introduced pencils as specific affine subsets of \mathbb{R}^{d+1} , Property 2.13 guarantees that any triplet (V, c_0, r_0^2) , with V a k -flat, $c_0 \in V$ and $r_0^2 \in \mathbb{R}$, fully determines a pencil. We discuss here how each member of the triplet influences the corresponding pencil. Specifically, we have a look at characteristics of $\mathcal{P} = \mathcal{P}(V, c_0, r_0^2)$ as an affine subspace in \mathbb{R}^{d+1} .

Because we require that $c_0 \in V$, note that those two parameters cannot be independent. In order to decouple their influence over \mathcal{P} , it is handy to consider \vec{V} instead of V . Indeed, we have $V = c_0 + \vec{V}$. Hence, given c_0 and \vec{V} , we can fully characterize V , and additionally c_0 and \vec{V} can be chosen independently. We discuss the influence of r_0^2 , c_0 and \vec{V} in that sequence.

Recall from the definition of $\mathcal{P}(V, c_0, r_0^2)$ that for all $b \in \mathcal{P}(V, c_0, r_0^2)$, we have

$$b = \left(c, \|c\|^2 - \|c - c_0\|^2 - r_0^2 \right)$$

It follows that changes in the value of r_0^2 will only induce a translation of $\mathcal{P}(V, c_0, r_0^2)$ along the vertical $(d+1)$ -th axis.

By Corollary 2.7, we know that when r_0^2 is null, then $\mathcal{P}(V, c_0, r_0^2)$ is parabolic, and that b_0 is the unique singleton ball of the pencil. In that case, $\mathcal{P}(V, c_0, r_0^2)$ is tangent to the paraboloid at $b_0 = (c_0, \|c_0\|^2)$. Obviously, c_0 fully determines the position of the tangency point. Because the paraboloid admits one unique tangent hyperplane $H(c_0)$ at that tangency point, and $\mathcal{P}(V, c_0, r_0^2)$ is an affine subspace, we deduce $\mathcal{P}(V, c_0, r_0^2) \subseteq H(c_0)$. Therefore c_0 determines the admissible direction of the pencil in \mathbb{R}^{d+1} . See Figure 2.4




Figure 2.4 – Schematic representation of the hyperplanes $H(c_0)$ and $H(c_0, r_0^2)$.

So far the discussion was independent from the dimension of $\mathcal{P}(V, c_0, r_0^2)$. Obviously, the information on the dimension is contained within \vec{V} . When $\mathcal{P}(V, c_0, r_0^2)$ is a d -pencil, that is when V has null dimension, then $\mathcal{P}(V, c_0, r_0^2)$ is a hyperplane in \mathbb{R}^{d+1} that coincide with

the tangent hyperplane $H(c_0)$ to the paraboloid at point $(c_0, \|c_0\|^2)$, vertically translated by the value $-r_0^2$. Denote that translated hyperplane by $H(c_0, r_0^2)$. The previous remark states that $\mathcal{P}(\{c_0\}, c_0, r_0^2) = H(c_0, r_0^2)$. Thus, c_0 and r_0^2 conjointly characterize the largest pencil (of dimension d) that admits $b_0 = b(c_0, r_0)$ as its radical ball. As for other, lower dimension pencils, they all are proper subpencils of $H(c_0, r_0^2)$. In fact, \vec{V} specifies which horizontal directions pencil $\mathcal{P}(V, c_0, r_0^2)$ does not span. See Figure 2.5.




Figure 2.5 – The flat in \mathbb{R}^d that goes through c_0 and whose direction is orthogonal to \vec{V} (represented in green) is equal to the vertical projection of $\mathcal{P}(V, c_0, r_0^2)$ in \mathbb{R}^d . Conversely we can obtain $\mathcal{P}(V, c_0, r_0^2)$ (red) by projecting that flat onto $H(c_0, r_0^2)$ (blue).

In summary, c_0 fixes the normal direction of a “maximal” hyperplane that contains $\mathcal{P}(V, c_0, r_0^2)$, r_0^2 specify the position of that hyperplane, while \vec{V} yields the actual direction spanned by the pencil within that hyperplane.

§ Families of balls containing a sphere Given a sphere s in \mathbb{R}^d , one can define the family of every balls that contain that sphere in their boundary:

$$\mathcal{B}(s) = \{b \text{ a } d\text{-ball} \mid s \subseteq \partial b\}$$

We here discuss the link between these families of balls and pencils. Recall that from Property 2.10, every ball in a pencil must contain the radical sphere in its boundary. The discussion will thus revolve around whether any sphere can be the radical sphere of a pencil, and also whether a pencil \mathcal{P} represents the entire family of balls that contains a particular radical sphere, $\mathcal{B}(s_0(\mathcal{P}))$, or if it is strictly smaller. Eventually, we will reach the conclusion below.

Property 2.16. *Let $s_0 \subseteq \mathbb{R}^d$ be a (non-imaginary) sphere.*

- *If s_0 has positive radius and dimension k , then $\mathcal{B}(s_0)$ is an elliptic pencil of dimension $d - k - 1$. It is the unique pencil that admits s_0 as its radical sphere.*
- *If s_0 is a singleton sphere, then $\mathcal{B}(s_0)$ is a parabolic d -pencil. A pencil admits s_0 as its radical sphere if and only if it is a parabolic subpencil of $\mathcal{B}(s_0)$.*

Remark 15. Consider a family of balls \mathcal{B} whose bounding spheres share a common point u . Then, $\text{aff}(\mathcal{B})$ is a subpencil of a parabolic d -pencil with radical sphere $\{u\}$. Indeed by Property 2.2, $\text{aff}(\mathcal{B})$ is a pencil, it may be either parabolic or elliptic. However, due to Property 2.16 above, we know that $\text{aff}(\mathcal{B})$ is a subset of $\mathcal{B}(s = \{u\})$ which is a parabolic d -pencil.

Because every sphere of \mathbb{R}^d can be defined from a triplet (V, c_0, r_0^2) that satisfies the conditions of Property 2.13, it follows that we can associate a pencil \mathcal{P} to any sphere s_0 . Note however that we defined spheres as subsets of \mathbb{R}^d , hence this association is ambiguous whenever the sphere, as a subset, is not uniquely determined by the triplet (V, c_0, r_0^2) . This ambiguity in particular arises for every singleton spheres. Given a singleton sphere $\{c_0\}$, there are indeed multiple affine subspaces V that satisfy $V \cap b(c_0, r_0 = 0) = \{c_0\}$. Regardless, any sphere is always the radical sphere of some pencil.

Proof. Let \mathcal{P} be a pencil and denote its radical sphere by s_0 . Assume that s_0 is non-empty. As observed previously, we have $\mathcal{P} \subseteq \mathcal{B}(s_0)$. To determine whether the reverse inclusion holds or not, we discuss the properties of the collection of centers of balls in $\mathcal{B}(s_0)$, that is $c(\mathcal{B}(s_0))$. First, notice that $\mathcal{B}(s_0)$ does not contain any concentric balls. Indeed the boundaries of any two such balls are disjoint, which would imply that s_0 must be empty, leading to a contradiction. There is thus a one to one mapping between $c(\mathcal{B}(s_0))$ and $\mathcal{B}(s_0)$. Consequently, \mathcal{P} coincides with $\mathcal{B}(s_0)$ if and only if $c(\mathcal{P})$ coincides with $c(\mathcal{B}(s_0))$.

We know that $c(\mathcal{P}) \subseteq c(\mathcal{B}(s_0))$. For the reverse inclusion, consider $b = b(c, r) \in \mathcal{B}(s_0)$. Necessarily, c is in the locus of points that are equidistant from every point in s_0 . It is fairly easy to see that this locus is the intersection of every bisector hyperplane of diameters of s_0 . That intersection of hyperplanes coincides with the affine subspace through the center of s_0 , and whose direction is the orthogonal complement of $\text{aff}(s_0)$. Said otherwise, let c_0 be the center of s_0 . We have $\vec{c}(\mathcal{B}(s_0)) \oplus \vec{\text{aff}}(s_0) = \mathbb{R}^d$, and $c(\mathcal{B}(s_0)) = c_0 + \vec{c}(\mathcal{B}(s_0))$. We easily deduce that the pencil \mathcal{P} coincides with the family $\mathcal{B}(s_0)$, if and only if the affine hull of s_0 coincides with the radical space of \mathcal{P} , that is $\text{aff}(s_0) = V(\mathcal{P})$.

The identity $\text{aff}(s_0) = V(\mathcal{P})$ is always verified for elliptic pencils. This encompasses every non-singleton spheres. Note that if s_0 is a k -sphere with positive radius, then $\text{aff}(s_0)$ is a $(k + 1)$ -flat. By Corollary 2.4 we have $\dim \mathcal{P} = d - \dim V(\mathcal{P}) = d - k - 1$.

For parabolic pencils, we observed previously an ambiguity in specifying a radical space that contains a given radical (singleton) sphere. Nevertheless, parabolic d -pencils coincide with families of balls that contain a singleton sphere, while lower dimension parabolic pencils are strict subsets of these families. It follows that families of balls that go through a sphere are always pencils. \square

§ Pencils from finite intersection of $(d - 1)$ -spheres Recall that we previously gave through Property 2.2 a sufficient condition for the affine hull $\text{aff}(\mathcal{B})$ of a finite collection of balls \mathcal{B} to define a pencil. That condition was for their bounding $(d - 1)$ -spheres to have non-empty intersection, $\bigcap_{b \in \mathcal{B}} \partial b \neq \emptyset$. We dedicate this paragraph to explicitly specify what types of pencils can be obtained through such collections \mathcal{B} . Indeed, every pencils we will come across later on can always be defined through such a collection. We thus give some more details about those pencils.

Property 2.17. *Consider a pencil \mathcal{P} and let \mathcal{B} be a finite collection of balls such that $\mathcal{P} = \text{aff}(\mathcal{B})$. Then, $\bigcap_{b \in \mathcal{B}} \partial b$ is non-empty if and only if the radical sphere $s_0(\mathcal{P})$ is non-empty.*

Remark 16. Recall that by Property 2.9, or also Property 2.16, pencils with non-empty radical sphere are either elliptic of dimension at most $d - 1$, or parabolic.

Proof. If \mathcal{B} satisfies $\bigcap_{b \in \mathcal{B}} \partial b \neq \emptyset$, then Theorem 2 on intersection of $(d - 1)$ -spheres implies that $s_0(\mathcal{P})$ must also be non-empty.

Conversely, suppose that the radical sphere is non-empty. Then any finite collection $\mathcal{B} \subseteq \mathcal{P}$ will satisfy $\bigcap_{b \in \mathcal{B}} \partial b \neq \emptyset$. It suffices to choose a finite \mathcal{B} so that $\text{aff}(\mathcal{B}) = \mathcal{P}$, which can always be done. \square

In other words, a collection of balls \mathcal{B} with $\bigcap_{b \in \mathcal{B}} \partial b \neq \emptyset$ defines a pencil with non-empty radical sphere, and every pencil with non-empty radical sphere can be defined through such a collection \mathcal{B} . Observe that contrary to associating a pencil to a sphere, there is no ambiguity when defining a pencil from a finite collection of balls, even when the intersection of their bounding spheres reduces to a singleton. Indeed, when starting from the singleton sphere, there are many possible choices of radical space for the associated pencil. With a finite collection \mathcal{B} , the pencil must coincide with the affine hull of \mathcal{B} in \mathbb{R}^{d+1} , which fully specifies the pencil. From the point of view of \mathbb{R}^d , the radical space is constrained to be orthogonal to the affine hull of the centers, which resolves the ambiguity presented by singleton spheres.

When dealing with finite unions of balls, collections of balls whose bounding spheres have non-empty intersection naturally arise. Indeed, many subsets of the collection that defines the union will satisfy that property. In Section 2.4 we will rely on the properties of the pencils

defined by those subsets to study the properties of finite unions of balls. We will investigate those pencils throughout Sections 2.2 and 2.3.

2.2 Convex subsets of pencils

§ Convex combinations of balls Similar to affine combinations of balls, we define convex combinations of balls. Consider n balls $b_1, \dots, b_n \in \mathbb{R}^{d+1}$, and n coefficients $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ such that $\sum_{i=1}^n \lambda_i = 1$. Recall that we denote by $\sum_{i=1}^n \lambda_i b_i$ the affine combination of the b_i 's with coefficient λ_i 's, that is the ball which corresponds to the affine combination of the b_i 's when seen as $(d+1)$ -dimensional points. If additionally, the coefficients satisfy the constraint $0 \leq \lambda_i \leq 1$, for all i , then we say that $\sum_{i=1}^n \lambda_i b_i$ is the convex combination of the b_i 's with coefficient λ_i 's.

§ Convex hull of a collection of balls Let \mathcal{B} be a finite collection of d -balls in \mathbb{R}^d . We denote by $\text{CH}(\mathcal{B})$ the set of all convex combinations of balls in \mathcal{B} . If balls are considered as points in \mathbb{R}^{d+1} , $\text{CH}(\mathcal{B})$ is the convex hull of \mathcal{B} .

Remark 17. When working with the convex hull of a finite collection of balls \mathcal{B} , we may assume without loss of generality that \mathcal{B} only contains the vertices of $\text{CH}(\mathcal{B})$.

Clearly, $\text{CH}(\mathcal{B})$ is a subset of $\text{aff}(\mathcal{B})$, the set of affine combinations of balls in \mathcal{B} . A classical result states the following [Ber87]:

Property 2.18. *If $b \in \text{CH}(\mathcal{B})$, then $b \subseteq \bigcup \mathcal{B}$.*

Proof. Let $b_1, \dots, b_n \in \mathcal{B}$ and $\lambda_1, \dots, \lambda_n \in [0, 1]$ such that $b = \sum_{i=1}^n \lambda_i b_i$ and $\sum_{i=1}^n \lambda_i = 1$. By Property 2.1 we have for all $x \in \mathbb{R}^d$

$$\begin{aligned} \text{pow}(x, b) &= \sum_{i=1}^n \lambda_i \text{pow}(x, b_i) \geq \sum_{i=1}^n \lambda_i \min_{1 \leq j \leq n} \text{pow}(x, b_j) = \min_{1 \leq j \leq n} \text{pow}(x, b_j) \\ &\geq \min_{b' \in \mathcal{B}} \text{pow}(x, b') \end{aligned} \tag{2.6}$$

When additionally, $x \in b$, we have $0 \geq \text{pow}(x, b)$. Therefore, there always exists some ball $b' \in \mathcal{B}$ such that the power of x to b' is non-positive, hence that $x \in b'$. It follows that $b \subseteq \bigcup \mathcal{B}$. \square

Remark 18. Equivalently, we have $\bigcup \text{CH}(\mathcal{B}) = \bigcup \mathcal{B}$. Indeed, if $b \in \text{CH}(\mathcal{B})$, then by Property 2.18 $b \subseteq \bigcup \mathcal{B}$. By taking the union over $\text{CH}(\mathcal{B})$ we obtain $\bigcup_{b' \in \text{CH}(\mathcal{B})} b' = \bigcup \text{CH}(\mathcal{B}) \subseteq \bigcup \mathcal{B}$. The reverse inclusion is obvious and we obtain equality. Conversely, if $\bigcup \text{CH}(\mathcal{B}) = \bigcup \mathcal{B}$, then consider $b \in \text{CH}(\mathcal{B})$. We have $b \subseteq \bigcup \text{CH}(\mathcal{B}) = \bigcup \mathcal{B}$, from which we obtain Property 2.18.

§ **Local and global inclusion** Consider two subsets $X, Y \subseteq \mathbb{R}^d$, and a point $u \in \mathbb{R}^d$. We say that X is locally contained in Y around point u if there exists a neighbourhood U of u such that $X \cap U \subseteq Y$. We say that $X \cap U \subseteq Y$ is a *local inclusion* of X in Y (around u), and $X \subseteq Y$ is a *global inclusion*.

§ **Result statement** In this section, we shall establish the converse of Property 2.18 in two situations:

- (1) when $\text{aff}(\mathcal{B})$ is an elliptic k -pencil with $k \leq d - 1$, and
- (2) when $\text{aff}(\mathcal{B})$ is a parabolic pencil.

Remark 19. By Remark 16, these are precisely every pencils with non-empty radical spheres.

Furthermore, we shall see that in the first situation, the local inclusion of b in $\bigcup \mathcal{B}$ is equivalent to a global inclusion, for all $b \in \text{aff}(\mathcal{B})$. Specifically, when $\text{aff}(\mathcal{B})$ is an elliptic pencil of dimension at most $d - 1$, if a ball $b \in \text{aff}(\mathcal{B})$ satisfies local inclusion in $\bigcup \mathcal{B}$ in the neighbourhood of any point of the radical sphere, we actually have global inclusion. As illustrated in Figure 2.6, this does not always hold when $\text{aff}(\mathcal{B})$ is parabolic.




Figure 2.6 – Let \mathcal{B} be the collection of the three black balls. Then $\text{aff}(\mathcal{B})$ is a parabolic pencil with radical sphere $\{c_0\}$. The red ball satisfies local inclusion in the neighbourhood of c_0 , but it is not globally included in $\text{CH}(\mathcal{B})$.

For elliptic pencils we have:

Lemma 2.19. *Let \mathcal{P} be an elliptic k -pencil, $k \leq d - 1$, and \mathcal{B} a finite collection of balls such that $\mathcal{P} = \text{aff}(\mathcal{B})$. Let $s_0 = s_0(\mathcal{P})$. For $b \in \mathcal{P}$, we have equivalence between:*

- (i) $b \in \text{CH}(\mathcal{B})$

(ii) $b \subseteq \bigcup \mathcal{B}$

(iii) there exists $u \in s_0$ and U a neighbourhood of u such that $b \cap U \subseteq \bigcup \mathcal{B}$

Whereas for parabolic pencils, we have:

Lemma 2.20. *Let \mathcal{P} be a parabolic pencil defined by the finite collection of balls \mathcal{B} . Denote by $b_0 = b_0(\mathcal{P})$ and let $\mathcal{B}_0 = \mathcal{B} \cup \{b_0\}$. For $b \in \mathcal{P}$, we have equivalence between:*

(i) $b \in \text{CH}(\mathcal{B}_0)$

(ii) $b \subseteq \bigcup \mathcal{B}$

Remark 20. Denote by $c_0 = c_0(\mathcal{P})$ the radical center of \mathcal{P} . When $\mathcal{P} = \text{aff}(\mathcal{B})$ is parabolic, notice that $b_0 = \{c_0\}$ is contained in every ball of pencil \mathcal{P} . Thus we have $\bigcup \mathcal{B}_0 = b_0 \cup (\bigcup \mathcal{B}) = \bigcup \mathcal{B}$. By Property 2.18 we deduce that for all $b \in \text{CH}(\mathcal{B}_0)$ we have $b \subseteq \bigcup \mathcal{B}$. Because $\text{CH}(\mathcal{B})$ and $\text{CH}(\mathcal{B}_0)$ may be strictly distinct, it is necessary to include b_0 to obtain equivalence. See Figure 2.7.

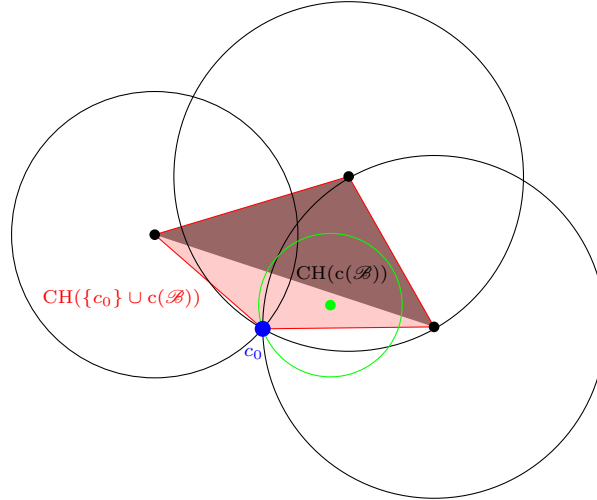


Figure 2.7 – Let \mathcal{B} be the collection of the three black balls, $\text{aff}(\mathcal{B})$ is a parabolic pencil with radical sphere $\{c_0\}$. The green ball does not belong to the convex hull of \mathcal{B} , but belongs to the complete pencil $\text{aff}(\mathcal{B})$. It satisfies global inclusion in $\text{CH}(\mathcal{B})$. To obtain equivalence for parabolic pencils, it is thus necessary to include the radical ball and consider the collection \mathcal{B}_0 instead of \mathcal{B} .

Note that from Remark 20, the two elliptic and parabolic cases are strictly distinct.

In order to prove these lemmas, we investigate the boundary of a convex subset of pencil in \mathbb{R}^d , $\partial(\bigcup \mathcal{B})$, and expose its dual relationship to the boundary polytope $\partial \text{CH}(c(\mathcal{B}))$.

2.2.1 Normal cones and power diagram

In order to state and prove the upcoming result on the boundary of convex subsets of pencils, we must first address the topic of normal cones, and how they relate to the power diagram of these pencil subsets. We first recall the definition of a power cell in this case, along with a few properties of the normal cones to a bounded convex polytope, before stating our result.

§ Power diagrams Let \mathcal{B} be a (potentially infinite) collection of balls, and $b \in \mathcal{B}$. Recall that the power cell of b with respect to \mathcal{B} is the collection of points whose power from all balls of \mathcal{B} is at least the power from b .

Remark 21. By definition, the radical space $V(\mathcal{B})$ of \mathcal{B} is the collection of points of \mathbb{R}^d with equal power from all balls of \mathcal{B} . Thus $V(\mathcal{B}) \subseteq \text{pcell}(b, \mathcal{B})$ for all $b \in \mathcal{B}$.

In this section, we will mostly consider the power diagram of $\text{CH}(\mathcal{B})$, with \mathcal{B} a finite collection of balls that defines a pencil. From inequality (2.6), observe that:

- if b is in the relative interior of $\text{CH}(\mathcal{B})$, then $\text{pcell}(b, \text{CH}(\mathcal{B})) = V(\mathcal{B})$
- if b belongs to an (open) face of $\text{CH}(\mathcal{B})$ with vertices b_1, \dots, b_k , then $\text{pcell}(b, \text{CH}(\mathcal{B})) = \bigcap_{i=1}^k \text{pcell}(b_i, \text{CH}(\mathcal{B}))$

Hence the power diagram of $\text{CH}(\mathcal{B})$ can be completely described by the power diagram of its vertices, \mathcal{B} .

§ Power cells and contribution to the global boundary For a ball $b \in \text{CH}(\mathcal{B})$, we call the *contribution of b to the (global) boundary of $\bigcup \mathcal{B}$* the points of $\partial(\bigcup \mathcal{B})$ that also belong to b , $\partial(\bigcup \mathcal{B}) \cap \partial b$. From the definition of a power cell, one easily deduces the inclusion $\partial(\bigcup \mathcal{B}) \cap \partial b \subseteq \text{pcell}(b, \text{CH}(\mathcal{B})) \cap \partial b$. Indeed, consider a point x in $\partial(\bigcup \mathcal{B}) \cap \partial b$. Then for all $b' \in \text{CH}(\mathcal{B})$, we have $\text{pow}(x, b') \geq 0 = \text{pow}(x, b)$ which implies by definition of power cells that $x \in \text{pcell}(b, \text{CH}(\mathcal{B}))$. In general, the reverse inclusion is false, see Figure 2.8. In order to partially bypass this issue, we require some insight on the structure of these power cells, to determine whether a given point in $\text{pcell}(b) \cap \partial b$ actually belongs to the boundary of $\bigcup \mathcal{B}$ or not. This structure will be given by the normal cones to the polytope $\text{CH}(\mathcal{B})$.

§ Normal cones to a convex set We now recall some properties of normal cones, in the context of convex sets. Note that the normal cone can be defined in a more general context, but we shall however focus on convex sets, and more particularly on convex polytopes as in [Ber87].

Definition 7 (Normal cone to a convex set). Let $X \subseteq \mathbb{R}^d$ be a convex set, and consider




Figure 2.8 – Let \mathcal{B} be the collection of the three balls. For ball b (green), its contribution to the boundary of $\cup \mathcal{B}$ is the red arc. However, the intersection of ∂b with $\text{pcell}(b, \text{CH}(\mathcal{B}))$ yields the extra point c_0 (blue). Thus, $\text{pcell}(b, \text{CH}(\mathcal{B})) \cap \partial b \not\subseteq \partial(\cup \mathcal{B}) \cap \partial b$.

$x \in \bar{X}$. We call *normal cone* to X at point x and denote $N(x, X)$ the collection of vectors

$$N(x, X) = \left\{ \mathbf{e} \in \mathbb{R}^d \mid \forall x' \in X, \langle \mathbf{e}, x' - x \rangle \leq 0 \right\}$$

Definition 8 (Unit normal cone to a convex set). Let $X \subseteq \mathbb{R}^d$ be a convex set, and consider $x \in \bar{X}$. We call *unit normal cone* to X at point x and denote $N_{\mathbf{1}}(x, X)$ the restriction of the normal cone to the unit (vector) sphere

$$N_{\mathbf{1}}(x, X) = \{ \mathbf{e} \in N(x, X) \mid \|\mathbf{e}\| = 1 \}$$

From the definition, it immediately follows that $N(x, X)$ is a convex cone that contains the null vector $\mathbf{0}$. Intuitively, the normal cone to a convex $X \subseteq \mathbb{R}^d$ at a point $x \in X$ can be seen as the set of outward normal directions to X at point x . Let $\mathbf{e} \in N(x, X)$ be non-null. Let H be the hyperplane with normal \mathbf{e} that goes through x , then X is contained in the closed half-space below H . Conversely, a non zero-vector \mathbf{e} belongs to the normal cone of X at point x only if the corresponding hyperplane through x is above X . See Figure 2.9.

Let $y \in X$ be a point distinct from x . The normal cones of x and y are equal if and only if x and y belong to the same open face of X . It thus makes sense to define the *normal cone of a face* as the normal cone of any of its interior point. For a face f , we denote its normal cone as $N(f, X)$.

If X has non-empty interior, its full dimensional (open) face coincides with its interior $\overset{\circ}{X}$. For this specific face, we have $N(\overset{\circ}{X}, X) = \{\mathbf{0}\}$. Let f_1 be a face of X and f_2 a coface of f_1 . Then $N(f_2, X) \subseteq N(f_1, X)$. More generally for two faces f_1 and f_2 , $N(f_1, X) \cap N(f_2, X)$ is

non-empty and coincides with the normal cone of the common coface of f_1 and f_2 of lowest dimensionality. From the above, it follows that for $x \in \partial X$, $\mathbf{e} \in N(x, X) \setminus \{\mathbf{0}\}$, and any $y \in X$, we have $\langle \mathbf{e}, y - x \rangle \leq 0$, with equality if and only if there is some boundary face f such that $x, y \in \bar{f} \subseteq H$, where H is the hyperplane through x with normal \mathbf{e} . Hence, a point x is a vertex of the polytope if and only if it possesses some *exclusive* outward normal direction, that is, there exists \mathbf{e} such that for all $y \in X \setminus \{x\}$, then $\langle \mathbf{e}, y - x \rangle < 0$.

Remark 22. When X is bounded, any non-zero direction is always the outward normal of some boundary face in X . Conversely, any boundary face always has a non-zero outward normal.

Because convex hulls are convex polytopes, we can consider the normal cones to either $\text{CH}(\mathcal{B})$ or $\text{CH}(c(\mathcal{B})) = c(\text{CH}(\mathcal{B}))$. Suppose $\text{aff}(\mathcal{B})$ is a pencil. The definition of a pencil guarantees that the normal cones to these two polytopes share a close relationship. In particular, $f_{\mathcal{B}}$ is a face of $\text{CH}(\mathcal{B})$ if and only if there is a face f_c of $\text{CH}(c(\mathcal{B}))$ such that $f_c = c(f_{\mathcal{B}})$, with $\dim f_{\mathcal{B}} = \dim f_c$. If \mathbf{e} is an outward normal to f_c , then it is also an outward normal to $f_{\mathcal{B}}$ and we thus have $N(f_c, \text{CH}(c(\mathcal{B}))) \subseteq N(f_{\mathcal{B}}, \text{CH}(\mathcal{B}))$. Strictly speaking, $\text{CH}(\mathcal{B})$ is a convex hull in \mathbb{R}^{d+1} whereas $\text{CH}(c(\mathcal{B}))$ is in \mathbb{R}^d . Hence the normal cone to $\text{CH}(\mathcal{B})$ at $f_{\mathcal{B}}$ should also span the $(d + 1)$ -th dimension. Because we will not use these vectors, we will mostly consider the normal cones to points in \mathbb{R}^d .




Figure 2.9 – The normal cone to polytope X at point x is the blue cone N_x . The red and green vectors along the boundary of N_x are shared by the points of the two adjacent edges to x .

§ **Normal cones and power cells** We now state the relationship between normal cones and power cells.

Lemma 2.21. Consider \mathcal{B} a finite collection of balls, $b \in \text{CH}(\mathcal{B})$ centered at c . Suppose \mathcal{B} spans a pencil and let $c_0 = c_0(\text{aff}(\mathcal{B}))$ the radical center. Then, the power cell of b to $\text{CH}(\mathcal{B})$

is

$$\begin{aligned} \text{pcell}(b, \text{CH}(\mathcal{B})) &= c_0 + N(c, \text{CH}(c(\mathcal{B}))) \\ &= \{c_0 + \mathbf{e} \mid \mathbf{e} \in N(c, \text{CH}(c(\mathcal{B})))\} \end{aligned}$$




Figure 2.10 – Consider the collection \mathcal{B} of the three balls b_1, b_2, b_3 . Its affine hull $\text{aff}(\mathcal{B})$ is a parabolic pencil with radical center c_0 . The normal cone to $\text{CH}(c(\mathcal{B}))$ at b_4 (black ball) has one unique element represented by the black arrow. The power cell of b_4 to $\text{CH}(\mathcal{B})$ is the half-line originating from c_0 in the direction of that black arrow, it satisfies Lemma 2.21. Likewise, the normal cone to $\text{CH}(c(\mathcal{B}))$ at b_1 (red ball) is represented by the red circular sector. The power cell of b_1 to $\text{CH}(\mathcal{B})$ is the red area, and Lemma 2.21 is again verified.

Refer to Figure 2.10 for an illustration. From this lemma, we obtain again the two introductory remarks we made on power cells, but through the properties of normal cones. Indeed when b is in the relative interior of $\text{CH}(\mathcal{B})$, then its normal cone N_b to $\text{CH}(\mathcal{B})$ coincides with the direction $\vec{V}(\mathcal{B})$ of the radical space $V(\mathcal{B})$; hence $\text{pcell}(b, \text{CH}(\mathcal{B})) = V(\mathcal{B})$. When b belongs to an open face of $\text{CH}(\mathcal{B})$ with vertices b_1, \dots, b_k , its normal cone is $N_b = \bigcap_{i=1}^k N(b_i, \text{CH}(\mathcal{B}))$ and thus $\text{pcell}(b, \text{CH}(\mathcal{B})) = \bigcap_{i=1}^k \text{pcell}(b_i, \text{CH}(\mathcal{B}))$.

In order to prove this result, we rely on the technical lemma below.

Lemma 2.22. Consider a pencil \mathcal{P} and let $c_0 = c_0(\mathcal{P})$. Let $b, b' \in \mathcal{P}$ and $x \in \mathbb{R}^d$. Then

$$\text{pow}(x, b) \leq \text{pow}(x, b') \iff \langle x - c_0, c(b') - c(b) \rangle \leq 0$$

Proof. Denote by $b_0 = b_0(\mathcal{P})$ the radical ball of \mathcal{P} . Following Lemma 2.8 we know that

$$\begin{aligned}\text{pow}(x, b) &= \text{pow}(x, b_0) - 2\langle x - c_0, c(b) - c_0 \rangle \\ \text{pow}(x, b') &= \text{pow}(x, b_0) - 2\langle x - c_0, c(b') - c_0 \rangle\end{aligned}$$

Thus

$$\text{pow}(x, b) - \text{pow}(x, b') = 2\langle x - c_0, c(b') - c(b) \rangle$$

These two expressions $\text{pow}(x, b) - \text{pow}(x, b')$ and $\langle x - c_0, c(b') - c(b) \rangle$ therefore have the same sign. \square

This lemma in fact implies that for every half-line originating from c_0 , all of its points belong to the same power cell(s).

Corollary 2.23. *Consider \mathcal{B} a collection of balls. Suppose \mathcal{B} spans a pencil and let $c_0 = c_0(\text{aff}(\mathcal{B}))$. Consider $c \in \text{CH}(c(\mathcal{B}))$ and the corresponding ball $b \in \text{CH}(\mathcal{B})$. Then*

$$\mathbf{e} \text{ is a unit normal vector to } \text{CH}(c(\mathcal{B})) \text{ at } c \iff c_0 + \mathbb{R}^+ \mathbf{e} \subseteq \text{pcell}(b, \text{CH}(\mathcal{B}))$$

Proof.

$$\begin{aligned}\mathbf{e} \in N_{\mathbf{1}}(c, \text{CH}(c(\mathcal{B}))) & \\ \iff \forall c' \in \text{CH}(c(\mathcal{B})), \langle \mathbf{e}, c' - c \rangle \leq 0 & \\ \iff \forall b' \in \text{CH}(\mathcal{B}), \langle \mathbf{e}, c(b') - c \rangle \leq 0 & \\ \iff \forall b' \in \text{CH}(\mathcal{B}), \forall x = c_0 + \lambda \mathbf{e}, \lambda \geq 0, \langle x - c_0, c(b') - c \rangle \leq 0 & \\ \iff \forall b' \in \text{CH}(\mathcal{B}), \forall x = c_0 + \lambda \mathbf{e}, \lambda \geq 0, \text{pow}(x, b) \leq \text{pow}(x, b') & \\ \iff \forall x = c_0 + \lambda \mathbf{e}, \lambda \geq 0, x \in \text{pcell}(b, \text{CH}(\mathcal{B})) & \\ \iff c_0 + \mathbb{R}^+ \mathbf{e} \subseteq \text{pcell}(b, \text{CH}(\mathcal{B})) & \square\end{aligned}$$

We can now prove Lemma 2.21.

Proof of Lemma 2.21. We want to prove that $\text{pcell}(b, \text{CH}(\mathcal{B})) = c_0 + N(c, \text{CH}(c(\mathcal{B})))$. We proceed by double inclusion. Following Corollary 2.23, we immediately deduce the inclusion

$$c_0 + N(c, \text{CH}(c(\mathcal{B}))) \subseteq \text{pcell}(b, \text{CH}(\mathcal{B}))$$

Conversely, consider a point $x \in \text{pcell}(b, \text{CH}(\mathcal{B}))$. Because c_0 belongs to every power cells of the power diagram of $\text{CH}(\mathcal{B})$, we can assume $x \neq c_0$. Let \mathbf{e} be the unit direction from c_0 to x , we show that \mathbf{e} is a normal to $\text{CH}(c(\mathcal{B}))$ at c , which then concludes the proof.

Suppose for a contradiction that \mathbf{e} is not a normal to $\text{CH}(c(\mathcal{B}))$ at c . Because \mathcal{B} is finite, $\text{CH}(\mathcal{B})$ is bounded. By Remark 22 there necessarily exists some point $c' \in \text{CH}(c(\mathcal{B}))$ such

that \mathbf{e} is a normal to $\text{CH}(c(\mathcal{B}))$ at c' . By definition of a normal, we have $\langle \mathbf{e}, c - c' \rangle \leq 0$. We claim that we in fact have the strict inequality $\langle \mathbf{e}, c - c' \rangle < 0$. Indeed if we had $\langle \mathbf{e}, c - c' \rangle = 0$, we would deduce that for all $c_* \in \text{CH}(c(\mathcal{B}))$, then $\langle \mathbf{e}, c_* - c \rangle = \langle \mathbf{e}, c_* - c' \rangle - \langle \mathbf{e}, c - c' \rangle \leq 0$. Vector \mathbf{e} would thus be a normal at c , which contradicts the assumption that it isn't.

Let b' the ball of $\text{CH}(\mathcal{B})$ centered at c' . Because $\langle \mathbf{e}, c - c' \rangle < 0$, we deduce by Lemma 2.22 that $\text{pow}(x, b') < \text{pow}(x, b)$ which implies $x \notin \text{pcell}(b, \text{CH}(\mathcal{B}))$. This contradicts our initial assumption, therefore \mathbf{e} must be a normal to $\text{CH}(c(\mathcal{B}))$ at c . \square

We are now ready to describe the boundary of $\bigcup \mathcal{B}$.

2.2.2 Structure of the boundary

We claimed in the previous section that having more insight on the structure of power cells of $\text{CH}(\mathcal{B})$ would help us investigate the boundary of the domain $\bigcup \mathcal{B}$. As per Lemma 2.21, normal cones are one way to structure these power cells. We now have a closer look at the boundary $\partial(\bigcup \mathcal{B})$.

To study this boundary, we shall define a function m_b , where b is a ball in a specific convex hull that we will specify soon after. Ultimately, we show that the image of the functions m_b are in fact related to the contribution of ball b to the boundary of $\bigcup \mathcal{B}$. On our way there, we will show several properties of these functions, making them desirable tools to manipulate the boundary.

§ The collection of balls \mathcal{B}_0 Throughout this section, \mathcal{B} denotes a finite collection of balls that spans a pencil with non-empty radical sphere. We let $b_0 = b_0(\text{aff}(\mathcal{B}))$, $c_0 = c_0(\text{aff}(\mathcal{B}))$ and $r_0^2 = r_0^2(\text{aff}(\mathcal{B}))$. If $\text{aff}(\mathcal{B})$ is elliptic, we let $\mathcal{B}_0 = \mathcal{B}$. If $\text{aff}(\mathcal{B})$ is parabolic, we let $\mathcal{B}_0 = \mathcal{B} \cup \{b_0\}$. In both cases, \mathcal{B}_0 remains a finite union of balls, and we have the two identities $\text{aff}(\mathcal{B}_0) = \text{aff}(\mathcal{B})$ and $\bigcup \mathcal{B}_0 = \bigcup \mathcal{B}$. However, the convex hulls of \mathcal{B}_0 and \mathcal{B} may be strictly distinct. That distinction will be important in our statements.

§ Definition of m_{b, \mathcal{B}_0} and first properties Let $b \in \text{CH}(\mathcal{B}_0)$ and consider a unit vector \mathbf{e} normal to $\text{CH}(c(\mathcal{B}_0))$ at $c(b)$. Because $c_0 \in b$, the intersection between the half-line $c_0 + \mathbb{R}^+ \mathbf{e}$ and ball b is a line segment that admits c_0 as an endpoint. We let $m_{b, \mathcal{B}_0}(\mathbf{e})$ be the other endpoint of that segment. When this segment reduces to the singleton $\{c_0\}$, we adopt the convention $m_{b, \mathcal{B}_0}(\mathbf{e}) = c_0$. To summarize, we thus define the function $m_{b, \mathcal{B}_0} : N_{\mathbf{1}}(c(b), \text{CH}(c(\mathcal{B}_0))) \rightarrow b$ such that

$$m_{b, \mathcal{B}_0}(\mathbf{e}) = \begin{cases} c_0 & \text{if } (c_0 + \mathbb{R}^+ \mathbf{e}) \cap b = \{c_0\} \\ (c_0 + \mathbb{R}^+ \mathbf{e}) \cap \partial b \setminus \{c_0\} & \text{otherwise} \end{cases}$$

Remark 23. Throughout this section, the collection \mathcal{B}_0 can be considered constant, hence we let $m_b = m_{b, \mathcal{B}_0}$. Later on, we will have to consider distinct collections \mathcal{B}_1 and \mathcal{B}_2 with

$b \in \mathcal{B}_1 \cap \mathcal{B}_2$. At that point the notation m_b will be ambiguous, and we will revert to the explicit notations m_{b, \mathcal{B}_1} and m_{b, \mathcal{B}_2} .

For most of the discussion regarding function m_b and its properties, refer to Figures 2.10 and 2.11.




Figure 2.11 – The three balls constitute the collection \mathcal{B} , and \mathcal{B}_0 additionally contains the singleton ball $\{c_0\}$. Consider the ball b centered at c (red). The unit normal cone N_1^c to $\text{CH}(c(\mathcal{B}_0))$ at point c is represented by the red arrows and circular region. The unit vector \mathbf{e} (green) belongs to that unit normal cone. Note that c_0 (blue) does not belong to the image of N_1^c by m_b (red arc), but still belongs to the contribution of b to the global boundary.

Remark 24. Note that the unit normal cone of a point in $\text{CH}(c(\mathcal{B}_0))$ may very well be empty, in which case the function m_b for the ball b corresponding to that point is the trivial empty function. From the properties of unit normal cones, it follows that the function m_b is non-trivial if and only if $c(b) \in \partial \text{CH}(c(\mathcal{B}_0))$. Note that $\partial \text{CH}(c(\mathcal{B}_0))$ is the boundary of polytope $\text{CH}(c(\mathcal{B}_0))$ and not its relative boundary. Hence when $\text{aff}(\mathcal{B}_0)$ is a k -flat, $k \leq d - 1$, every ball $b \in \text{CH}(\mathcal{B}_0)$ has its center $c(b)$ in the boundary of $\text{CH}(c(\mathcal{B}_0))$.

Remark 25. Suppose that $\text{aff}(\mathcal{B}_0)$ is a parabolic pencil, and that $b_0 \in \partial \text{CH}(\mathcal{B}_0)$. Then m_{b_0} is the constant function that maps every admissible vector to point c_0 . To see this, notice that $m_{b_0}(\mathbf{e}) \in b_0 = \{c_0\}$.

Remark 26. We always have $m_b(\mathbf{e}) \in \partial b$ for all admissible pairs b, \mathbf{e} . Indeed, the definition ensures that the line segment (possibly a singleton) $[c_0 m_b(\mathbf{e})]$ coincides with $(c_0 + \mathbb{R}^+\mathbf{e}) \cap b$. Thus any neighbourhood of $m_b(\mathbf{e})$ always contains points in the complement of b .

Remark 27. Suppose $\text{aff}(\mathcal{B}_0)$ is elliptic, we know that $c_0 \notin \partial b$ for all $b \in \text{CH}(\mathcal{B}_0)$. By Remark 26 we deduce that we actually have $m_b(\mathbf{e}) = (c_0 + \mathbb{R}^+\mathbf{e}) \cap \partial b$, for all admissible pairs b, \mathbf{e} .

One key property of the function m_b , $b \in \text{CH}(\mathcal{B}_0)$, is that its image is contained within the global boundary of $\bigcup \mathcal{B}_0$, and not just the boundary of ball b . Indeed, from the relationship between power cells and normal cones, we easily deduce the lemma below.

Lemma 2.24. *Let \mathcal{B} be a finite collection of balls. Suppose \mathcal{B} spans a pencil and let $c_0 = c_0(\text{aff}(\mathcal{B}))$. Let $b \in \text{CH}(\mathcal{B})$, and \mathbf{e} a unit normal vector to $\text{CH}(c(b))$ at point $c(b)$. Then the half-line $c_0 + \mathbb{R}^+\mathbf{e}$ satisfies*

$$(c_0 + \mathbb{R}^+\mathbf{e}) \cap b = (c_0 + \mathbb{R}^+\mathbf{e}) \cap \left(\bigcup \mathcal{B} \right)$$

Proof. By Corollary 2.23 we know that $c_0 + \mathbb{R}^+\mathbf{e} \subseteq \text{pcell}(b, \text{CH}(\mathcal{B}))$. The identity then follows from the definition of a power cell. \square

Property 2.25. $m_b(\mathbf{e}) \in \partial(\bigcup \mathcal{B})$

Proof. Immediate from Lemma 2.24 and Remark 26. \square

Although the various functions m_b for distinct balls $b \in \text{CH}(\mathcal{B}_0)$ do not in general share the same definition domain, part of these domains may overlap. For those shared unit vectors, the functions coincide.

Property 2.26. *Consider two balls $b_1, b_2 \in \text{CH}(\mathcal{B}_0)$, and \mathbf{e} a unit vector. Suppose \mathbf{e} is normal to $\text{CH}(c(\mathcal{B}_0))$ both at point $c(b_1)$ and point $c(b_2)$. Then $m_{b_1}(\mathbf{e}) = m_{b_2}(\mathbf{e})$.*

Proof. Also immediate from Lemma 2.24. \square

§ Continuity and injectivity We now show that m_b is continuous for all $b \in \text{CH}(\mathcal{B}_0)$. We first seek a more explicit expression of $m_b(\mathbf{e})$. We let $\rho_b(\mathbf{e}) = \|c_0 - m_b(\mathbf{e})\|$. By definition $m_b(\mathbf{e})$ lies on the half-line $c_0 + \mathbb{R}^+\mathbf{e}$, thus it follows that $m_b(\mathbf{e}) = c_0 + \rho_b(\mathbf{e})\mathbf{e}$. We now proceed to compute the length $\rho_b(\mathbf{e})$.

Lemma 2.27. *Let $b = b(c, r)$ and $b_0 = b(c_0, r_0)$ be two balls such that $c_0 \in b$ and $r_0^2 = r^2 - \|c - c_0\|^2$. Then $b = \{c_0 + \rho\mathbf{e} \mid 0 \leq \rho \leq \rho_{\mathbf{e}}, \|\mathbf{e}\| = 1\}$, where $\rho_{\mathbf{e}} = \langle \mathbf{e}, c - c_0 \rangle + \sqrt{r_0^2 + \langle \mathbf{e}, c - c_0 \rangle^2}$.*

Proof. Because $c_0 \in b$, we have $r_0^2 \geq 0$ and therefore $\rho_{\mathbf{e}}$ is always a non-negative real number. Consider a unit vector \mathbf{e} . We seek the limit value $\rho_{\text{lim}} \geq 0$ such that $c_0 + \rho_{\text{lim}}\mathbf{e} \in \partial b$ and show that it is indeed equal to $\rho_{\mathbf{e}}$. Consider $x = c_0 + \rho\mathbf{e}$ with $\rho \geq 0$. Note that the constraints on b and b_0 guarantee the existence of a pencil \mathcal{P} such that $b, b_0 \in \mathcal{P}$, and $b_0(\mathcal{P}) = b_0$. By

Lemma 2.8 we obtain

$$\begin{aligned}
x \in b &\iff \text{pow}(x, b) \leq 0 \\
&\iff \text{pow}(x, b_0) - 2\langle x - c_0, c - c_0 \rangle \leq 0 \\
&\iff \rho^2 - r_0^2 - 2\rho\langle \mathbf{e}, c - c_0 \rangle \leq 0 \\
&\iff (\rho - \langle \mathbf{e}, c - c_0 \rangle)^2 \leq r_0^2 + \langle \mathbf{e}, c - c_0 \rangle^2 \\
&\iff \langle \mathbf{e}, c - c_0 \rangle - \sqrt{r_0^2 + \langle \mathbf{e}, c - c_0 \rangle^2} \leq \rho \leq \langle \mathbf{e}, c - c_0 \rangle + \sqrt{r_0^2 + \langle \mathbf{e}, c - c_0 \rangle^2} \\
&\iff 0 \leq \rho \leq \rho_{\mathbf{e}}
\end{aligned}
\quad \square$$

Corollary 2.28. $\rho_b(\mathbf{e}) = \|c_0 - m_b(\mathbf{e})\| = \langle \mathbf{e}, c - c_0 \rangle + \sqrt{r_0^2 + \langle \mathbf{e}, c - c_0 \rangle^2}$

Property 2.29. *The function m_b is continuous for all $b \in \partial \text{CH}(\mathcal{B}_0)$.*

Proof. Given a ball b , we deduce from Corollary 2.28 that the mapping $\mathbf{e} \mapsto \rho_{\mathbf{e}} = \|c_0 - m_b(\mathbf{e})\|$ is continuous with respect to the unit vector \mathbf{e} . It follows that m_b is the composition of continuous functions and is thus continuous itself. \square

We now argue that the function m_b can mostly be considered injective.

Property 2.30. *If $\text{aff}(\mathcal{B}_0)$ is an elliptic pencil, then m_b is injective. If $\text{aff}(\mathcal{B}_0)$ is a parabolic pencil, then the restriction of m_b to $N_{\mathbf{1}}(c(b), \text{CH}(c(\mathcal{B}_0))) \setminus N_{\mathbf{1}}(c_0, \text{CH}(c(\mathcal{B}_0)))$ is injective.*

Proof. The principle of the proof relies on the fact that for any two distinct unit vectors \mathbf{e}_1 and \mathbf{e}_2 , the two half-lines $c_0 + \mathbb{R}^+\mathbf{e}_1$ and $c_0 + \mathbb{R}^+\mathbf{e}_2$ only intersect at their origin c_0 . It follows that if $m_b(\mathbf{e}_1) = m_b(\mathbf{e}_2)$, then necessarily the point of collision is c_0 , and $\rho_{\mathbf{e}_1} = 0 = \rho_{\mathbf{e}_2}$. In order to prove that m_b is an injection on the considered domain, it thus suffices to prove that for all admissible unit vector \mathbf{e} , we have $\|c_0 - m_b(\mathbf{e})\| = \rho_{\mathbf{e}} > 0$.

We start with the elliptic case. Recall as per Corollary 2.7 that for elliptic pencils, we have $r_0^2 > 0$. It follows that for every ball $b \in \partial \text{CH}(\mathcal{B}_0)$ and unit vector $\mathbf{e} \in N_{\mathbf{1}}(c(b), \text{CH}(c(\mathcal{B}_0)))$, we have $\rho_{\mathbf{e}} = \|c_0 - m_b(\mathbf{e})\| > 0$. Thus, m_b is an injection.

We now address the parabolic case. Consider $b \in \partial \text{CH}(\mathcal{B}_0)$ and let \mathbf{e} be a unit vector that is normal to $\text{CH}(c(\mathcal{B}_0))$ at point $c(b)$, but not at point c_0 . Because \mathbf{e} is normal to the convex hull at $c(b)$, we deduce that for all $c' \in \text{CH}(c(\mathcal{B}_0))$ we have $\langle \mathbf{e}, c' - c(b) \rangle \leq 0$. Because \mathbf{e} is not normal to the convex hull at c_0 , we must have $\langle \mathbf{e}, c_0 - c(b) \rangle < 0$. Indeed, if we instead had $\langle \mathbf{e}, c_0 - c(b) \rangle = 0$, we would then deduce that $\langle \mathbf{e}, c' - c_0 \rangle \leq 0$ for all $c' \in \text{CH}(c(\mathcal{B}_0))$, thus contradicting the assumption that \mathbf{e} is not normal at c_0 . Recall by Corollary 2.7 that for parabolic pencils, we have $r_0^2 = 0$. We can thus simplify the expression $\rho_{\mathbf{e}} = 2\langle \mathbf{e}, c(b) - c_0 \rangle$. This immediately implies that $\rho_{\mathbf{e}} > 0$. Because this holds for every $\mathbf{e} \in N_{\mathbf{1}}(c(b), \text{CH}(c(\mathcal{B}_0))) \setminus N_{\mathbf{1}}(c_0, \text{CH}(c(\mathcal{B}_0)))$, thus the function m_b restricted to that domain is injective. \square

Remark 28. Note that when $\text{aff}(\mathcal{B}_0)$ is a parabolic pencil and $b = b_0$, we then restrict m_{b_0} to an empty domain.

Remark 29. When $\text{aff}(\mathcal{B}_0)$ is a parabolic pencil of dimension at most $d - 1$, its radical space V has dimension at least 1. Because \vec{V} is contained in every normal cone to $\text{CH}(c(\mathcal{B}_0))$, it follows that $N_{\mathbb{1}}(c(b), \text{CH}(c(\mathcal{B}_0))) \cap N_{\mathbb{1}}(c_0, \text{CH}(c(\mathcal{B}_0)))$ always contain at least two distinct elements for which $m_b(e) = c_0$. Thus m_b is in general not injective on its complete domain.

Remark 30. When $\text{aff}(\mathcal{B}_0)$ is a parabolic d -pencil, then $N_{\mathbb{1}}(c_0, \text{CH}(c(\mathcal{B}_0)))$ can possibly be empty or reduced to a single element. In those particular cases, m_b is an injection for all $b \in \partial \text{CH}(\mathcal{B}_0)$.

Remark 31. Let $b \in \text{CH}(\mathcal{B}_0)$, and suppose $\text{aff}(\mathcal{B}_0)$ is parabolic. The proof above also shows that if $m_b(\mathbf{e}) = c_0$, then $\mathbf{e} \in N_{\mathbb{1}}(c_0, \text{CH}(c(\mathcal{B}_0)))$.

Notice from Property 2.26, Remarks 25 and 31 that we obtain the property below.

Property 2.31. *Suppose $\text{aff}(\mathcal{B}_0)$ is parabolic. Let $b \in \text{CH}(\mathcal{B}_0)$ be centered at c . Let $N_{\mathbb{1}}^c$ (resp. $N_{\mathbb{1}}^{c_0}$) be the unit normal cone to $\text{CH}(c(\mathcal{B}_0))$ at point c (resp. c_0). Consider $\mathbf{e} \in N_{\mathbb{1}}^c$. Then*

$$m_b(\mathbf{e}) = c_0 \iff \mathbf{e} \in N_{\mathbb{1}}^{c_0}$$

We state now two corollaries to the properties above that will be useful later on, when we discuss the relationship between the image of m_b and the contribution of b to the boundary of $\bigcup \mathcal{B}_0$.

Corollary 2.32. *Consider $b \in \text{CH}(\mathcal{B}_0)$ centered at c , and let $N_{\mathbb{1}}^c$ be the unit normal vector to $\text{CH}(c(\mathcal{B}_0))$ at c .*

- *If $\text{aff}(\mathcal{B}_0)$ is an elliptic $(d - 1)$ -pencil and b lies in the relative interior of $\text{CH}(\mathcal{B}_0)$, then $m_b(N_{\mathbb{1}}^c)$ is the 0-sphere $s_0(\text{aff}(\mathcal{B}_0))$ and consists of two singletons.*
- *Otherwise, $m_b(N_{\mathbb{1}}^c)$ is path-connected.*

Proof. We investigate the situations when $N_{\mathbb{1}}^c$ is not path-connected. Indeed, if $N_{\mathbb{1}}^c$ is path-connected, the continuity of m_b ensures that $m_b(N_{\mathbb{1}}^c)$ will also be path-connected. First note that the normal cone $N(c, \text{CH}(c(\mathcal{B}_0)))$ is convex, therefore it is always path-connected. The unit normal cone can only contain several path-connected components when $N(c, \text{CH}(c(\mathcal{B}_0)))$ is a one-dimensional vector space. In that case, the unit normal cone $N_{\mathbb{1}}^c$ contains two opposite, collinear vectors. Every path in $N(c, \text{CH}(c(\mathcal{B}_0)))$ connecting these vectors go through the null vector, and thus there are no corresponding path in the unit normal cone.

This situation only occurs when $\dim \text{aff}(\mathcal{B}_0) = d - 1$, and b is in the relative interior of $\text{CH}(\mathcal{B}_0)$. In particular, we deduce that the normal cone at c must coincide with the direction of the radical space of $\text{aff}(\mathcal{B}_0)$, that is $N(c, \text{CH}(c(\mathcal{B}_0))) = \vec{V}(\text{aff}(\mathcal{B}_0))$. We now distinguish two cases.

(a) When $\text{aff}(\mathcal{B}_0)$ is parabolic, recall that by definition $m_b(\mathbf{e}) \in (c_0 + \mathbb{R}^+\mathbf{e}) \cap \partial b$. Thus, $m_b(N_{\mathbb{1}}^c) \subseteq V(\text{aff}(\mathcal{B}_0)) \cap \partial b = s_0(\text{aff}(\mathcal{B}_0)) = \{c_0\}$. Therefore the image of m_b is still necessarily path-connected.

(b) When $\text{aff}(\mathcal{B}_0)$ is elliptic, recall that by Remark 27 we have $m_b(\mathbf{e}) = (c_0 + \mathbb{R}^+\mathbf{e}) \cap \partial b$. Thus $m_b(N_{\mathbb{1}}^c) = V(\text{aff}(\mathcal{B}_0)) \cap \partial b = s_0(\text{aff}(\mathcal{B}_0))$. In other words, the image of m_b coincides with the radical sphere of $\text{aff}(\mathcal{B}_0)$. That radical sphere is a 0-sphere that consists of two singletons.

In conclusion, the image of m_b is path-connected, except when b is in the relative interior of $\text{CH}(\mathcal{B}_0)$ and $\text{aff}(\mathcal{B}_0)$ is elliptic with dimension $d - 1$. \square

Corollary 2.33. *Suppose b belongs to an open boundary k -face of $\text{CH}(\mathcal{B}_0)$. When $\text{aff}(\mathcal{B}_0)$ is parabolic, suppose additionally that this open k -face does not contain b_0 in its closure. Let c be the center of b and $N_{\mathbb{1}}^c$ the unit normal cone to $\text{CH}(c(\mathcal{B}_0))$ at c . Then, $m_b(N_{\mathbb{1}}^c)$ has intrinsic dimension $d - 1 - k$.*

Proof. We start with the elliptic case. As per Property 2.30, we know that m_b is injective. It is thus bijective from $N_{\mathbb{1}}^c$ to $m_b(N_{\mathbb{1}}^c)$. Because $N_{\mathbb{1}}^c$ is a subset of the unit vector sphere with intrinsic dimension $d - 1 - k$, the result follows.

Assume that $\text{aff}(\mathcal{B}_0)$ is parabolic. Let $N_{\mathbb{1}}^{c_0}$ be the unit normal cone to $\text{CH}(c(\mathcal{B}_0))$ at c_0 . As per Property 2.30, m_b is injective when restricted to $N_{\mathbb{1}}^c \setminus N_{\mathbb{1}}^{c_0}$. By assumption, b is in an open k -face of $\text{CH}(\mathcal{B}_0)$ that does not contain b_0 in its closure. Thus, $N_{\mathbb{1}}^c \setminus N_{\mathbb{1}}^{c_0}$ is non-empty with intrinsic dimension $d - 1 - k$. The restriction of m_b to $N_{\mathbb{1}}^c \setminus N_{\mathbb{1}}^{c_0}$ is thus a subset of ∂b with intrinsic dimension $d - 1 - k$. In addition, notice as per Remark 25 and Property 2.26 that $m_b(N_{\mathbb{1}}^c \cap N_{\mathbb{1}}^{c_0}) = \{c_0\}$. Thus, the image of m_b over its whole domain is still a spherical patch of dimension $d - 1 - k$. \square

Remark 32. When $\text{aff}(\mathcal{B}_0)$ is parabolic and the open k -face contains b_0 in its closure, then $m_b(N_{\mathbb{1}}^c)$ is the singleton sphere $\{c_0\}$.

§ Exclusivity In the paragraph above, we saw that for any $b \in \text{CH}(\mathcal{B}_0)$, the function m_b can for the most part be considered injective if we set aside the image point c_0 . We here establish a stronger result.

Let b be a ball in $\text{CH}(\mathcal{B}_0)$ and suppose the normal cone at $c(b)$ contains an exclusive normal vector. That is, there exists \mathbf{e} such that for all $c' \in \text{CH}(c(\mathcal{B}_0))$, we have $\langle \mathbf{e}, c' - c(b) \rangle < 0$. Then we argue that the corresponding boundary point $m_b(\mathbf{e})$ is only covered by ball b .

Obviously, for all $b' \neq b$ in the convex hull $\text{CH}(\mathcal{B}_0)$, the function $m_{b'}$ is undefined for \mathbf{e} . This however does not prevent b' from covering other boundary points, that may not belong to the image of $m_{b'}$. This property thus extends beyond simple injectivity of the functions m_b , and in a sense indicates a one-to-one mapping for the overall shape $\bigcup \mathcal{B}_0$.

First we address a technical lemma.

Lemma 2.34. *Let $b \in \text{CH}(\mathcal{B}_0)$ and $x \in \partial(\bigcup \mathcal{B}_0)$. Then*

$$x \in b \iff x - c_0 \in N(c(b), \text{CH}(c(\mathcal{B}_0)))$$

Proof. We start by showing that with the assumption $x \in \partial(\bigcup \mathcal{B}_0)$, we have equivalence between $x \in b$ and $x \in \text{pcell}(b, \text{CH}(\mathcal{B}_0))$. Indeed, suppose $x \in b$. Because $x \in \partial(\bigcup \mathcal{B}_0)$, we deduce that for all $b' \in \text{CH}(\mathcal{B}_0)$ we have $\text{pow}(x, b') \geq 0$. If there was some ball $b' \in \text{CH}(\mathcal{B}_0)$ for which $\text{pow}(x, b') < 0$, we would deduce by Property 2.18 that $x \in \overset{\circ}{b'} \subseteq (\bigcup \mathcal{B}_0)^\circ$. Thus if $x \in b$, we have for all $b' \in \text{CH}(\mathcal{B}_0)$ that $\text{pow}(x, b) \leq 0 \leq \text{pow}(x, b')$ and thus $x \in \text{pcell}(b, \text{CH}(\mathcal{B}_0))$. Conversely suppose that $x \in \text{pcell}(b, \text{CH}(\mathcal{B}_0))$. Because $x \in \bigcup \mathcal{B}_0$, there is some ball $b' \in \mathcal{B}_0$ such that $\text{pow}(x, b') \leq 0$. Since $\mathcal{B}_0 \subseteq \text{CH}(\mathcal{B}_0)$, we deduce that $\text{pow}(x, b) \leq \text{pow}(x, b') \leq 0$ and $x \in b$. Therefore, we have equivalence between $x \in b$ and $x \in \text{pcell}(b, \text{CH}(\mathcal{B}_0))$.

Using Lemma 2.21 we obtain

$$\begin{aligned} x \in b &\iff x \in \text{pcell}(b, \text{CH}(\mathcal{B}_0)) \\ &\iff x \in c_0 + N(c(b), \text{CH}(c(\mathcal{B}_0))) \\ &\iff x - c_0 \in N(c(b), \text{CH}(c(\mathcal{B}_0))) \quad \square \end{aligned}$$

Property 2.35. *Consider $b \in \text{CH}(\mathcal{B}_0)$. Let \mathbf{e} be a unit normal vector to $\text{CH}(c(\mathcal{B}_0))$ at point $c(b)$. If \mathbf{e} is an exclusive normal for point $c(b)$, then b is the only ball that contains $\mathbf{m}_b(\mathbf{e})$.*

Proof. Immediate by application of Lemma 2.34 to point $x = \mathbf{m}_b(\mathbf{e})$. □

§ Image of \mathbf{m}_b and contribution to the global boundary We now describe the image of function \mathbf{m}_b . We are in particular interested in its link with the contribution of a ball b to the global boundary of $\bigcup \mathcal{B}_0$.

Lemma 2.36. *Let $b \in \text{CH}(\mathcal{B}_0)$ centered at c . Let $N_{\mathbb{1}}^c$ be the unit normal cone to $\text{CH}(c(\mathcal{B}_0))$ at c .*

$$\mathbf{m}_b(N_{\mathbb{1}}^c) \subseteq \partial b \cap \partial(\bigcup \mathcal{B}_0) \subseteq \{c_0\} \cup \mathbf{m}_b(N_{\mathbb{1}}^c)$$

Proof. The first inclusion is immediate from Property 2.25. For the second inclusion, we show that $\partial b \cap \partial(\bigcup \mathcal{B}_0) \setminus \{c_0\} \subseteq \mathbf{m}_b(N_{\mathbb{1}}^c)$. Consider $x \in \partial b \cap \partial(\bigcup \mathcal{B}_0)$, $x \neq c_0$. By Lemma 2.34 we have $x - c_0 \in N(c, \text{CH}(c(\mathcal{B}_0)))$. Let \mathbf{e} be the unit vector from c_0 to x . Because $x - c_0$ is in the normal cone to $\text{CH}(c(\mathcal{B}_0))$ at c , thus \mathbf{e} is in the unit normal cone and we can consider $\mathbf{m}_b(\mathbf{e})$. Observe that we necessarily have $\mathbf{m}_b(\mathbf{e}) = x$. Indeed, $x \in (c_0 + \mathbb{R}^+\mathbf{e}) \cap \partial b$ and $x \neq c_0$. Therefore, $x \in \mathbf{m}_b(N_{\mathbb{1}}^c)$, which concludes the proof. □

Property 2.37. *Let $b \in \text{CH}(\mathcal{B}_0)$ centered at c . Let $N_{\mathbb{1}}^c$ be the unit normal cone to $\text{CH}(c(\mathcal{B}_0))$ at c .*

- If $c_0 \notin \partial(\bigcup \mathcal{B}_0)$, or $c_0 \in m_b(N_{\mathbb{1}}^c)$, then the contribution of b to the global boundary is

$$\partial b \cap \partial\left(\bigcup \mathcal{B}_0\right) = m_b(N_{\mathbb{1}}^c)$$

- Otherwise if $c_0 \in \partial(\bigcup \mathcal{B}_0)$ and $c_0 \notin m_b(N_{\mathbb{1}}^c)$, then the contribution of b to the global boundary is

$$\partial b \cap \partial\left(\bigcup \mathcal{B}_0\right) = \{c_0\} \cup m_b(N_{\mathbb{1}}^c)$$

Proof. The first case is a direct consequence of Lemma 2.36. In the second case, notice that c_0 may contribute to the global boundary $\partial(\bigcup \mathcal{B}_0)$ only if $\text{aff}(\mathcal{B}_0)$ is parabolic. Recall that in parabolic pencils, every ball contains c_0 in its boundary. From there, Lemma 2.36 again yields the equality. \square

Remark 33. Assume that $\text{aff}(\mathcal{B}_0)$ is a parabolic pencil. There are cases where c_0 does not belong to the image of m_b , see for instance Figure 2.11. It follows that the property above highlights one key difference between $\partial b \cap \text{pcell}(b, \text{CH}(\mathcal{B}_0))$ and the image $m_b(N_{\mathbb{1}}^c)$. The former will always contain the radical center c_0 of the pencil, contrary to the latter.

Remark 34. In the second case, $\{c_0\}$ and the image of m_b are separated. Indeed, unit normal cones are closed and m_b is continuous, thus the image of m_b is also closed. Two closed disjoint subsets are separated, thus $\{c_0\}$ and the image of m_b are separated.

We summarize below several properties regarding the image of m_b .

Property 2.38. *Let $b \in \text{CH}(\mathcal{B}_0)$ centered at c . Let $N_{\mathbb{1}}^c$ be the unit normal cone to $\text{CH}(c(\mathcal{B}_0))$ at c .*

- If $c_0 \in \partial(\bigcup \mathcal{B}_0)$ and $c_0 \notin m_b(N_{\mathbb{1}}^c)$, then the contribution of b to the global boundary consists of two path-connected components. One is the spherical patch $m_b(N_{\mathbb{1}}^c)$, the other is the singleton $\{c_0\}$.
- If b lies in the relative interior of $\text{CH}(\mathcal{B}_0)$ and $\text{aff}(\mathcal{B}_0)$ is an elliptic $(d-1)$ -pencil, then the contribution of b to the global boundary has two (path-)connected components. It is the 0-sphere $s_0(\text{aff}(\mathcal{B}_0))$.
- Otherwise the contribution of b to the global boundary is path-connected. It is the spherical patch $m_b(N_{\mathbb{1}}^c)$.

Proof. We proceed in order. As a general remark, recall that the image of m_b is always a “spherical patch”: it is the subset of a sphere, whose dimension is given by Corollary 2.33.

Assume that $c_0 \in \partial(\bigcup \mathcal{B}_0)$ and $c_0 \notin m_b(N_{\mathbb{1}}^c)$. By Property 2.37 above, the contribution of b to the boundary of $\bigcup \mathcal{B}_0$ is $\{c_0\} \cup m_b(N_{\mathbb{1}}^c)$. As per Remark 34, the singleton $\{c_0\}$ and $m_b(N_{\mathbb{1}}^c)$ are separated. Thus $\{c_0\}$ constitutes its own path-connected component. As for the spherical patch $m_b(N_{\mathbb{1}}^c)$, recall that c_0 may contribute to the boundary only if pencil $\text{aff}(\mathcal{B}_0)$ is parabolic. Corollary 2.32 thus guarantees that it is one single path-connected component.

Suppose now that either $c_0 \notin \partial(\bigcup \mathcal{B}_0)$ or $c_0 \in m_b(N_{\mathbf{1}}^c)$. Property 2.37 implies that necessarily, the contribution of b to the global boundary of $\bigcup \mathcal{B}_0$ is $m_b(N_{\mathbf{1}}^c)$. By Corollary 2.32, we have two more subcases. If $\text{aff}(\mathcal{B}_0)$ is an elliptic $(d-1)$ -pencil, and b is in the relative interior of $\text{CH}(\mathcal{B}_0)$, then the image of m_b is the radical sphere of $\text{aff}(\mathcal{B}_0)$, which is indeed a 0-sphere with two path-connected components.

Otherwise, in every remaining case the image of m_b is path-connected. \square

§ On the distinction between \mathcal{B}_0 and \mathcal{B} We started this section by introducing the collection \mathcal{B}_0 , and stated our results in terms of \mathcal{B}_0 instead of \mathcal{B} . Although the difference between \mathcal{B} and \mathcal{B}_0 may seem trivial, it is in fact important in the parabolic case.

Indeed, when $\text{aff}(\mathcal{B}) = \text{aff}(\mathcal{B}_0)$ is a parabolic pencil, the point c_0 plays a central role in the results above. When c_0 is part of the global boundary $\partial(\bigcup \mathcal{B}_0)$, one must indeed be careful of whether c_0 is part of the image of m_b or not. By nature of the pencil, c_0 will be part of the boundary of every ball in the pencil, but Property 2.37 explicitly shows that there are two distinct behaviours depending on the membership of c_0 to the image of m_b .

In order to more efficiently characterize which ball b of the collection may or may not cover c_0 in its image of m_b , it is helpful to rely on Property 2.31. This property directly states that any unit vector whose image by m_b is c_0 must belong to $N_{\mathbf{1}}^{c_0}$, the unit normal cone to $\text{CH}(c(\mathcal{B}_0))$ at point c_0 . This not only gives a characterization of $N_{\mathbf{1}}^{c_0}$, but also yields a test of the membership of c_0 that only relies on normal cones. Note that this unit normal cone $N_{\mathbf{1}}^{c_0}$ is only guaranteed to be properly defined if c_0 belongs to the convex hull of centers, hence the use of the collection \mathcal{B}_0 .

For completeness, we now show that the point c_0 is part of the global boundary if and only if the unit normal cone $N_{\mathbf{1}}^{c_0}$ is non-empty. Indeed, so far c_0 was always treated as an exception. We only know through Lemma 2.34 that every point $x \in \partial(\bigcup \mathcal{B}_0) \setminus \{c_0\}$ always admit some pair b, \mathbf{e} for which $x = m_b(\mathbf{e})$. However, we do not yet have such a result for c_0 itself.

Property 2.39. *Suppose $\text{aff}(\mathcal{B}_0)$ is parabolic and let $N_{\mathbf{1}}^{c_0}$ be the unit normal vector to $\text{CH}(c(\mathcal{B}_0))$ at point c_0 . Then,*

$$c_0 \in \partial\left(\bigcup \mathcal{B}_0\right) \iff N_{\mathbf{1}}^{c_0} \neq \emptyset$$

Proof. By Properties 2.31 and 2.25 we immediately have the implication $N_{\mathbf{1}}^{c_0} \neq \emptyset \implies c_0 \in \partial(\bigcup \mathcal{B}_0)$. Conversely, suppose that $N_{\mathbf{1}}^{c_0} = \emptyset$. This implies that c_0 lies in the interior of $\text{CH}(c(\mathcal{B}_0))$. By definition, $\text{CH}(c(\mathcal{B}_0))$ is the collection of centers of balls in $\text{CH}(\mathcal{B}_0)$, it is thus entirely covered by the union of ball itself. Therefore, $c_0 \in (\text{CH}(c(\mathcal{B}_0)))^\circ \subseteq (\bigcup \text{CH}(\mathcal{B}_0))^\circ = (\bigcup \mathcal{B}_0)^\circ$. It follows that c_0 cannot be part of the boundary of $\bigcup \mathcal{B}_0$. \square

Following this last property we now obviously have

$$\partial\left(\bigcup \mathcal{B}_0\right) = \bigcup_{b \in \text{CH}(\mathcal{B}_0)} m_b(N_{\mathbf{1}}(c(b), \text{CH}(c(\mathcal{B}_0))))$$

Remark 35. For any ball $b \in \text{CH}(\mathcal{B}_0)$, there always exists a vertex ball b' of in $\text{CH}(\mathcal{B}_0)$ such that $N(c(b), \text{CH}(c(\mathcal{B}_0))) \subseteq N(c(b'), \text{CH}(c(\mathcal{B}_0)))$. Therefore we actually have

$$\partial\left(\bigcup \mathcal{B}_0\right) = \bigcup_{b \in \mathcal{B}_0} m_b(N_{\mathbf{1}}(c(b), \text{CH}(c(\mathcal{B}_0))))$$

It follows that the boundary of $\bigcup \mathcal{B}_0$ can be described exactly with the images of m_b . Because the function m_b is built from the normal cones to $\text{CH}(c(\mathcal{B}_0))$, these images inherit many structural properties of normal cones as shown above. We shall thus rely on these properties in the upcoming results that involve the boundary of $\bigcup \mathcal{B}_0$.

2.2.3 Proofs of results

We are now ready to prove the two Lemmas 2.19 and 2.20. We start with the parabolic case.

Proof of Lemma 2.20. By Property 2.18, if $b \in \text{CH}(\mathcal{B}_0)$ we indeed have $b \subseteq \bigcup \mathcal{B}$. Conversely, let $b \in \mathcal{P}$ be a ball such that $b \subseteq \bigcup \mathcal{B}$. We show that $b \in \text{CH}(\mathcal{B}_0)$.

By contradiction, suppose that $b \notin \text{CH}(\mathcal{B}_0)$. Consider $\mathcal{B}' = \mathcal{B}_0 \cup \{b\}$. Then b is a vertex of $\text{CH}(\mathcal{B}')$ and its center c is a vertex of $\text{CH}(c(\mathcal{B}'))$. Hence c has an exclusive unit normal vector \mathbf{e} to the convex polytope $\text{CH}(c(\mathcal{B}'))$. Because $c_0 \in \text{CH}(c(\mathcal{B}'))$, we are guaranteed that $\rho_b(\mathbf{e}) > 0$. Indeed, because \mathbf{e} is exclusive to c we have $\langle \mathbf{e}, c_0 - c \rangle < 0$, which implies $2\langle \mathbf{e}, c - c_0 \rangle = \rho_b(\mathbf{e}) > 0$. Consider $m_b(\mathbf{e}) = c_0 + \rho_b(\mathbf{e})\mathbf{e}$, we have $m_b(\mathbf{e}) \neq c_0$. Because \mathbf{e} only belongs to the normal cone of b , Property 2.35 implies that b is the only ball in $\text{CH}(\mathcal{B}')$ that contains the point $m_b(\mathbf{e})$. This contradicts the assumption $b \subseteq \bigcup \mathcal{B}$ and is thus impossible. \square

Proof of Lemma 2.19. The equivalence (i) \iff (ii) is mostly identical to the parabolic case, the only change is that we obtain $\rho_b(\mathbf{e}) > 0$ because $r_0 > 0$. The implication (ii) \implies (iii) being obvious, we here prove that (iii) \implies (i).

Assume that (iii) holds and by contradiction suppose that $b \notin \text{CH}(\mathcal{B})$. As previously consider the collection of balls $\mathcal{B}' = \mathcal{B} \cup \{b\}$. Ball b is again a vertex in $\text{CH}(\mathcal{B}')$, and its center c a vertex in $\text{CH}(c(\mathcal{B}'))$. It has an exclusive outward normal vector \mathbf{e} . Let $u \in s_0(\text{aff}(\mathcal{B}'))$ a point of the radical sphere, and \mathbf{n} the unit direction from c_0 to u . Note that because u belongs to the radical sphere of pencil $\text{aff}(\mathcal{B}')$, necessarily $\mathbf{n} \in \vec{V}$, with $V = V(\text{aff}(\mathcal{B}'))$ the radical space of the pencil. In particular, $\mathbb{R}\mathbf{n} \subseteq \vec{V}$. We deduce that \mathbf{e} and \mathbf{n} cannot be collinear. For $0 \leq \lambda \leq 1$, let $\mathbf{e}_\lambda = \lambda\mathbf{e} + (1 - \lambda)\mathbf{n}$. For all $c' \in \text{CH}(c(\mathcal{B}'))$ we have

$$\langle \mathbf{e}_\lambda, c' - c \rangle = \lambda\langle \mathbf{e}, c' - c \rangle + (1 - \lambda)\langle \mathbf{n}, c' - c \rangle = \lambda\langle \mathbf{e}, c' - c \rangle \leq 0$$

with equality if and only if $\lambda = 0$, which occurs when \mathbf{e}_λ coincides with \mathbf{n} . Therefore, when $\lambda > 0$, \mathbf{e}_λ is an exclusive outward normal vector for vertex c . Since \mathbf{e} and \mathbf{n} are not collinear, the \mathbf{e}_λ 's are non-zero vectors. We can consider their normalization, and subsequently their image by function m_b . Because the \mathbf{e}_λ 's form a path in the unit vector sphere, and m_b is continuous, the images of their normalization is also a path γ . Because the \mathbf{e}_λ , $\lambda > 0$, are exclusive normals at c , the set of points $\gamma \setminus \{m_b(\mathbf{n})\}$ is exclusively covered by ball b in $\bigcup \mathcal{B}$. However, $m_b(\mathbf{n}) = u$, and γ is a path containing u . It follows that in any neighbourhood of u , ball b exclusively covers several points of $\bigcup \mathcal{B}$. This contradicts (iii) and is absurd. \square

2.3 Local inclusion in parabolic pencils

In the previous section, we have seen that the local inclusion of a ball b in $\bigcup \mathcal{B}$ is equivalent to its global inclusion whenever \mathcal{B} is a finite collection of balls in \mathbb{R}^d , and both \mathcal{B} and b belong to some pencil whose radical sphere has positive radius. In other words, $\text{aff}(\mathcal{B} \cup \{b\})$ must span an elliptic pencil of dimension at most $d - 1$.

In this section, we have a look at local inclusion properties, but relax the constraints on \mathcal{B} and b . We assume \mathcal{B} to be a finite collection of balls whose boundaries share a common point u . Recall as per Remark 15 that \mathcal{B} thus spans a pencil of balls, which can always be seen as a (parabolic or elliptic) subpencil of a parabolic pencil with radical (singleton) sphere $\{u\}$.

We show that a ball b going through u satisfies local inclusion in $\bigcup \mathcal{B}$ around u , if and only if the center of b is constrained to belong to a cone defined by u and the centers of balls in \mathcal{B} . In other words, we provide a characterization of the belonging of ball b in $\bigcup \mathcal{B}$ locally around u . To define the cone, we need some definitions.

Let $C = \{c_1, \dots, c_k\}$ be a finite collection of points in \mathbb{R}^d . Recall that the affine conical hull of C with apex $u \in \mathbb{R}^d$ is

$$\text{coni}(u, C) = \left\{ u + \sum_{i=1}^k \lambda_i (c_i - u) \mid \forall i, \lambda_i \in \mathbb{R}, \lambda_i \geq 0 \right\}$$

For technical reasons, we need to exclude some points on the boundary of these affine conical hulls. This motivates the following definition:

$$\text{coni}^*(u, C) = \overset{\circ}{\text{coni}}(u, C) \cup \text{CH}(\{u\} \cup C)$$

Proposition 2.40. *Let $u, c_1, \dots, c_k, c_{k+1} \in \mathbb{R}^d$ and $C = \{c_1, \dots, c_k\}$. Let b_i be the d -ball centered at c_i whose bounding $(d - 1)$ -sphere goes through u . We have equivalence between:*

(i) $c_{k+1} \in \text{coni}^*(u, C)$

(ii) $\exists U$ a neighbourhood of u such that $b_{k+1} \cap U \subseteq \bigcup_{i=1}^k b_i$




Figure 2.12 – Illustration of Proposition 2.40 for $k = 2$. The inclusion normal cone $\text{coni}(u, C)^*$ is represented in orange. Note that the two green balls have their center in the inclusion normal cone, and indeed satisfy local inclusion in the blue shape around u . The black ball covers the red area and does not satisfy local inclusion, thus showing why we have to remove some points on the boundary of the conical hull.

Due to the equivalence above, we refer to the set of points $\text{coni}(u, C)^*$ as the *inclusion normal cone of u in $\bigcup_{i=1}^k b_i$* . Refer to Figure 2.12 for an illustration. This inclusion normal cone nearly coincides with the normal cone of the complement of $\bigcup_{i=1}^k b_i$ at u . Recall that intuitively, the normal cone of a subset X at a point $x \in \partial X$ is the set of directions along which distance to X increases the fastest if x were moved. In other words for an arbitrarily small distance ρ and a direction \mathbf{e} in the normal cone, the point $x + \rho\mathbf{e}$ admits x as unique closest point in $\overline{X^c}$. This implies that the ball centered at $x + \rho\mathbf{e}$ and radius ρ is locally included in $\overline{X^c}$ around x . The main distinction with our inclusion normal cone is that the normal cone of the complement gives a superset of the directions of centers of balls that satisfy local inclusion, whereas our inclusion normal cone yields the actual centers of these balls.

The remainder of this section is dedicated to the proof of the proposition above.

2.3.1 Technical lemmas

In this section we present two technical lemmas. Let $u \in \mathbb{R}^d$ and consider a family of balls whose bounding $(d - 1)$ -sphere goes through u . The balls in the family defines a subpencil to a parabolic d -pencil, as per Remark 15. Given a finite collection of balls in that family, the first lemma states that when one center c_{k+1} lies in the interior of the conical hull formed by the other centers, then the ball centered at c_{k+1} is locally contained around u in

the union of the other balls. The second lemma is a weaker version of the converse.

Lemma 2.41. *Let $u \in \mathbb{R}^d$ and k points $c_1, c_2, \dots, c_k \in \mathbb{R}^d$ not all lying on the same affine hyperplane through u . Let $c_{k+1} \in \text{coni}(u, \{c_i, 1 \leq i \leq k\})$. For $1 \leq i \leq k+1$, let b_i be the ball centered at c_i whose bounding $(d-1)$ -sphere goes through u . Then, there exists $r > 0$ such that $b_{k+1} \cap b(u, r) \subseteq \bigcup_{i=1}^k b_i$.*




Figure 2.13 – Notation for the proof of Lemma 2.41 with $k = 3$. The bold curve represents the set of points at distance $R_{\mathbf{e}}$ from u in directions \mathbf{e} such that $\langle \mathbf{e}, c - u \rangle \geq 0$.

Proof. Consider the closed half-space H through u containing b_{k+1} . We prove the result by establishing that $H \cap b(u, r) \subseteq \bigcup_{i=1}^k b_i$.

By Lemma 2.27, notice that for all unit vectors \mathbf{e} and real numbers ρ , the point $u + \rho \mathbf{e}$ belongs to b_i if and only if $0 \leq \rho \leq 2\langle \mathbf{e}, c_i - u \rangle$. It follows that $u + \rho \mathbf{e}$ belongs to the union $\bigcup_{i=1}^k b_i$ if and only if $0 \leq \rho \leq R_{\mathbf{e}}$ where $R_{\mathbf{e}} = 2 \max_{1 \leq i \leq k} \langle \mathbf{e}, c_i - u \rangle$; see Figure 2.13.

Let $r = \inf\{R_{\mathbf{e}} \mid \langle \mathbf{e}, c_{k+1} - u \rangle \geq 0, \|\mathbf{e}\| = 1\}$. Note that r is an infimum over the unit directions \mathbf{e} , with $u + \mathbf{e} \in H$. We next prove that $r > 0$. Because the closed half unit sphere is compact, r is actually a minimum. Consider the unit vector \mathbf{e} for which $r = R_{\mathbf{e}}$. It satisfies $\langle \mathbf{e}, c_{k+1} - u \rangle \geq 0$. Because c_{k+1} lies in the interior of $\text{coni}(u, \{c_i, 1 \leq i \leq k\})$, there exist k real positive numbers $\lambda_1, \lambda_2, \dots, \lambda_k$ such that $c_{k+1} = u + \sum_{i=1}^k \lambda_i (c_i - u)$; see Appendix A.2. By linearity,

$$0 \leq \sum_{i=1}^k \lambda_i \langle \mathbf{e}, c_i - u \rangle \leq \left(\sum_{i=1}^k \lambda_i \right) \max_{1 \leq i \leq k} \langle \mathbf{e}, c_i - u \rangle.$$

Assume for a contradiction that $\max_{1 \leq i \leq k} \langle \mathbf{e}, c_i - u \rangle \leq 0$. Since $\lambda_i > 0$ for all $1 \leq i \leq k$, this would imply that $\langle \mathbf{e}, c_i - u \rangle = 0$ for all $1 \leq i \leq k$, showing that all vectors $c_i - u$ lie in the same hyperplane and contradicting our hypothesis.

To conclude, consider a point $x \in H \cap b(u, r)$. Let $\rho = \|x - u\|$ and \mathbf{e} the unit vector such that $x = u + \rho\mathbf{e}$. Because $x \in H$, $\langle \mathbf{e}, c_{k+1} - u \rangle \geq 0$. Because $x \in b(u, r)$, $0 \leq \rho \leq r \leq R_{\mathbf{e}}$. Hence, $x \in \bigcup_{i=1}^k b_i$ and $H \cap b(u, r) \subseteq \bigcup_{i=1}^k b_i$. \square

Lemma 2.42. *Let $u \in \mathbb{R}^d$ and $k + 1$ points $c_1, \dots, c_k, c_{k+1} \in \mathbb{R}^d \setminus \{u\}$. For $1 \leq i \leq k + 1$, let b_i be the ball centered at c_i whose bounding $(d - 1)$ -sphere goes through u . If there exists $r > 0$ such that $b_{k+1} \cap b(u, r) \subseteq \bigcup_{i=1}^k b_i$, then $c_{k+1} \in \text{coni}(u, \{c_i, 1 \leq i \leq k\})$.*

Proof. Denote by b_0 the singleton ball $\{u\}$ and $C = \{c_i \mid 1 \leq i \leq k\}$. Let $\mathcal{B}_0 = \{b_i \mid 0 \leq i \leq k\}$ and notice that $\bigcup_{i=1}^k b_i = \bigcup \mathcal{B}_0$. The interiors of b_{k+1} and $b(u, r)$ intersect, thus there is a ball $\hat{b} \subseteq b_{k+1} \cap b(u, r)$ with center \hat{c} on the half-line from u through c_{k+1} and such that $u \in \partial \hat{b}$. We show that $\hat{b} \in \text{CH}(\mathcal{B}_0)$, this implies that $\hat{c} \in \text{CH}(\{u\} \cup C) \subseteq \text{coni}(u, C)$. Hence we obtain $c_{k+1} \in \text{coni}(u, \{\hat{c}\}) \subseteq \text{coni}(u, C)$.

Because $b_0 = \{u\} \in \mathcal{B}_0$, $\text{aff}(\mathcal{B}_0)$ is parabolic with radical ball b_0 . To prove that $\hat{b} \in \text{CH}(\mathcal{B}_0)$, it suffices to prove that $\hat{b} \in \text{aff}(\mathcal{B}_0)$. Indeed by definition $\hat{b} \subseteq b_{k+1} \cap b(u, r) \subseteq \bigcup \mathcal{B}_0$. We can thus apply Lemma 2.20. By contradiction, assume that $\hat{b} \notin \text{aff}(\mathcal{B}_0)$. Because the radical sphere $\{u\}$ of $\text{aff}(\mathcal{B}_0)$ is in the boundary of \hat{b} , we know that \hat{b} belongs to the pencil $\text{aff}(\mathcal{B}_0)$ if and only if its center \hat{c} belongs to the centers' space $c(\text{aff}(\mathcal{B}_0))$. Our assumption thus implies that $\hat{c} \notin c(\text{aff}(\mathcal{B}_0))$. Therefore, there are $c \in c(\text{aff}(\mathcal{B}_0))$, $h > 0$ and $\mathbf{e} \in \vec{V}(\text{aff}(\mathcal{B}_0))$, $\|\mathbf{e}\| = 1$, such that $\hat{c} = c + h\mathbf{e}$. Because $\mathbf{e} \in \vec{V}(\text{aff}(\mathcal{B}_0))$, for all $1 \leq i \leq k$, $\langle \mathbf{e}, c_i - u \rangle = 0$. By Lemma 2.27, this implies that for all $\rho > 0$, we have $u + \rho\mathbf{e} \notin \bigcup_{i=1}^k b_i = \bigcup \mathcal{B}_0$. However $\langle \mathbf{e}, \hat{c} - u \rangle = h\|\mathbf{e}\|^2 > 0$. This implies $2\langle \mathbf{e}, \hat{c} - u \rangle = \rho_{\hat{b}}(\mathbf{e}) > 0$. Therefore Lemma 2.27 implies that for any $\rho > 0$ small enough, we have $u + \rho\mathbf{e} \in \hat{b} \setminus (\bigcup \mathcal{B}_0) = \emptyset$, which is impossible. Hence, $\hat{b} \in \mathcal{P}$ and $c_{k+1} \in \text{coni}(u, C)$. \square

2.3.2 Proving the property of local inclusion cones

We can finally prove Proposition 2.40.

Proof of Proposition 2.40. Let $b_0 = \{u\}$ and $\mathcal{B}_0 = \{b_i \mid 0 \leq i \leq k\}$. The affine hull $\text{aff}(\mathcal{B}_0)$ is a parabolic pencil. Its centers' space $c(\text{aff}(\mathcal{B}_0))$ coincides with $\text{aff}(\{u\} \cup C)$.

Assume that $c_{k+1} \in \overset{\star}{\text{coni}}(u, C)$. If $c_{k+1} \in \text{CH}(\{u\} \cup C)$ we deduce by Property 2.18 that $b_{k+1} \subseteq \bigcup \mathcal{B}_0$. If $c_{k+1} \in \overset{\circ}{\text{coni}}(u, C)$ we instead rely on Lemma 2.41 and thus b_{k+1} is locally contained in $\bigcup \mathcal{B}_0$ locally around u .

Conversely assume that b_{k+1} is locally included in $\bigcup \mathcal{B}_0$ around u . By Lemma 2.42 we know that $c_{k+1} \in \text{coni}(u, C)$. If $c_{k+1} \in \overset{\circ}{\text{coni}}(u, C)$, then we already have $c_{k+1} \in \overset{\star}{\text{coni}}(u, C)$.

Otherwise, $c_{k+1} \in \partial \text{coni}(u, C)$. We show that then, c_{k+1} must belong to the convex hull of $\{u\} \cup C$.

The boundary of $\text{coni}(u, C)$ consists of portions of affine subspaces that are supported by u and points of C . It follows that for any face f_{coni} of $\text{coni}(u, C)$, there is some face f_{ch} of $\text{CH}(\{u\} \cup C)$ such that $\text{coni}(u, f_{\text{ch}}) = f_{\text{coni}}$. Without loss of generality we can assume that f_{ch} contains u . Indeed, $\text{CH}(\{u\} \cup f_{\text{ch}})$ must be a face of $\text{CH}(\{u\} \cup C)$, otherwise $\text{coni}(u, f_{\text{ch}})$ cannot be a face of $\text{coni}(u, C)$. Below, f_{coni} denotes a face of $\text{coni}(u, C)$ that contains c_{k+1} and f_{ch} its associated face of $\text{CH}(\{u\} \cup C)$.

Assume by contradiction that c_{k+1} does not belong to $\text{CH}(\{u\} \cup C)$. It follows that c_{k+1} is a vertex in $\text{CH}(\{u, c_{k+1}\} \cup C)$. Therefore b_{k+1} is also a vertex in $\text{CH}(\mathcal{B}_0 \cup \{b_{k+1}\})$. Let \mathbf{e} be an exclusive outward normal vector of c_{k+1} , and \mathbf{n} an outward normal vector of face f_{ch} . Since f_{ch} contains both u and c_{k+1} , \mathbf{n} is a common outward normal to u and c_{k+1} . In particular, \mathbf{e} and \mathbf{n} cannot be collinear. For $0 \leq \lambda \leq 1$, let $\mathbf{e}_\lambda = \lambda \mathbf{e} + (1 - \lambda) \mathbf{n}$. Proceeding as in the proof of Lemma 2.19, we can show that when $\lambda > 0$, \mathbf{e}_λ is an exclusive, non-zero, outward normal vector to c_{k+1} . We can consider the normalization of the \mathbf{e}_λ 's, they constitute a continuous path of unit vectors in the unit normal cone to $\text{CH}(\{u, c_{k+1}\} \cup C)$ at point c_{k+1} . By continuity of $m_{b_{k+1}}$, their image by $m_{b_{k+1}}$ forms a continuous path in $\partial b_{k+1} \cap \partial(b_{k+1} \cup (\bigcup \mathcal{B}_0))$. In addition, recall that \mathbf{n} is a normal vector to $\text{CH}(\{u, c_{k+1}\} \cup C)$ at both c_{k+1} and u . Therefore, $m_{b_{k+1}}(\mathbf{n}) = m_{b_0}(\mathbf{n}) = u$. Because the \mathbf{e}_λ 's, for $\lambda > 0$, are normal vectors exclusive to c_{k+1} , the points $m_{b_{k+1}}(\mathbf{e}_\lambda)$, $\lambda > 0$, are exclusively covered by b_{k+1} among all balls of the union $b_{k+1} \cup (\bigcup \mathcal{B}_0)$. This contradicts the local inclusion assumption and is thus impossible. \square

2.4 Boundary of a finite union of balls

Given a finite set of balls \mathcal{S} in \mathbb{R}^d , we decompose the boundary of the union $\bigcup \mathcal{S}$ into a finite set of k -dimensional elements that we call k -faces. Intuitively, k -faces are maximal open connected spherical patches of dimension k that partition the boundary. To define them formally, we associate to each point $u \in \mathbb{R}^d$ the subsets of balls that have the property to represent $\bigcup \mathcal{S}$ locally around u (see Definition 9 below). From these local representations, we then define the degree of a boundary point and rely on this to distinguish the different faces of $\partial(\bigcup \mathcal{S})$. When assuming balls in \mathcal{S} to be in general position, defining the degree of a boundary point is easy: this is the number of balls in \mathcal{S} containing the boundary point. We however require a definition that still holds when degenerate situations occur. Indeed, in the later Chapter 4, we will require collections of balls that do not satisfy the general position assumption. It is thus desirable to have a terminology that remains relevant within our own framework.

Definition 9 (Local representation). Given a subset $X \subseteq \mathbb{R}^d$ and a point $u \in \mathbb{R}^d$, we say that a set of balls \mathcal{B} in \mathbb{R}^d represents X locally around u if:

- (i) for all $b \in \mathcal{B}$, $u \in b$;

(ii) there exists $r > 0$ such that $X \cap b(u, r) = \bigcup \mathcal{B} \cap b(u, r)$.

For an arbitrary X and u , such a set \mathcal{B} may be infinite or even not exist. For instance, take for X the region enclosed by an ellipse and for u any of its boundary points: a collection of balls that satisfies condition (ii) requires infinitely many balls that do not contain the point being locally described, and thus infringes (i); see Figure 2.14. We first establish that such a local representation always exists and is finite when X is the union of a finite collection of balls.




Figure 2.14 – Consider the point u on the boundary of the ellipse, it has no local representation. Because the boundary is smooth, balls of a local representation have to be tangent to the ellipse at point u , they must thus have their center along the blue arrow. These tangent balls are a nested family of balls, any ball whose center is beyond c meets the complement of the ellipse around u , and balls whose center is between u and c are contained in the red ball. However the red ball does not capture the ellipse in the neighbourhood of u . A “local representation” thus requires balls that do not contain u .

Recall that \mathcal{S} is a finite collection of d -balls in \mathbb{R}^d . We have the property below.

Property 2.43. *The set $\mathcal{S}_u = \{b \in \mathcal{S} \mid u \in b\}$ represents $\bigcup \mathcal{S}$ locally around u , for all $u \in \mathbb{R}^d$.*

Proof. For any ball $b \in \mathcal{S}$, either $u \in b$ or the distance between u and b is positive. Because \mathcal{S} is finite, $\mathcal{S} \setminus \mathcal{S}_u$ is also finite and we let r be the smallest distance between u and a ball in $\mathcal{S} \setminus \mathcal{S}_u$. We have $r > 0$, and by construction $b(u, r) \cap \bigcup \mathcal{S} \subseteq \bigcup \mathcal{S} \setminus \bigcup (\mathcal{S} \setminus \mathcal{S}_u) \subseteq \bigcup \mathcal{S}_u \subseteq \bigcup \mathcal{S}$. Thus $\bigcup \mathcal{S} \cap b(u, r) = \bigcup \mathcal{S}_u \cap b(u, r)$. \square

Let $S = \bigcup \mathcal{S}$ be the union of a finite collection of balls \mathcal{S} in \mathbb{R}^d . For any point $u \in \partial S$, we are now ready to define the *degree of u in S* . Consider \mathcal{B} a set of balls that represents S locally around u . Notice that for all $b \in \mathcal{B}$, u is on the boundary of b . Hence the intersection of the bounding $(d - 1)$ -spheres of \mathcal{B} is non-empty and \mathcal{B} defines a pencil of balls \mathcal{P} by Property 2.2.

Definition 10 (Degree of boundary points). Let S be a finite union of balls, $u \in \partial S$, and \mathcal{B} that represents S locally around u . Assume that \mathcal{B} achieves minimum cardinality among the sets that represents S locally around u . We define the *degree* of u in S to be $\deg u = d - 1 - \dim \text{aff}(\mathcal{B})$ when $\text{aff}(\mathcal{B})$ is elliptic, and $\deg u = 0$ when $\text{aff}(\mathcal{B})$ is parabolic.

Remark 36. If we consider all singleton spheres to have dimension 0, the degree of a point is in fact the dimension of the radical sphere of $\text{aff}(\mathcal{B})$.

It is not obvious that such a definition is consistent, because there may be several minimal sets \mathcal{B} , that may define distinct pencils with different dimensions. We show through Property 2.44 that this cannot occur. First, notice that regardless of which set locally represents S , they must all define the same inclusion normal cone for u in S . The local inclusion property is indeed a shape property. Consider two collections of balls \mathcal{B} and \mathcal{B}' that locally represent S around u . Recall that their respective collections of centers are $c(\mathcal{B})$ and $c(\mathcal{B}')$. Then one easily proves the equality $\text{coni}^*(u, c(\mathcal{B})) = \text{coni}^*(u, c(\mathcal{B}'))$. The inclusion normal cone is thus an invariant across the various local representations, and we will use the notation $\text{coni}^*(u, S)$.

Property 2.44. Consider $X \subseteq \mathbb{R}^d$ and $u \in \partial X$. Let $\mathcal{P}(u, X)$ be the collection of pencils $\text{aff}(\mathcal{B})$ where \mathcal{B} is a finite set of balls that locally represents X around u . If it is non-empty, then $\mathcal{P}(u, X)$ admits a unique inclusion-wise minimum pencil, defined by a cardinal-wise minimal collection \mathcal{B} .

Proof. For any finite local representation \mathcal{B} of X around u , $\text{coni}^*(u, c(\mathcal{B})) = \text{coni}^*(u, X)$. We distinguish three cases:

(a) There is some collection \mathcal{B} such that $\dim(\text{aff}(\{u\} \cup c(\mathcal{B}))) \leq d-1$. Thus $\text{coni}^*(u, c(\mathcal{B})) = \emptyset$ and it follows that $\text{coni}^*(u, X) = \text{CH}(\{u\} \cup c(\mathcal{B}))$. Let C_{\min} be the vertices of $\text{CH}(c(\mathcal{B}))$ and \mathcal{B}_{\min} the set of balls centered at points of C_{\min} whose bounding $(d-1)$ -spheres go through u . We have $\text{coni}^*(u, C_{\min}) = \text{coni}^*(u, X)$, hence \mathcal{B}_{\min} defines a pencil $\text{aff}(\mathcal{B}_{\min}) = \mathcal{P}_{\min} \in \mathcal{P}(u, X)$. Any set \mathcal{B} that locally describes X around u must contain \mathcal{B}_{\min} , otherwise they cannot satisfy $\text{coni}^*(u, c(\mathcal{B})) = \text{coni}^*(u, X)$. Thus, \mathcal{B}_{\min} achieves minimum cardinality. Therefore, for all $\mathcal{P} \in \mathcal{P}(u, X)$ we have $\mathcal{P}_{\min} \subseteq \mathcal{P}$.

(b) Otherwise, for all local representation \mathcal{B} , $\dim(\text{aff}(\{u\} \cup c(\mathcal{B}))) = d$, which implies that for all \mathcal{B} , $\dim(\text{aff}(c(\mathcal{B}))) \geq d-1$, and $\text{coni}^*(u, c(\mathcal{B})) \neq \emptyset$. We have two subcases.

(i) Suppose there is some collection \mathcal{B} such that $\dim(\text{aff}(c(\mathcal{B}))) = d-1$. Necessarily, u does not belong to the affine space supported by $c(\mathcal{B})$, hence $\text{aff}(\mathcal{B})$ is an elliptic pencil of dimension $d-1$. As in case (a), consider C_{\min} and \mathcal{B}_{\min} . Any collection of balls defining a pencil in $\mathcal{P}(u, X)$ must contain every ball in \mathcal{B}_{\min} , otherwise we would have the strict inclusion $\text{coni}^*(u, c(\mathcal{B})) \subsetneq \text{coni}^*(u, X)$. Let $\mathcal{P}_{\min} = \text{aff}(\mathcal{B}_{\min})$.

If \mathcal{P}_{\min} belongs to $\mathcal{P}(u, X)$ then it is inclusion-wise minimum in $\mathcal{P}(u, X)$. However, contrary to the previous situation (a), it is not straightforward to see that $\text{coni}^*(u, C_{\min}) =$

$\text{coni}^*(u, X)$, because the component $\text{coni}^*(u, C_{\min})$ is now non-empty. But we can still show that indeed, \mathcal{P}_{\min} belongs to $\mathcal{P}(u, X)$. To do so, we show that $\mathcal{P}_{\min} = \text{aff}(\mathcal{B})$. Because $\mathcal{B}_{\min} \subseteq \mathcal{B}$, we have $\mathcal{P}_{\min} \subseteq \text{aff}(\mathcal{B})$. Additionally, we have $\text{aff}(C_{\min}) = \text{aff}(c(\mathcal{B}))$, therefore the two pencils \mathcal{P}_{\min} and $\text{aff}(\mathcal{B})$ must have the same dimension $d - 1$. It follows that $\mathcal{P}_{\min} = \text{aff}(\mathcal{B})$ and therefore \mathcal{P}_{\min} is the unique inclusion minimum pencil of $\mathcal{P}(u, X)$.

- (ii) For all collection \mathcal{B} , $\dim(\text{aff}(c(\mathcal{B}))) = d$. Any such collection defines a parabolic d -pencil that admits $\{u\}$ as its radical ball. Hence $\mathcal{P}(u, X)$ contains only one pencil. \square

In the above proof, the two cases (a) and (b)-(i) have an explicit construction of the (only) cardinal minimum collection \mathcal{B}_{\min} that locally describes a shape around any boundary point u . In the remaining case (b)-(ii), although we could still build a similar collection from the vertices of $\text{CH}(\{u\} \cup c(\mathcal{B}))$, we would have no guarantee that it actually achieves minimum cardinality. Indeed, assume that u is not a vertex of $\text{CH}(\{u\} \cup c(\mathcal{B}))$ but belongs to a $(d - 1)$ -face of the polytope. The affine conical hull supported by the vertices of that face is an affine hyperplane through u . Because $\dim(\text{aff}(c(\mathcal{B}))) = d$ there exists another center lying in one of the open half-spaces delimited by that hyperplane. Because u belongs to the boundary, the conical hull cannot cover the whole space \mathbb{R}^d , hence $\text{coni}^*(u, c(\mathcal{B}))$ is a half-space. In particular, any single center in that open half-space suffices to fully determine $\text{coni}^*(u, X)$, regardless of whether the center is a vertex or not, or how many vertices lie in the half-space. See Figure 2.15 for an illustration.

Property 2.44 guarantees that every point of the boundary has a well-defined degree. We then call a k -face of ∂S any maximal connected component of $\partial_k S = \{u \in \partial S \mid \deg u = k\}$. In particular we refer to 0-faces as *vertices* of the boundary, and denote the vertex set by \mathcal{V} . Vertices come in many varieties.

The two cases (b)-(i) and (b)-(ii) of the proof above only occur when u is a vertex of the boundary. A common point between these two situations is that the inclusion normal cone of S at u has non-empty interior, we call these vertices *simple*. In the remaining case (a), the inclusion normal cone is always flat with empty interior. Either the pencil is elliptic, and the boundary point u belongs to some open k -face, $k \geq 1$, or the pencil is parabolic, and the boundary point u is a vertex. In the vertex case, we say that u is a *non-simple* vertex.

Note that the general position assumption usually means that non-simple vertices from case (a) and simple vertices from case (b)-(ii) are both forbidden. Thus the only vertices allowed in general position are the simple vertices from case (b)-(i). The classification of simple and non-simple vertices thus deviates from the usual “regular” and “degenerate” classification of vertices from the general position assumption. Refer to Figure 2.16 for an illustration. For easy reference, we re-state the definitions of simple and non-simple vertices.

Definition 11 (Simple and non-simple vertices of the boundary). Let $S = \bigcup \mathcal{S}$ be a finite union of balls in \mathbb{R}^d and $u \in \partial S$ a vertex. Let \mathcal{B} be a cardinal-minimum local representation of S around u . We say that u is a *simple* vertex if $\text{coni}^*(u, c(\mathcal{B})) \neq \emptyset$, otherwise we say u is a *non-simple* vertex.

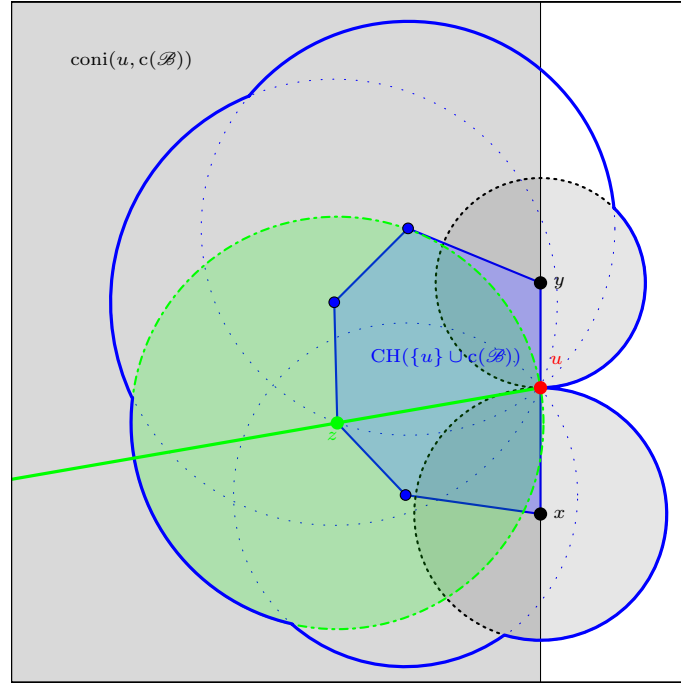


Figure 2.15 – The inclusion normal cone of u is the left half-space. Any single point in the interior of that half-space, in addition to x and y (black points), are sufficient to determine the interior normal cone. For instance we can take point z (green). We have $\text{con}_i^*(u, c(\mathcal{B})) = \text{con}_i^*(u, \{x, y, z\})$. Notice that the transparent green and black balls indeed suffice to locally describe the overall union of balls around u .

We now show that the other k -faces, $k \geq 1$, are open connected patches of k -spheres.

Property 2.45. Consider $u \in \partial_k S$, $k \geq 1$, and let \mathcal{B} be the unique cardinal-wise minimum collection of balls that represents S locally around u . Then, there exists $r > 0$ such that \mathcal{B} is the unique cardinal-wise minimum collection of balls that represents S locally around v for all $v \in \mathring{b}(u, r) \cap \partial_k S$.

Proof. Let $r > 0$ such that $S \cap b(u, r) = (\bigcup \mathcal{B}) \cap b(u, r)$. Let $\rho_v = r - \|u - v\|$ and observe that $\mathring{b}(v, \rho_v) \subseteq b(u, r)$. For all $v \in \mathring{b}(u, r) \cap S$, we thus have that $S \cap \mathring{b}(v, \rho_v) = (\bigcup \mathcal{B}) \cap \mathring{b}(v, \rho_v)$. By restricting \mathcal{B} to the set of balls that contain v , $\mathcal{B}_v = \{b \in \mathcal{B} \mid v \in b\}$, it follows that for all points $v \in \mathring{b}(u, r) \cap S$, \mathcal{B}_v represents S locally around v . Let \mathcal{P}_v be the pencil defined by \mathcal{B}_v . In particular, since $v \in \mathring{b}(u, r) \cap \partial_k S$, we have $\dim \mathcal{P}_v = d - 1 - k = \dim \mathcal{P}$ by definition of the degree. Since $\mathcal{B}_v \subseteq \mathcal{B}$, $\mathcal{P}_v \subseteq \mathcal{P}$ and we actually have $\mathcal{P}_v = \mathcal{P}$, therefore $\mathcal{B}_v = \mathcal{B}$. Thus \mathcal{B} is indeed a local representation of S around v . In addition, case (a) in the proof of Property 2.44 ensures minimality and uniqueness of \mathcal{B} when $k \geq 1$. \square

We deduce immediately that any k -face is supported by a unique k -sphere in \mathcal{S} . Formally:

Property 2.46. Let f be a k -face of ∂S , $k \geq 1$. Then there is a unique set $\mathcal{B} \subseteq \mathcal{S}$ such




Figure 2.16 – The three possible configurations of a vertex u in \mathbb{R}^2 , corresponding to the three cases in the proof of Property 2.44. The general position assumption forbids the two cases (a) and (b)-(ii).

that (1) $s = \bigcap_{b \in \mathcal{B}} \partial b$ is a non-degenerate k -sphere; (2) $f \subseteq s$; (3) \mathcal{B} is the cardinal minimum representation of S locally around each point of f .

Proof. Consider $g : \partial_k S \rightarrow \{\mathcal{B} \mid \mathcal{B} \subseteq \mathcal{S}\}$ that associates to each point $x \in \partial_k S$ the unique inclusion-minimum set that represents S locally around x . By Property 2.45, f is locally constant and therefore constant over each connected component of $\partial_k S$. \square

As a concluding remark, note that there are finitely many faces of any dimension. We outline a possible proof by induction below.

Consider a vertex u of ∂S , there exists some subset $\mathcal{B} \subseteq \mathcal{S}$ such that $\text{aff}(\mathcal{B})$ is either an elliptic $(d-1)$ -pencil, or a parabolic pencil, and u belongs to the radical sphere of that pencil. In both cases, the radical sphere contains finitely many points, because it is either a singleton, or a 0-sphere with 2 points. Because \mathcal{S} is finite, there are finitely many possible pencils that can yield a vertex of ∂S , therefore there are finitely many vertices.

Let $k \geq 1$ and suppose that there are finitely many faces of ∂S with dimension at most $k-1$. We claim there are finitely many k -faces of ∂S . By the above properties, for any open k -face f of ∂S , there is some subset $\mathcal{B} \subseteq \mathcal{S}$ such that $\text{aff}(\mathcal{B})$ is an elliptic $(d-k-1)$ -pencil and $f \subseteq s_0(\text{aff}(\mathcal{B}))$. Recall that we defined $\partial_k S = \{u \in \partial S \mid \deg u = k\}$, and that f is a maximal connected component of $\partial_k S$. Because the relative boundary of f is a collection of faces of ∂S of dimension at most $k-1$, we actually have that f is a maximal connected component of $s_0(\text{aff}(\mathcal{B})) \setminus \left(\bigcup_{\ell \leq k-1} \partial_\ell S\right)$. Because there are finitely many possible radical k -spheres $s_0(\text{aff}(\mathcal{B}))$, and finitely many faces of dimension $\ell \leq k-1$, there are only finitely many k -faces.

By induction, we thus deduce that S has finitely many faces of any dimension.

2.5 Medial axis, skeleton, boundary and pencils

An important result regarding unions of balls concerns the structure of their medial axis. Unions of balls are in fact one of the few classes of shapes for which we can exactly compute the medial axis. A characterization suited for an algorithmic implementation was given by Amenta and Kolluri [AK01]. It was inspired from a previous characterization from Attali and Montanvert [AM97] and the so called dual shape of a union of balls from Edelsbrunner [Ede92]. One caveat of these characterizations is that they assume general position.

Although relaxing this assumption makes the previous characterizations invalid, some properties remain true. Indeed from Amenta and Kolluri's result (that we recall in the following Section 2.5.1), we easily deduce that in general position, the medial axis of a union of balls is a (flat) polytope in which every vertex is either the center of a ball in \mathcal{S} , or a vertex in the Voronoi diagram of the set \mathcal{V} of vertices of the boundary ∂S . In particular, it shows that the medial axis is closed and thus the two notions of medial and maximal balls coincide.

We argue that this particular property still holds without the general position assumption, specifically:

Theorem 3. The medial axis of a finite union of balls in \mathbb{R}^d can be decomposed into a finite collection of faces. Each face is the centers of balls which, when interpreted as $(d + 1)$ -dimensional points, are closed convex subsets of elliptic pencils with a finite number of extremal points.

Note that this property directly implies that the medial axis is closed. Because we have in general the series of inclusions $\text{MA}(X) \subseteq \text{Sk}(X) \subseteq \overline{\text{MA}(X)}$ for any subset $X \subseteq \mathbb{R}^d$ ([Mat88]), we deduce that the medial axis and the skeleton of a finite union of balls coincide. It follows that the two notions of medial ball and (inclusion-wise) maximal ball are one and the same. We can thus use the two appellations interchangeably in finite unions of balls.

The remainder of this section is dedicated to proving this result. In order to do so, we provide a characterization of the medial axis without the general position assumption, and adopt an approach similar to the one of Edelsbrunner.

2.5.1 State of the art when assuming general position

We start by recalling a few definitions from [Ede93] and [AK01]. Given any finite collection of balls \mathcal{S} , we can always consider its power diagram. It naturally decomposes the space into various power cells, and therefore induces a decomposition of the union of balls $S = \bigcup \mathcal{S}$ itself. This diagram is dual to a Delaunay triangulation. It consists of the simplices connecting the centers of every subset of balls \mathcal{B} whose power cells have a non-empty common intersection. The *dual complex* is a subcomplex of this triangulation in which we keep a simplex only if the intersection of power cells also intersects the shape S . The *dual shape* is then the underlying space of the dual complex.

When in general position, the space spanned by a k -simplex of the dual complex has dimensionality k . Furthermore, it is of interest to distinguish between regular and singular simplices of the dual complex. A simplex is called *singular* if it has dimension at most $d - 1$ and is not the face of a d -dimensional simplex of the dual complex. In contrast, the remaining simplices are called *regular* and intuitively correspond to the full-dimensional components of the dual shape.

Without the general position assumption, a k -simplex may degenerate into an affine subset of dimension strictly less than k , and its faces may partially overlap. For instance, for $k = 2$, three collinear points yield a “flat”, degenerate triangle. The notion of regular and singular simplices also becomes somewhat ambiguous, as there may be k -simplices, $k > d$, that span a d -dimensional subset of \mathbb{R}^d . We adopt here the convention that a regular simplex must be d -dimensional in \mathbb{R}^d , or be the face of such a d -dimensional simplex.

From this notion of dual shape, Edelsbrunner then goes on to show that the boundary of a union of balls is dual to the boundary of its dual complex, and then defines the *join* of a

simplex of the dual complex with its dual face in ∂S . Given two subsets X and Y , the *join of X and Y* is defined as $X \star Y = \bigcup_{x \in X} \bigcup_{y \in Y} [xy]$. The join of two sets is well-defined whenever any pair of segments from one set to the other are either disjoint, or meet at a common endpoint. For convenience, when one set is empty we have $X \star \emptyset = X$.

These particular joins of a finite union of balls have several properties, of which we cite only a few: the collection of joins forms a partition of the shape S ; and a point in a join admits as a closest point in the boundary ∂S some point u that lies on the face defining the join.

Amenta and Kolluri restate the properties of these joins and make use of the structure of the dual shape to then prove the following:

Theorem 4 ([AK01]). Assuming general position, let $S = \bigcup \mathcal{S}$ be a finite union of balls and \mathcal{V} its vertex set. Then its medial axis is composed of:

- (i) singular simplices of the dual complex,
- (ii) portions of the Voronoi diagram of \mathcal{V} that intersect regular simplices of the dual complex.




Figure 2.17 – Consider the collection of balls $\mathcal{S} = \{b_1, b_2\}$ and $S = \bigcup \mathcal{S}$. Because the power cells of b_1 and b_2 intersect at $u \in S$, the dual complex of S contains the singular simplex $\{c_1, c_2\}$. However, the medial axis of S only contains the two points c_1 and c_2 , and not the open segment $]c_1c_2[$. Theorem 4 does not apply without the general position assumption.

Observe in Figure 2.17 that without the general position assumption, the theorem above does not apply. To provide our own characterization without this general position assumption, we start by defining our own collection of “joins”, based on a division of the medial axis. We then argue that these joins coincide with the cells of the Voronoi diagram of ∂S restricted to S , and then that they are actually polytopes we can easily compute from \mathcal{S} . From there, we argue that we can retrieve the medial axis by inspecting specific links of the joins.

2.5.2 Link of a polytope

Before introducing the joins we will be working with, we first address the notion of link. This notion is already well-defined in the context of simplicial complexes, we here use a slight adaptation better suited for polytopes. Indeed, in the context of simplicial complexes, the link of a face is usually defined based on the combinatorial structure of the simplices incident to that face, and it is a collection of faces. Formally, given a simplicial complex K , the *star* of a simplex σ in K is the collection of simplices σ' which have σ as a face. Then the *link* of σ in K is the closure of its star, minus its star. So far, we intuitively used the notion of a face of a polytope. We now specify formally the meaning we give to this notion of face when applied to a convex polytope, as opposed to a simplicial complex. Indeed, a k -simplex is a collection of $k + 1$ vertices, its faces are any subset of these $k + 1$ vertices. In contrast, the face f of a convex polytope can be seen as the convex hull of a finite number of vertices. Contrary to a simplex however, any subset of these vertices does not necessarily define a face of f in the intuitive sense. We say that the convex hull of such a subset defines a proper face of f if and only if its affine hull has dimension strictly less than the dimension of the affine hull of f . With this specific notion of face, we can then define the geometric star and link in a similar way to their simplicial counterparts. Refer to Figure 2.18 for an illustration of Definitions 12 and 13.




Figure 2.18 – Star (green) and link (blue) for various points (and faces) of a convex polytope X . (a) Interior point x (and interior face of X); (b) point x on an edge (and edge a); (c) vertex x .

Definition 12 (Star and link for polytope faces). Consider a polytope $X \subseteq \mathbb{R}^d$ and a face f of X . The *star of f in X* is the collection of (open) faces of X that have f as a face. Then the *geometric star of f in X* is the underlying space of its star. We denote it as $\text{star}(f, X) = \overline{\{y \in \mathbb{R}^d \mid \exists f' \text{ face of } X, y \in f', f \subseteq \overline{f'}\}}$. The *geometric link of f in X* is then $\text{lnk}(f, X) = \text{star}(f, X) \setminus \text{star}(f, X)$.

We however need to extend this definition slightly so that the geometric link of any point of the polytope is well-defined. Indeed we will later have to consider the geometric link of any ball b , possibly not a vertex of $\text{CH}(\mathcal{B})$. This motivates the definition below.

Definition 13 (Star and link for polytope points). Consider a polytope $X \subseteq \mathbb{R}^d$ and a point $x \in X$. The *star of x in X* is the collection of (open) faces of X whose closure contains

x . Then the *geometric star of x in X* is the underlying space of its star. We denote it as $\text{star}(x, X) = \{y \in \mathbb{R}^d \mid \exists f' \text{ face of } X, y \in f', x \in \overline{f'}\}$. The *geometric link of x in X* is then $\text{lnk}(x, X) = \text{star}(x, \overline{X}) \setminus \text{star}(x, X)$.

Because this extension coincides with the previous definition when x is a 0-face of the polytope, we thus use the same appellation and notation for both definitions. From here on, we will refer to the geometric link of a point x in a polytope simply as the link of x .

Remark 37. Let $x \in X$, there is some open face f_x that contains x . We in fact have $\text{star}(x, X) = \text{star}(f_x, X)$ and thus $\text{lnk}(x, X) = \text{lnk}(f_x, X)$.

Remark 38. So long as the notion of faces is properly defined, the definitions above extend to non-convex polytopes. In particular if X is a convex polytope, its boundary ∂X is in general a non-convex polytope. However ∂X inherits its faces from X and we can apply our notions of star and link to ∂X .

By definition, the link of a boundary point x is the collection of points that share a face with x , and also belong to another face whose closure does not contain x . If we consider its link in the boundary polytope ∂X instead of X , it is then the collection of points that share a boundary face with x , and still lie on another boundary face. Then, the link of x in ∂X is also the collection of points y such that $N_{\mathbb{1}}^x \cap N_{\mathbb{1}}^y \neq \emptyset$ and $N_{\mathbb{1}}^y \setminus N_{\mathbb{1}}^x \neq \emptyset$, where $N_{\mathbb{1}}^x$ (resp. $N_{\mathbb{1}}^y$) denotes the unit normal cone to polytope X at x (resp. y). Alternatively, given a direction \mathbf{e} supported by a face that contains x , consider the half-line originating from x in the direction of \mathbf{e} . The link $\text{lnk}(x, \partial X)$ is also the collection of points in the polytope that are the furthest away from x along one of these half-lines. See Figure 2.19.

Property 2.47. Consider a closed bounded convex polytope $X \subseteq \mathbb{R}^d$ and $x \in \partial X$. We have equivalence between:

- (i) $y \in \text{lnk}(x, \partial X)$
- (ii) $N_{\mathbb{1}}(x, X) \cap N_{\mathbb{1}}(y, X) \neq \emptyset$ and $N_{\mathbb{1}}(y, X) \setminus N_{\mathbb{1}}(x, X) \neq \emptyset$.
- (iii) $\exists \mathbf{e}, (x + \mathbb{R}^+ \mathbf{e}) \cap \partial X = [xy]$ and $x \neq y$

Proof. We prove the series of implication (i) \implies (ii) \implies (iii) \implies (i).

Let $y \in \text{lnk}(x, \partial X)$. Because the link is a subset of ∂X , necessarily there exists some open boundary face f_1 of X that contains y . Because y does not belong to $\text{star}(x, \partial X)$ however, that open face cannot contain x in its closure. It follows that f_1 admits some unit normal vector $\mathbf{e} \notin N_{\mathbb{1}}(x, X)$. However $y \in f_1$ and thus $\mathbf{e} \in N_{\mathbb{1}}(y, X)$. Therefore $N_{\mathbb{1}}(y, X) \setminus N_{\mathbb{1}}(x, X) \neq \emptyset$. Because $y \in \text{star}(x, \partial X)$, there exists an open boundary face $f_2 \subseteq \text{star}(x, \partial X)$ such that $y \in \overline{f_2}$. Hence, x and y lie on a common (closed) face. The outward unit normal vectors of f_2 all belong to $N_{\mathbb{1}}(x, X) \cap N_{\mathbb{1}}(y, X)$, thus $N_{\mathbb{1}}(x, X) \cap N_{\mathbb{1}}(y, X) \neq \emptyset$. It follows that (ii) holds.

By contradiction suppose that (ii) holds but (iii) does not. The condition $N_{\mathbb{1}}(y, X) \setminus N_{\mathbb{1}}(x, X) \neq \emptyset$ implies that x and y must be distinct, let \mathbf{e} be the unit direction from x

to y . The constraint $N_{\mathbb{1}}(x, X) \cap N_{\mathbb{1}}(y, X) \neq \emptyset$ implies that there exists an open boundary face f_2 such that $x, y \in \overline{f_2}$. Indeed, because X is bounded, every non null vector is the outward normal to some face. Because f_2 must be convex, $[xy] \subseteq \overline{f_2} \subseteq \partial X$. Let z such that $(x + \mathbb{R}^+ \mathbf{e}) \cap \partial X = [xz]$. We have $y \in]xz[$. Let $\mathbf{n} \in N_{\mathbb{1}}(y, X) \setminus N_{\mathbb{1}}(x, X)$ and f_1 the face of X with greatest dimensionality that admits \mathbf{n} as an outward normal. By definition of an outward normal vector, for all $c \in \overline{f_1}$ and $c' \in X \setminus \overline{f_1}$ we have $\langle \mathbf{n}, c' - c \rangle < 0$. Because $x \notin \overline{f_1}$ and $y \in f_1$, $\langle \mathbf{n}, x - y \rangle = -\|x - y\| \langle \mathbf{n}, \mathbf{e} \rangle < 0$. In other words, $\langle \mathbf{n}, \mathbf{e} \rangle > 0$. If z belongs to $\overline{f_1}$, then $\langle \mathbf{n}, z - y \rangle = \|z - y\| \langle \mathbf{n}, \mathbf{e} \rangle = 0$, which contradicts $\langle \mathbf{n}, \mathbf{e} \rangle > 0$. However if z does not belong to $\overline{f_1}$, then $\langle \mathbf{n}, z - y \rangle = \|z - y\| \langle \mathbf{n}, \mathbf{e} \rangle < 0$, which is also impossible. Therefore condition (iii) must hold.

Assume that (iii) holds. Since $[xy] \subseteq \partial X$, we have $y \in \overline{\text{star}(x, \partial X)}$. We show that $y \notin \text{star}(x, \partial X)$, therefore $y \in \text{lnk}(x, \partial X)$ and (i) also holds. By contradiction, assume there exists an open boundary face f such that $y \in f$ and $x \in \overline{f}$. Because f is convex, $]xy[\subseteq f$ and \mathbf{e} is parallel to f . In addition, f is an open face, hence there exists $r > 0$ such that $b(y, r) \cap \text{aff}(f) \subseteq f$. In particular this implies that the line segment between y and $y + r\mathbf{e}$ belongs to f and thus to ∂X . This contradicts (iii) and is thus impossible. Therefore, no such face f exists and $y \notin \text{star}(x, \partial X)$, (i) also holds. \square




Figure 2.19 – The star of x in ∂X is composed of point x and the two open green segments. Its link is $\{y, y'\}$.

Remark 39. Suppose $y \in \text{lnk}(x, \partial X)$. Then any pair of unit vectors $\mathbf{n} \in N_{\mathbb{1}}(x, X) \cap N_{\mathbb{1}}(y, X)$ and $\mathbf{e} \in N_{\mathbb{1}}(y, X) \setminus N_{\mathbb{1}}(x, X)$, is non-collinear. To see this, suppose by contradiction that $\mathbf{e} = -\mathbf{n}$. We deduce that for all $z \in X$, we have $\langle \mathbf{e}, z - y \rangle = 0$. This implies $\langle \mathbf{e}, z - x \rangle = \langle \mathbf{e}, z - y \rangle + \langle \mathbf{e}, y - x \rangle = 0$, and thus \mathbf{e} is also a normal at point x , which is absurd.

2.5.3 From joins to Voronoi cells...

Recall that $S = \bigcup \mathcal{S}$ is a finite union of balls. Consider a point $u \in \partial S$, and let $J_u \subseteq \text{MA}(S)$ the points of the medial axis that admit u as a closest point in ∂S . The union $\bigcup_u J_u$ covers the medial axis, but these sets overlap one another. In the next few sections, we show that

it is possible to compute these J_u , and that a finite number of them suffices to retrieve the entire medial axis. In order to do so, we investigate the properties of the join of u with J_u ,

$$\text{join}(u, S) = \{u\} \star J_u$$

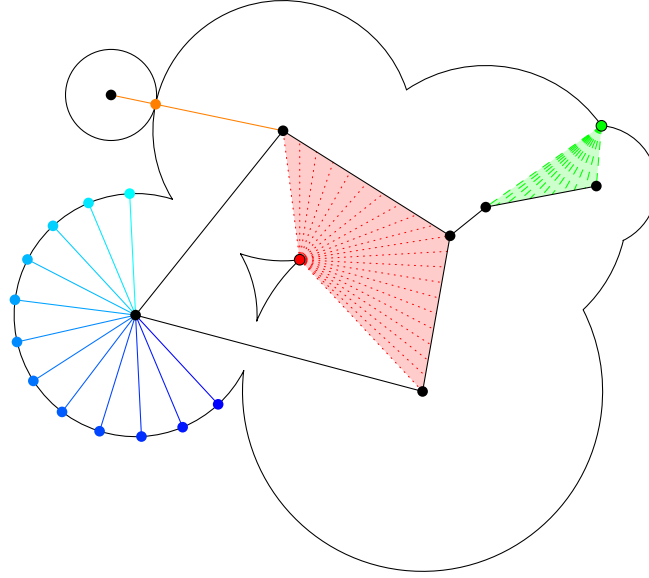


Figure 2.20 – Joins for various boundary points. The medial axis of the shape is the collection of black line segments and vertices.

See Figure 2.20 for an illustration. First we argue that this join is well-defined.

Property 2.48. *The join $\text{join}(u, S)$ is well-defined.*

Proof. Let $x, y \in J_u$, then either u, x, y are not collinear, or $u \in [xy]$. Indeed, assume by contradiction that x lies in the open segment $]uy[$. Denote by b_x the ball centered at x and whose bounding $(d - 1)$ -sphere goes through u , and likewise b_y centered at y . Then $b_x \subseteq b_y$, and the two balls are tangent at u . Therefore x admits u as unique closest point in ∂S and cannot belong to the medial axis, which contradicts the definition of J_u . \square

Consider the Voronoi diagram of ∂S restricted to S . For $u \in \partial S$, we denote its cell in this diagram by $\text{vor}_S(u, \partial S) = \{x \in S \mid \forall v \in \partial S, \|x - u\| \leq \|x - v\|\}$. The purpose of this section is to prove the equality:

$$\text{join}(u, S) = \text{vor}_S(u, \partial S) \tag{2.7}$$

Remark 40. Because Voronoi cells are closed and convex, Equality 2.7 implies that $\text{join}(u, S)$ is closed and convex.

We proceed by double inclusion. Note that the argument used in the proof of Property 2.48 also shows that we must have the inclusion $\text{join}(u, S) \subseteq \text{vor}_S(u, \partial S)$. Indeed, consider a point $x \in \text{join}(u, S)$. By definition of the join, there exists $y \in J_u$ such that $x \in [uy]$. The previous logic with b_x and b_y still applies, hence we deduce that u is the unique closest point to x in ∂S . Thus, $x \in \text{vor}_S(u, \partial S)$.

The remainder of this section is dedicated to the proof of the reverse inclusion. Let $x \in \text{vor}_S(u, \partial S)$ and consider b_x the ball centered at x and whose bounding $(d-1)$ -sphere goes through u . Point x belongs to $\text{join}(u, S)$ if and only if there exists $y \in J_u$ such that $x \in [uy]$. Necessarily, $b_x \subseteq S$ and in particular, b_x satisfies local inclusion in S locally around u . Therefore we have $x \in \overset{\star}{\text{coni}}(u, S) = \overset{\star}{\text{coni}}(u, C)$, where C is the collection of centers of a local representation of S around u (see Definition 9 and Section 2.4). Recall that $\overset{\star}{\text{coni}}(u, C) = \overset{\circ}{\text{coni}}(u, C) \cup \text{CH}(\{u\} \cup C)$. We distinguish two cases:

- (1) $x \in \overset{\circ}{\text{coni}}(u, C)$, and
- (2) $x \in \overset{\star}{\text{coni}}(u, C) \setminus \overset{\circ}{\text{coni}}(u, C)$.

In each case, we rely on a technical lemmas to prove the existence of an appropriate $y \in J_u$ such that $x \in [uy]$. We start with the case (2).

Lemma 2.49. *Let $u \in \partial S$, \mathcal{B} a cardinal minimum representation of S locally around u and $C = c(\mathcal{B})$ the collection of centers of \mathcal{B} . Then the link of u in $\partial \text{CH}(\{u\} \cup C)$ belongs to J_u .*

Proof. Let x in the link of u in polytope $\partial \text{CH}(\{u\} \cup C)$. Recall that b_x is the ball centered at x whose bounding $(d-1)$ -sphere goes through x . We show that $x \in J_u$ by proving that the ball b_x is medial in S . Let $N_{\mathbb{1}}^x$ (resp. $N_{\mathbb{1}}^u$) be the unit normal cone to polytope $\text{CH}(\{u\} \cup C)$ at point x (resp. u). By Property 2.47, we have $N_{\mathbb{1}}^u \cap N_{\mathbb{1}}^x \neq \emptyset$ and $N_{\mathbb{1}}^x \setminus N_{\mathbb{1}}^u \neq \emptyset$. Let $\mathbf{n} \in N_{\mathbb{1}}^u \cap N_{\mathbb{1}}^x$ and $\mathbf{e} \in N_{\mathbb{1}}^x \setminus N_{\mathbb{1}}^u$ be two unit vectors. Recall by Remark 39 that \mathbf{e} and \mathbf{n} are non-collinear. For $0 \leq \lambda \leq 1$, let $\mathbf{e}_\lambda = \lambda \mathbf{e} + (1 - \lambda) \mathbf{n}$. Since \mathbf{e} and \mathbf{n} are non collinear, \mathbf{e}_λ is non null for all λ . By convexity of normal cones, the \mathbf{e}_λ all are normal to $\text{CH}(\{u\} \cup C)$ at x . We can thus consider the image of their normalization by m_{b_x} . By continuity of m_{b_x} , it follows that in every neighbourhood of u , ∂b_x contributes several points to the boundary of $\bigcup \mathcal{B}$. Because \mathcal{B} locally represents S around u , we thus deduce that b_x is a medial ball in S and hence $x \in J_u$. See Figure 2.21 for an illustration. \square

When $x \in \overset{\star}{\text{coni}}(u, C) \setminus \overset{\circ}{\text{coni}}(u, C)$, we necessarily have $x \in \partial \text{CH}(\{u\} \cup C)$. Thus the segment $[xu]$ is included in $\partial \text{CH}(\{u\} \cup C)$. Then, let \mathbf{e} be such that $x \in u + \mathbb{R}^+ \mathbf{e}$, and y be the point such that $(u + \mathbb{R}^+ \mathbf{e}) \cap \partial \text{CH}(\{u\} \cup C) = [uy]$. By Property 2.47, $y \in \text{lnk}(u, \partial \text{CH}(\{u\} \cup C))$, which implies $y \in J_u$ by Lemma 2.49, and thus $x \in \text{join}(u, S)$. This concludes the proof in case (2).

For case (1), we have $x \in \overset{\circ}{\text{coni}}(u, C)$ and we rely on this other lemma instead.




Figure 2.21 – The link of u in the boundary polytope $\partial \text{CH}(\{u\} \cup C)$ is the two points x and y . The vectors between \mathbf{e} and \mathbf{n} corresponds to the red circular arc, supported by ∂b_x . Every open neighbourhood of u contains infinitely many points of that arc.

Lemma 2.50. *Let $u \in \partial S$, \mathcal{B} a cardinal minimum representation of S locally around u and $C = c(\mathcal{B})$ its collection of centers. Let \mathbf{e} be a unit vector such that the half-line $u + \mathbb{R}^+ \mathbf{e}$ is contained in $\mathring{\text{coni}}(u, C)$. Then this half-line intersects J_u at some point y , and $(u + \mathbb{R}^+ \mathbf{e}) \cap \text{vor}_S(u, \partial S) = [uy]$.*

Proof. We first prove the existence of an intersection point y between the half-line and J_u . For $t \geq 0$, let $x_t = u + t\mathbf{e}$. Consider b_t the ball centered at x_t whose bounding $(d-1)$ -sphere goes through u , $b_t = b(x_t, t)$. The b_t 's form a nested collection of balls that are pairwise tangent at u . See Figure 2.22. For small values of $t > 0$, we have $b_t \subseteq \bigcup \mathcal{B} \subseteq S$. Since S is bounded and b_∞ is not, there exists a limit value $T > 0$ such that $b_T \subseteq S$, and for all $t > T$, $b_t \cap S^c \neq \emptyset$. For such $t > T$, let $z_t \in b_t \cap S^c$. As t goes down to T , the sequence (z_t) admits a limit point $z \in \partial b_T$. We claim that $z \neq u$. Indeed for all $t > 0$, $x_t \in (u + \mathbb{R}^+ \mathbf{e}) \subseteq \mathring{\text{coni}}(u, S)$ and b_t satisfies local inclusion in S in the neighbourhood of u . Let $r > 0$ such that $b_T \cap b(u, r) \subseteq \bigcup \mathcal{B}$. Necessarily, $z \notin b(u, r)$ and thus $z \neq u$. It follows that b_T is medial and its center x_T admits two distinct closest points $u, z \in \partial S$. In particular, $x_T \in J_u$ is the intersection point y that we seek.

For $t \leq T$, we have $x_t \in [uy]$. Hence $b_t \subseteq b_T$ and x_t admits u as a closest point in ∂S . Therefore $[uy] \subseteq \text{vor}_S(u, \partial S)$. For $t > T$ however, $b_T \subsetneq b_t$. In particular, we then have $z \in \mathring{b}_t$ and thus $\|x_t - z\| < \|x_t - u\|$. It follows that $x_t \notin \text{vor}_S(u, \partial S)$, which concludes the proof. \square

Given $x \in \mathring{\text{coni}}(u, C) \cap \text{vor}_S(u, \partial S)$, distinct from u , we can always consider the unit vector \mathbf{e} from u to x . This vector satisfies the assumption of the above lemma, from which we deduce that $x \in \text{join}(u, S)$. This concludes the proof in case (1), and thus Equality (2.7) is true.

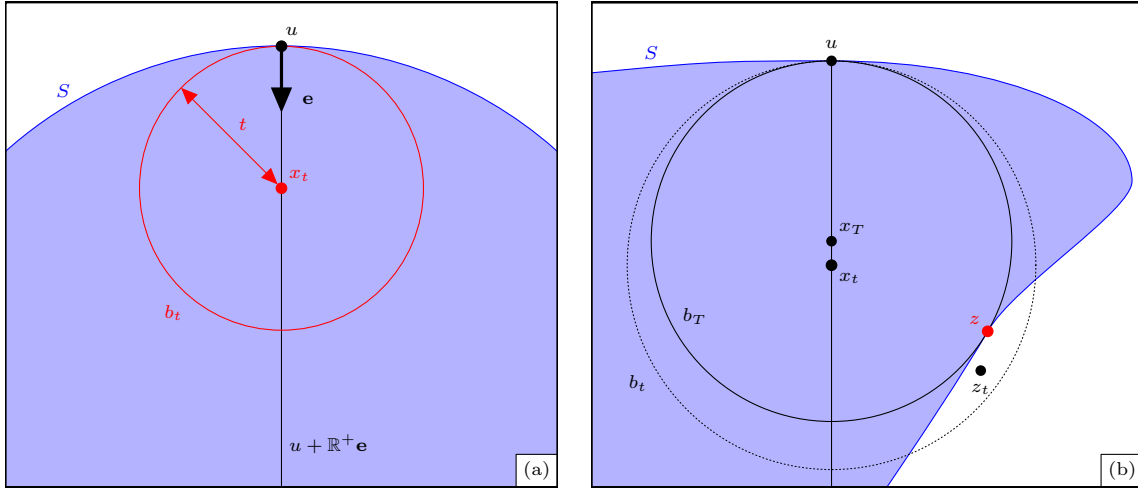


Figure 2.22 – Illustration for the proof of Lemma 2.50. (a) Construction of x_t and b_t ; (b) second boundary point z .

2.5.4 ... to polytopes...

In this section we show that our joins are in fact convex polytopes. To do so, we let \mathcal{V} be the set of simple vertices of ∂S and prove the identity below

$$\text{join}(u, S) = \begin{cases} \text{coni}(u, S) \cap \text{vor}(u, \mathcal{V}) & \text{if } u \text{ is a simple vertex,} \\ \text{coni}(u, S) & \text{otherwise} \end{cases} \quad (2.8)$$

where $\text{vor}(u, \mathcal{V})$ denotes the cell of u in the usual Voronoi diagram of \mathcal{V} . First, let us show that this suffices to conclude that $\text{join}(u, S)$ is a convex polytope.

Lemma 2.51. *join(u, S) is a convex polytope.*

Proof. We assume identity (2.8) to be true.

If u is a non-simple vertex, $\text{coni}(u, S)$ is a convex polytope and the result is straightforward. Otherwise, let u be a simple vertex. Let \mathcal{B} a cardinal minimum local representation of S around u , and $C = c(\mathcal{B})$ its collection of centers. We have $\text{join}(u, S) = (\text{CH}(\{u\} \cup C) \cap \text{vor}(u, \mathcal{V})) \cup (\text{coni}(u, C) \cap \text{vor}(u, \mathcal{V}))$. Denote by $A = \text{CH}(\{u\} \cup C) \cap \text{vor}(u, \mathcal{V})$ and $B = \text{coni}(u, C) \cap \text{vor}(u, \mathcal{V})$. Note that by Equation (2.7) we already know that $\text{join}(u, S)$ is closed and convex. Therefore $\text{join}(u, S) = \overline{A \cup B} = \overline{A} \cup \overline{B}$. Since A and \overline{B} are both closed polytopes, $\text{join}(u, S)$ is also a polytope. We already know from Equation (2.7) that $\text{join}(u, S)$ is convex, therefore this concludes the proof. \square

Remark 41. The proof above also shows that when u is a simple vertex, we have $\text{coni}(u, S) \cap \text{vor}(u, \mathcal{V}) = \text{coni}(u, C) \cap \text{vor}(u, \mathcal{V})$, where C is the collection of centers of a local representation of S around u .

The remainder of this section is dedicated to the proof of identity (2.8). We proceed by double inclusion, depending on whether u is a simple vertex or not.

In the previous section, we already used the inclusion $\text{join}(u, S) \subseteq \text{coni}^*(u, S)$ implicitly. Indeed, let $x \in \text{join}(u, S)$ and $y \in J_u$ such that $x \in [uy]$. Consider the ball b_x (resp. b_y) centered at x (resp. y) and whose bounding $(d-1)$ -sphere goes through u . Because b_y is medial in S , we have $b_x \subseteq b_y \subseteq S$. In particular, b_x satisfies local inclusion in S around u , thus $x \in \text{coni}^*(u, S)$. Hence the inclusion

$$\text{join}(u, S) \subseteq \text{coni}^*(u, S) \quad (2.9)$$

Again, let \mathcal{B} be a cardinal minimum local representation of S around point u , and let $C = c(\mathcal{B})$ be its collection of centers. We have $\text{coni}^*(u, S) = \text{coni}^*(u, C)$. Consider $x \in \text{CH}(\{u\} \cup C)$ and b_x the ball centered at x whose bounding $(d-1)$ -sphere goes through u . We have $u \in b_x \subseteq \bigcup \mathcal{B} \subseteq S$, thus x admits u as a closest point in ∂S . Therefore

$$\text{CH}(\{u\} \cup C) \subseteq \text{vor}_S(u, \partial S) = \text{join}(u, S) \quad (2.10)$$

When u is not a simple vertex, then $\text{coni}^\circ(u, S) = \emptyset$. The two inclusions (2.9) and (2.10) thus yield the desired equality $\text{join}(u, S) = \text{CH}(\{u\} \cup C) = \text{coni}^*(u, S)$.

Otherwise, u is a simple vertex and $\text{coni}^\circ(u, S) \neq \emptyset$. Note that $\text{vor}_S(u, \partial S) \subseteq \text{vor}(u, \mathcal{V})$ and hence

$$\text{join}(u, S) \subseteq \text{coni}^*(u, S) \cap \text{vor}(u, \mathcal{V})$$

In addition, inclusion (2.10) still holds. Remains to prove the inclusion $\text{coni}^\circ(u, C) \cap \text{vor}(u, \mathcal{V}) \subseteq \text{join}(u, S)$. We rely on the technical lemma below.

Lemma 2.52. *Let $u \in \partial S$ a simple vertex, \mathcal{B} a cardinal minimum local representation of S around u and $C = c(\mathcal{B})$ its collection of centers. Let \mathbf{e} be a unit vector such that the half-line $u + \mathbb{R}^+\mathbf{e}$ is contained in $\text{coni}^\circ(u, C)$. This half-line intersects J_u at a point y that admits a simple vertex $v \neq u$ as closest point in ∂S .*

Proof. The existence of an intersection point y has already been established in Lemma 2.50. Note that it also implies the unicity of y . We already know that u is a closest point of y in ∂S . Denote by b_y the ball centered at y whose bounding $(d-1)$ -sphere goes through u . We show that there exists a simple vertex $v \neq u$ such that $v \in \partial b_y$. To do so, we explicitly construct a candidate point $v \in \partial b_y \cap \partial S$, and then show that this point v must be a simple vertex of ∂S .

Construction of v First we argue that in a small neighbourhood of u , point u is in fact the only contribution of ∂b_y to the boundary ∂S . To prove this, our strategy is to rely on Property 2.38. Thus we must show that for a certain collection of balls \mathcal{B}_0^y , the image of m_{b_y, \mathcal{B}_0^y} does not contain u . For the following paragraphs, refer to Figure 2.23.

Let $b_0 = \{u\}$ and $\mathcal{B}_0 = \mathcal{B} \cup \{b_0\}$. Let $\mathcal{B}_0^y = \mathcal{B}_0 \cup \{b_y\}$ and $C_0^y = c(\mathcal{B}_0^y)$. Consider the polytope $\text{CH}(C_0^y)$. Let $N_{\mathbb{1}}^y$ (resp. $N_{\mathbb{1}}^u$) be the unit normal cone to polytope $\text{CH}(C_0^y)$ at point y (resp. u). Because $u \in \partial S$, we have $u \in \partial \bigcup \mathcal{B}_0^y$. By Property 2.39, we have $N_{\mathbb{1}}^u \neq \emptyset$. We want to show that $N_{\mathbb{1}}^y \cap N_{\mathbb{1}}^u = \emptyset$. Thus by Property 2.31 we deduce that the image of m_{b_y, \mathcal{B}_0^y} does not contain u , which concludes the proof that u is an isolated point in $\partial b_y \cap \partial S$.

By contradiction, suppose that $N_{\mathbb{1}}^y \cap N_{\mathbb{1}}^u$ is non-empty, and let \mathbf{e}' be a common outward normal to y and u . Because $\text{CH}(C_0^y)$ is bounded, Remark 22 guarantees that \mathbf{e}' is the outward normal to some open boundary face f of the polytope. Since $\text{CH}(C_0^y)$ is convex, f is also convex, and the open segment $]uy[$ thus lies in the open face f . However, because of our assumption on \mathbf{e} , the direction from u to y , the open segment $]uy[$ must lie in the interior of the polytope $\text{CH}(C_0^y)$. This is impossible, therefore \mathbf{e}' cannot exist, and $N_{\mathbb{1}}^y \cap N_{\mathbb{1}}^u = \emptyset$.

Let v be a closest point to u in $\partial b_y \cap \partial S \setminus \{u\}$. By the above, we know that v is well defined and distinct from u . We must now show that v is a simple vertex.

The candidate point v is a simple vertex Let \mathcal{B}^v be a cardinal minimum local representation of S around v , and $C^v = c(\mathcal{B}^v)$ its collection of centers. The collection \mathcal{B}^v defines a pencil \mathcal{P}_v with radical space V_v , radical ball b_0^v , and radical sphere s_0^v . Denote the center and radius of b_0^v by c_0^v and by r_0^v .

By contradiction, assume that v is not a simple vertex, then it is either a non-simple vertex or part of an open k -face, $k \geq 1$. We first show that y belongs to a certain polytope X_v , and that its unit normal cone $N_{\mathbb{1}}(y, X_v)$ to that polytope contains two unit vectors: a unit vector $\mathbf{n} \in \vec{V}_v$ which satisfies $v = c_0^v + r_0^v \mathbf{n}$, and the unit vector \mathbf{e}' from c_0^v to u . We then derive a contradiction from those two vectors. Since the polytope differs depending on the properties of v and \mathcal{P}_v , we address the two cases separately: (a) v is a non-simple vertex, or (b) v belongs to an open k -face, $k \geq 1$.

(a) Suppose v is a non-simple vertex. Let $\mathcal{B}_0^v = \mathcal{B}^v \cup \{b_0^v\}$ and $C_0^v = c(\mathcal{B}_0^v)$ the collection of centers. By definition, \mathcal{P}_v is parabolic and $\text{coni}(v, C^v)$ is flat with empty interior. That is, \mathcal{P}_v has dimension at most $d - 1$, or equivalently its radical space V_v has dimension at least 1. Also, $r_0^v = 0$ and $\text{coni}(v, C^v) = \text{CH}(\{v\} \cup C^v) = \text{CH}(C_0^v)$. Refer to Figure 2.24 for an illustration in \mathbb{R}^2 .

Notice that $y \in \text{CH}(C_0^v)$. Indeed, although b_y may not belong to \mathcal{B}^v , ball b_y must still satisfy local inclusion in $\bigcup \mathcal{B}^v = \bigcup \mathcal{B}_0^v$ around v . Hence $y \in \text{coni}(v, C^v) = \text{CH}(C_0^v)$. We can thus consider $X_v = \text{CH}(C_0^v)$ and $N_{\mathbb{1}}(y, X_v)$ the unit normal cone to that polytope at y .

Since V_v has dimension at least 1, let \mathbf{n} be any unit vector of \vec{V}_v . We have $v = c_0^v = c_0^v + r_0^v \mathbf{n}$. Recall that $\vec{V}_v \subseteq N(y, X_v)$. Indeed, because \mathcal{P}_v is parabolic, its centers' space satisfies $\text{CH}(C_0^v) \subseteq c(\mathcal{P}_v)$, and by Corollary 2.4, $\vec{c}(\mathcal{P}_v) \oplus \vec{V}_v$. It follows that \vec{V}_v belongs to the normal cone of every point of polytope $\text{CH}(C_0^v)$. Hence $\mathbf{n} \in N_{\mathbb{1}}(y, X_v)$. As for \mathbf{e}' , recall that $u \neq v = c_0^v$, this unit vector is thus well defined. In addition, u belongs to the contribution of b_y to the global boundary, and is distinct from the radical center of \mathcal{P}_v . Thus Property 2.37 ensures that the pre-image of u by m_{b_y, \mathcal{B}_0^v} is non-empty. Necessarily, \mathbf{e}' belongs to that pre-image. Thus $\mathbf{e}' \in N_{\mathbb{1}}(y, X_v)$.

(b) Suppose that v belongs to an open k -face, $k \geq 1$. By definition, \mathcal{P}_v must then be elliptic and $\text{coni}^*(v, C^v)$ must be flat with empty interior. That is, \mathcal{P}_v has dimension at most $d-2$ and its radical space dimension at least 2. Also, $r_0^v > 0$ and $\text{coni}^*(v, C^v) = \text{CH}(\{v\} \cup C^v)$. Refer to Figure 2.25 for an illustration in \mathbb{R}^3 .

As above, b_y must satisfy local inclusion in $\bigcup \mathcal{B}_v$ around v and we deduce $y \in \text{CH}(\{v\} \cup C^v)$. We argue that in fact, $y \in \text{CH}(C^v)$. Because \mathcal{P}_v is elliptic, and v belongs to its radical sphere, then v cannot lie in the centers' space $\text{aff}(C^v)$ of \mathcal{P}_v . It follows that v is a vertex of $\text{CH}(\{v\} \cup C^v)$. By contradiction, if $y \notin \text{CH}(C^v)$ then there exists $z \in \text{CH}(C^v)$ such that $y \in]zv[$. Let b_z be the ball centered at z whose boundary $(d-1)$ -sphere goes through v , then $b_y \setminus \{v\} \subseteq b_z$. Because $v \neq u \in \partial b_y \cap \partial S$, this is impossible and thus we must have $y \in \text{CH}(C^v)$. It follows that we can consider $X_v = \text{CH}(C^v)$ and $N_{\mathbb{1}}(y, X_v)$, the unit normal cone to that polytope at y .

Let \mathbf{n} be the unit vector from c_0^v to v , we have $v = c_0^v + r_0^v \mathbf{n}$ and $\mathbf{n} \in \vec{V}_v$. Again, $\vec{V}_v \subseteq N(y, X_v)$ hence $\mathbf{n} \in N_{\mathbb{1}}(y, X_v)$. Because \mathcal{P}_v is elliptic, $c_0^v \in (\bigcup \mathcal{B}^v)^\circ \subseteq \dot{S}$. Since $u \in \partial S$, we have $u \neq c_0^v$ and \mathbf{e}' is again well defined. A similar argument relying on Property 2.37 yields $\mathbf{e}' \in N_{\mathbb{1}}(y, X_v)$.

Whether v is a non-simple vertex or part of an open k -face, we just proved that there exists a polytope X_v in which the unit normal cone $N_{\mathbb{1}}(y, X_v)$ contains the two unit vectors \mathbf{n} and \mathbf{e}' with the claimed properties. We now argue that \mathbf{n} and \mathbf{e}' cannot be collinear. By contradiction, suppose that they are, for instance as in Figure 2.26. The intersection between b_y and the line of direction \mathbf{n} that goes through c_0^v is the line segment between $v = c_0^v + r_0^v \mathbf{n}$ and $c_0^v - r_0^v \mathbf{n}$. Indeed, the line $c_0^v + \mathbb{R}\mathbf{n}$ must be in the radical space V_v , and b_y always belongs to the complete pencil \mathcal{P}_v . Because u lies on the boundary of b_y , u necessarily coincides with one endpoint of that line segment. We know that $u \neq v$, hence $\mathbf{e}' = -\mathbf{n}$ and $u = c_0^v - r_0^v \mathbf{n}$. In addition, r_0^v must be positive. Hence if \mathbf{e}' and \mathbf{n} are collinear, \mathcal{P}_v must be an elliptic pencil, and v belongs to an open face of S of dimension at least 1. By definition of v , it is the closest point to u in $\partial b_y \cap \partial S$. Yet, u and v are diametrical opposite in s_0^v . Thus $s_0^v \cap \partial S = \{u, v\}$. In particular in the neighbourhood of v , v is the only point of s_0^v in the boundary of S . Because s_0^v supports the open face of ∂S that contains v , this is absurd. Therefore, \mathbf{e}' and \mathbf{n} are not collinear.

For $0 \leq \lambda \leq 1$, let $\mathbf{e}_\lambda = \lambda \mathbf{e}' + (1 - \lambda) \mathbf{n}$. By non-collinearity of \mathbf{e}' and \mathbf{n} , the \mathbf{e}_λ 's are all non-null, and by convexity they all belong to $N(y, X_v)$. Also, they form a path connecting \mathbf{e}' to \mathbf{n} in $N(y, X_v)$. Using m_{b_y, \mathcal{B}_0^v} in case (a), and m_{b_y, \mathcal{B}^v} in case (b), we map each \mathbf{e}_λ to a point $x_\lambda \in \partial b_y \cap \partial(\bigcup \mathcal{B}^v)$, and these x_λ 's form a path connecting v to u . This path is in fact a circular arc and for $\lambda < \mu$, we have $\|u - x_\lambda\| > \|u - x_\mu\|$. See Figure 2.27. However, because \mathcal{B}^v is a local representation of S in the neighbourhood of v , for $\lambda > 0$ small enough, we have $x_\lambda \in \partial S \cap \partial b_y$. Hence there exists some x_λ , $\|u - x_\lambda\| < \|u - v\|$, which contradicts the definition of v , and is thus impossible. Therefore, v must be a simple vertex of ∂S . \square

From the lemma above, we deduce that if a point x belongs to $\text{coni}(u, C) \cap \text{vor}(u, \mathcal{V})$, then there exists $y \in J_u$ such that $x \in [uy]$. It follows that $x \in \text{join}(u, S)$ and hence we have the inclusion $\text{coni}(u, C) \cap \text{vor}(u, \mathcal{V}) \subseteq \text{join}(u, S)$. Therefore, the identity (2.8) is true for every boundary point $u \in \partial S$.

For easy reference, we summarize the various expressions of $\text{join}(u, S)$ below.

Property 2.53. *Let S a finite union of balls and \mathcal{V} the collection of its simple vertices. Consider $u \in \partial S$, then*

$$\begin{aligned} \text{join}(u, S) &= \{u\} \star J_u \\ &= \text{vor}_S(u, \partial S) \\ &= \begin{cases} \text{coni}(u, S) \cap \text{vor}(u, \mathcal{V}) & \text{if } u \text{ is a simple vertex,} \\ \text{coni}(u, S) & \text{otherwise} \end{cases} \end{aligned}$$




Figure 2.23 – Illustration for the construction of v in the proof of Lemma 2.52. In this figure, $\text{CH}(c(\mathcal{B}_0)) = \text{CH}(C_0^y)$. Because $N_{\mathbb{1}}^y \cap N_{\mathbb{1}}^u = \emptyset$, point u is isolated in the contribution of b_y to the boundary, and we can define v as a closest point to u in $\partial b_y \cap \partial S \setminus \{u\}$.

2.5.5 ... to the medial axis

Now that we established that joins are convex polytopes, we can consider their links. Specifically, we consider the link of u in $\text{join}(u, S)$.

Lemma 2.54. *Let $u \in \partial S$, then $\text{lnk}(u, \text{join}(u, S)) = J_u$.*

Proof. We proceed by double inclusion. Consider $x \in \text{lnk}(u, \text{join}(u, S))$, we show that $x \in J_u$. Note that $x \in \text{join}(u, S)$ and $x \neq u$. By definition of the join, there is $y \in J_u$ such that $x \in [uy]$. In addition let \mathbf{e} be the unit vector from u to x , and consider the half-line originating from




Figure 2.24 – Illustration for case (a) in the proof of Lemma 2.52. Both \mathbf{n} and \mathbf{e}' belong to the normal cone to $\text{CH}(C_0^v)$ at point y .

u supported by \mathbf{e} . Then, $(u + \mathbb{R}^+\mathbf{e}) \cap \text{join}(u, S) =]uy]$. In other words, x belongs to J_u if and only if $x = y$. Recall that $\text{join}(u, S)$ is a convex polytope, hence there exists an open face f that contains the open line segment $]uy[$. Assume by contradiction that $x \in]uy[$, then we have $x \in f$ and $u \in \bar{f}$. Thus, $x \in \text{star}(u, \text{join}(u, S))$ which contradicts the assumption that x belongs to the link. This is thus impossible and we conclude that $x = y$. Therefore, $x \in J_u$ and $\text{lnk}(u, \text{join}(u, S)) \subseteq J_u$.

Conversely, let $x \in J_u$, we now prove that $x \in \text{lnk}(u, \text{join}(u, S))$. By convexity of $\text{join}(u, S)$, x necessarily belongs to the closure of $\text{star}(u, \text{join}(u, S))$. Indeed, the relative interior face of $\text{join}(u, S)$ is in the star, hence the whole polytope is in the closure of the star. Now assume by contradiction that x is in the star, there exists f an open face of $\text{join}(u, S)$ such that $x \in f$ and $u \in \bar{f}$. Let \mathbf{e} be the unit vector from u to x . Because $x \in J_u$, necessarily $x \neq u$ and \mathbf{e} is well-defined. Because f is open, for $h > 0$ sufficiently small, we deduce $x + h\mathbf{e} \in f$. This implies that there is $y \in J_u$ such that $x \in]uy[$, and hence that the join between $\{u\}$ and J_u is ill-defined, which is absurd. Therefore, $x \in \text{lnk}(u, \text{join}(u, S))$ and $J_u \subseteq \text{lnk}(u, \text{join}(u, S))$. \square

Given $u \in \partial S$, we can easily compute $\text{join}(u, S)$ by identity (2.8). Indeed, it is sufficient to have any local representation of S around u to obtain $\text{coni}^*(u, S)$, and we can also compute the Voronoi diagram of the set of simple vertices. See for example Figure 2.28 when u is a (simple) vertex. From $\text{join}(u, S)$, we can then compute J_u as the link of u in that polytope. Because $\bigcup_u J_u = \text{MA}(S)$, computing the J_u 's allows us to investigate the medial axis. In fact, we claim there are only finitely many distinct J_u 's, for $u \in \partial S$. Specifically, if u and v belong to the same open face, we have $J_u = J_v$.

Property 2.55. *Let f be an open k -face of S , with $k \geq 1$. Then for all $u, v \in f$, $J_u = J_v$.*




Figure 2.25 – Illustration for case (b) in the proof of Lemma 2.52. Both \mathbf{n} and \mathbf{e}' belong to the normal cone to $\text{CH}(C^v)$ (line segment between the blue and red centers) at point y .

Proof. By Property 2.46, let \mathcal{B} a finite collection of balls such that it is the unique cardinal minimum local representation of S at every point of f . Denote by $C = c(\mathcal{B})$ the collection of centers of \mathcal{B} , and $\mathcal{P} = \text{aff}(\mathcal{B})$ the pencil it defines. Because f is a k -face, $k \geq 1$, \mathcal{P} is elliptic of dimension at most $d - 2$. In particular for all $u \in f$, we deduce from identity (2.8) that $\text{join}(u, S) = \text{coni}(u, S) = \text{CH}(\{u\} \cup C)$. Because u does not belong to $\text{aff}(C)$, we have $J_u = \text{lnk}(u, \text{join}(u, S)) = \text{CH}(C)$. Hence, $\forall u, v \in f$, we have $J_u = J_v = \text{CH}(C)$. \square

Owing to the above, we can thus define $J_f \subseteq \text{MA}(S)$ as the collection of points of the medial axis that have a closest point in the open face f .

Remark 42. The proof of Property 2.55 also yields an explicit expression of J_f as the convex hull of $c(\mathcal{B})$, where \mathcal{B} denotes the unique cardinal minimum local representation of S in the neighbourhood of f .

Remark 43. Following the previous remark, we know that $f \subseteq s_0(\text{aff}(\mathcal{B}))$ and $J_f \subseteq c(\text{aff}(\mathcal{B}))$. It follows from Corollary 2.4 that $\text{aff}(f)$ and $\text{aff}(J_f)$ are orthogonal, and that $f \star J_f = \bigcup_{u \in f} \{u\} \star J_f = \bigcup_{u \in f} \text{join}(u, S)$ is in fact a join. We call it the *face join* of f in S , and denote by $\text{join}(f, S)$.

Because there are finitely many faces and vertices, there are only finitely many distinct J_u/J_f . Also, note that distinct faces may be supported by the same sphere and hence share a common contribution J_f to the medial axis.




Figure 2.26 – Illustration for the proof of Lemma 2.52, \mathbf{n} and \mathbf{e}' cannot be collinear. If they were collinear, the definition of v would imply $s_0^v \cap \partial S = \{u, v\}$. Because s_0^v supports an open face of ∂S in the neighbourhood of v , this is impossible.

2.5.6 Proof of Theorem 3

In order to conclude the proof of Theorem 3, we now argue that each J_f corresponds to a finite union of faces with the desired properties. To each point in J_f , consider the corresponding ball and interpret it as a $(d+1)$ -dimensional point. For the purpose of notations, we still denote by J_f the resulting subset of points. To prove Theorem 3, we show that the $(d+1)$ -dimensional J_f can be split into finitely many flat faces, and that each face is a closed convex subset of an elliptic pencil. We distinguish three cases depending on the nature of the face f . Let \mathcal{B} be a cardinal minimum local representation of S in the neighbourhood of f , and $C = c(\mathcal{B})$ its collection of centers. Note that such a \mathcal{B} is finite and always exist as per Properties 2.43 and 2.46.

(a) f is an open k -face, $k \geq 1$. As per Remark 42, we know that $f \subseteq s_0(\text{aff}(\mathcal{B}))$ and $J_f = \text{CH}(C)$. Because f is non-empty of dimension at least 1, $\text{aff}(\mathcal{B})$ is an elliptic pencil (of dimension at most $d - 2$). Thus J_f itself is a face with the claimed properties.

(b) f is a non-simple vertex u , by definition $\overset{\circ}{\text{coni}}(u, C) = \emptyset$, $\overset{\star}{\text{coni}}(u, C) = \text{CH}(\{u\} \cup C)$ and $\text{aff}(\mathcal{B})$ is parabolic with dimension at most $d - 1$. Because u is the radical center of $\text{aff}(\mathcal{B})$, $u \in \text{aff}(C)$ and $\text{CH}(\{u\} \cup C)$ is a flat polytope of dimension at most $\dim \text{aff}(C) \leq d - 1$. Let $b_0 = \{u\}$ and $\mathcal{B}_0 = \mathcal{B} \cup \{b_0\}$. Because the radical space of $\text{aff}(\mathcal{B})$ has dimension at least 1, we know that the radical center u is in the image of $\text{mb}_{b, \mathcal{B}_0}$ for all $b \in \text{CH}(\mathcal{B}_0)$. In particular, this is true for all b such that $c(b) \in J_u \subseteq \text{CH}(c(\mathcal{B}_0))$. By Property 2.38, every ball b with $c(b) \in J_u$ contributes a single path-connected spherical patch to the boundary of $\bigcup \mathcal{B}_0$, that connects to the radical sphere $\{u\}$. Because \mathcal{B}_0 is a local representation of S in the neighbourhood of u , this implies that each of these balls actually contribute an open k -face, $k \geq 1$, to the global boundary ∂S . Hence there exists faces f_1, \dots, f_n such that $J_u \subseteq \bigcup_{i=1}^n J_{f_i}$. When it comes to computing the medial axis, we can hence ignore the contribution of non-simple vertices, since it will be covered by some open faces.

(c) f is a simple vertex u , by definition $\overset{\circ}{\text{coni}}(u, C) \neq \emptyset$. Let \mathcal{V} be the set of simple vertices of ∂S , by Remark 41 we have $\text{join}(u, S) = \text{coni}(u, C) \cap \text{vor}(u, \mathcal{V})$. Because $\text{join}(u, S)$ is a convex polytope, the star of u in $\text{join}(u, S)$ is $\text{coni}(u, C) \cap \text{v}\hat{\text{or}}(u, \mathcal{V})$. It follows that the link of u in $\text{join}(u, S)$ must be contained in $\text{coni}(u, C) \cap \partial \text{vor}(u, \mathcal{V})$. In particular $\partial \text{vor}(u, \mathcal{V})$ is a finite collection of convex subsets of hyperplanes. For each of these hyperplanes H , there always exists $v \in \mathcal{V}$ such that H is the mediator hyperplane of u and v . It follows that H is in fact the centers' space of the elliptic pencil that admits the 0-sphere $\{u, v\}$ as its radical sphere. Thus J_u is a finite union of faces with the desired properties. This concludes the proof of Theorem 3.

Note that in case (a) when f is an open k -face, $k \geq 1$, it can be proved that J_f is in fact a singular simplex of the dual complex. However, recall from Figure 2.17 that not every singular simplex is part of the medial axis. In case (c) when f is a simple vertex u , it is presently unknown whether J_u can be easily expressed in terms of regular simplices of the dual complex and Voronoi cells of \mathcal{V} . It can be shown that the union of these J_u 's is a subset of the portion of the Voronoi diagram of \mathcal{V} that meets the regular simplices, but we do not know whether the reverse inclusion holds or not.

2.6 Conclusion

As we have seen in this chapter, pencils are families of balls with many properties. By taking advantage of the existence of local representations in a finite union of balls S , we can describe the neighbourhood of every boundary point of S with pencils. In doing so, degenerate configurations of balls in \mathbb{R}^d correspond to parabolic pencils. Handling degeneracies in finite unions of balls can thus be reduced to studying the properties and configurations of parabolic pencils. Due to the many properties of pencils, and their simple definition as flats of \mathbb{R}^{d+1} , this

way of describing finite unions of balls can simplify the study of degenerate configurations. In particular, we have obtained a robust description of both the boundary and the medial axis of finite unions of balls. These descriptions in fact show that there exists a dual relationship between the boundary and the medial axis.

In continuity with the current chapter, this dual relationship could be explored in more depth. We surmise that it could lead to an efficient algorithm to compute the medial axis of a finite union of balls. Based on the results of this chapter, it is indeed possible to compute the medial axis, even in degenerate configurations, though we have yet to explore the best ways to perform this computation.




Figure 2.27 – The plane H (blue-cyan) is parallel to \mathbf{n} and \mathbf{e}' , and goes through c_0 . Its intersection with b_y is a circular arc. The red portion of that arc corresponds to the x_λ 's.




Figure 2.28 – Consider the red simple vertex. The Voronoi diagram of the simple vertices is displayed in blue. Note that the black vertex is non-simple. The two green balls are a local representation of the shape in the neighbourhood of the red vertex. This red vertex thus admits the green area as its inclusion normal cone. The intersection with its Voronoi cell yields the red area, which coincides with the join displayed in the previous Figure 2.20.

(δ, ε) -ball approximation

Contents

3.1 Problem statement	85
3.2 Boundary of the erosion	89
3.2.1 Boundary of the erosion and wavefront	89
3.2.2 Computational description of the wavefront in \mathbb{R}^2	91
3.3 The rational (δ, ε)-ball approximation is in NP	92
3.3.1 Capturing the erosion through an arrangement	93
3.3.2 Checking set inclusion	95
3.4 Rational (δ, ε)-ball approximation is NP-hard	95
3.4.1 Reduction from vertex cover	96
3.4.2 Practical construction	99
3.5 Conclusion	100

Now that we introduced our main tools related to union of balls, we address the problem that we will study for the remainder of this thesis, namely (δ, ε) -ball approximations. Given a set S in \mathbb{R}^d , a (δ, ε) -ball approximation of S is defined as a collection of balls that covers the morphological erosion of S (by a ball of radius ε) and remains inside the morphological dilation of S (by a ball of radius δ). We study the problem of computing such an approximation for certain classes of shapes (precisely shapes that are themselves finite unions of balls). This problem relates to geometric set cover problems but is however different in nature. We shall see that it offers a new framework for simplifying a collection of balls while controlling both the inner and outer distance to the collection. In this chapter, we go over some general properties of these approximations, and in particular prove that computing a (δ, ε) -ball approximation of minimum cardinality is NP-complete already for $d = 2$. In later chapters, we will investigate polynomial algorithms to compute such approximations in some restricted cases.

3.1 Problem statement

§ Context In many applications, shapes are often represented as the union of a finite set of balls. Such a representation is used in biochemistry to model molecules [EK05]; [Caz+14] or in computer graphics to detect collisions between objects [BO04]. Unions of balls have become

ubiquitous representations of shapes, largely due to the existence of provably good conversion algorithms that allow to build them from other representations such as point clouds [AK00], polygonal meshes [CKM99]; [Hub96] or digital shapes [CM07]. Unfortunately, conversion algorithms generally output very large (and sometimes infinite) collections of balls. This may be the case when conversion algorithms aim at representing the shape as the union of its medial balls. As we have seen in Chapter 1, a ball b is said to be *medial* for a given shape if every ball that contains b and is contained in the shape is equal to b . Assuming that balls are closed and the shape is compact, the union of medial balls provides an exact description of the shape, whose size however may be infinite. Recall that the centers of medial balls form the *medial axis*, whose stability and computation (either exact or approximated) have been widely studied; see [ABE09] for a state-of-the-art report. For computationally demanding tasks such as the simulation of physical processes [Fei+15], shape interpolation or shape matching [RF96]; [Cab+09], it is desirable that the collection of balls that describes a shape has a small size while its union still provides an accurate approximation of the shape. The quality of the approximation is usually measured by various quantities such as the Hausdorff distance or the difference in volume between the shape and its representation. In other situations such as collision detection, it may be important that the union of balls satisfies some *geometric constraints*. For instance, we may want the balls to cover the shape or at least a prescribed set of points.

§ Ball approximations In this work, we introduce a novel way of constraining a collection of balls so that its union remains close to the shape. The idea is to force balls to cover a subset of the shape while remaining contained in a superset of the shape, hence allowing balls neither to cover exactly nor to be contained perfectly in the shape. Given a subset $S \subseteq \mathbb{R}^d$ and a real number $r \geq 0$, we recall that the *dilation* of S by a ball of radius r is $S^{\oplus r} = \cup_{x \in S} b(x, r)$ and the *erosion* of S by a ball of radius r is $S^{\ominus r} = \{x \mid b(x, r) \subseteq S\}$. Note that if S is closed, then both $S^{\ominus r}$ and $S^{\oplus r}$ are closed. We shall refer to $S^{\ominus r}$ and $S^{\oplus r}$ as the r -erosion and r -dilation.

Definition 14. Let $S \subseteq \mathbb{R}^d$, $\varepsilon \geq 0$ and $\delta \geq 0$. A collection of balls \mathcal{B} is a (δ, ε) -ball approximation of S if $S^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B} \subseteq S^{\oplus \delta}$.

Let us make a few remarks. Consider a (δ, ε) -ball approximation \mathcal{B} of S ; see Figure 3.1 for an example. Setting $B = \bigcup \mathcal{B}$, we thus have $B \subseteq S^{\oplus \delta}$ and $(S^{\ominus \varepsilon})^{\oplus \varepsilon} \subseteq B^{\oplus \varepsilon}$. It follows that whenever $S = (S^{\ominus \varepsilon})^{\oplus \varepsilon}$, we have $B \subseteq S^{\oplus \delta}$ and $S \subseteq B^{\oplus \varepsilon}$, or equivalently the Hausdorff distance between B and S is smaller than or equal to the maximum of δ and ε , *i.e.* $d_H(S, B) \leq \max\{\delta, \varepsilon\}$. The condition $S = (S^{\ominus \varepsilon})^{\oplus \varepsilon}$ can be seen as some regularity¹ assumption on S . Indeed, the equality $S = (S^{\ominus \varepsilon})^{\oplus \varepsilon}$ is satisfied if every point of S can be covered by a ball of radius ε . In particular, the condition holds whenever medial balls have a radius larger than or equal to ε , which in turn holds whenever the reach of S is larger than or equal to ε . We recall that the reach of S was first introduced by Federer [Fed59] and is the infimum of distances between points in the medial axis of S and points in the complement of S . Denoting the reach of S by $\text{reach}(S)$, we summarize our findings in the following remark:

¹A set S is called r -regular if $S = (S^{\ominus r})^{\oplus r} = (S^{\oplus r})^{\ominus r}$ [Ser82]. In other words, the shape S is r -regular if it does not change neither under a morphological opening nor a morphological closing by a ball of radius r .




Figure 3.1 – A shape (in blue), its ε -erosion (in orange), its δ -dilation (bounded by the dashed blue curve) and a (δ, ε) -ball approximation (the three green balls).

Remark 44. Let \mathcal{B} be a (δ, ε) -ball approximation of S . If $\text{reach}(S) \geq \max\{\delta, \varepsilon\}$ then $d_H(\bigcup \mathcal{B}, S) \leq \max\{\delta, \varepsilon\}$.

Remark 45. S possesses (δ, ε) -ball approximations of finite cardinality whenever S is compact and $\delta + \varepsilon > 0$.

Suppose S is compact and $\delta + \varepsilon > 0$. To build a finite (δ, ε) -ball approximation of S , consider the collection of open balls with radius $\delta + \varepsilon$ centered at points in $S^{\ominus \varepsilon}$. By construction, the collection covers $S^{\ominus \varepsilon}$ which is compact because S is compact, implying that the collection has a finite subcover. By taking the closure of balls in that subcover, we get a collection of closed balls that are contained in $S^{\oplus \delta}$ and thus form a finite (δ, ε) -ball approximation of S .

§ Main result In this section, we are interested in a discrete variant of the following problem:

Problem 1 ((δ, ε)-ball approximation). Given a subset $S \subseteq \mathbb{R}^d$, two non-negative real numbers δ and ε , find a (δ, ε) -ball approximation of S of minimum cardinality.

Our main result is that a discrete decision version of this problem is NP-complete. To state the problem, we call any collection of balls whose centers have rational coordinates and whose radii are rational *rational*.

Problem 2 (Rational (δ, ε) -ball decision problem). Given a subset $S \subseteq \mathbb{R}^d$ which is the union of a rational finite collection of balls, two non-negative rational real numbers ε and δ and an integer $k > 0$, answer the following question: does S have a rational (δ, ε) -ball approximation with cardinality lower than or equal to k ?

We show that this problem is already difficult in dimension 2:

Theorem 5. The rational (δ, ε) -ball approximation problem is NP-complete in \mathbb{R}^2 .

§ **Related works** As mentioned above, any shape S can be exactly described by the union of its medial balls \mathcal{S} . A (δ, ε) -ball approximation of S of minimum cardinality is a collection of balls \mathcal{B} such that $|\mathcal{B}| \leq |\mathcal{S}|$. Indeed, note that, by Definition 14, \mathcal{S} itself is a (δ, ε) -ball approximation of S . Therefore, solving Problem 1 offers a new framework for *simplifying unions of balls* with a control of the distance to the shape. Various algorithms have been proposed to compute simpler collections of balls. A family of algorithms was designed to deal with the sensitivity of the medial axis to small perturbations on the shape S : the general idea is to compute a geometric quantity for each medial ball and prune the balls whose quantity is below a given threshold (see for instance [CL05]; [ABE09]; [CCT11]; [DZ04] to cite only a few). A slightly different approach has been presented in [MGP10] with the notion of scale axis, where balls of the medial axis are first inflated, the medial axis updated (resulting in a simplification) and its balls deflated. Last, hierarchical structures have also been proposed to obtain simpler collections of balls. In the context of collision detection, sphere-trees are built by iteratively merging and/or displacing balls while ensuring that the simpler set of balls includes the boundary of the input one [Hub95]; [BO04]; [BCS09].

The (δ, ε) -ball approximation problem is also closely related to *geometric set cover* problems. Following [AP14] or [HP11], let $\Sigma = (K, R)$ be a finite geometric range space where K is a *finite* set of points in \mathbb{R}^d and R , the range, is a *finite* family of simple shaped regions, like for instance rectangles, or disks. A subset $C \subseteq R$ is a cover of K if any point of K belongs to at least one element of C . Computing a set cover of minimum cardinality is known as the set cover problem. The corresponding decision problem is well-known to be NP-complete, and various approximation algorithms have been proposed, from the simple greedy one [Chv79]; [Joh74], to more convoluted algorithms using ε -nets or multiplicative weights methods [BG95]; [HP11]; [AP14]. Several variations of the geometric set cover problem have also been introduced and investigated, notably when the range R is a family of balls. Instead of minimizing the number of balls used to cover the whole set K , a related problem is to cover the set K at best using a prescribed number of balls (see for instance a greedy algorithm in [Caz+14]). For more applicative purposes, especially in sensor networks, the balls of the range R can be of prescribed centers, but undefined radii, and the goal is to find a cover that minimizes a function of the radii of the balls [RMS13].

The key difference between (δ, ε) -ball approximation and set cover problems lies in the assumption of a finite range space for the latter one. If the shape K to cover and the range R are infinite, another family of results exists for which the range is the (infinite) set of congruent balls of a given radius. Results on that matter go back as far as [Ker39], that establishes a relation between the (Lebesgue) measure of a bounded set of points K in \mathbb{R}^2 and the number of congruent disks of radius r necessary to cover K . When K is a simple geometric shape (disk, square, equilateral triangle), results on the maximum size (radius, side length) of a shape K that can be covered by n unit disks are documented in [T04].

§ **Chapter overview** We establish NP-completeness of the rational (δ, ε) -ball approximation problem (Theorem 5) by successively proving that:

- The rational (δ, ε) -ball approximation problem in \mathbb{R}^2 is in NP, meaning that one can verify in polynomial time that a given rational collection of disks is indeed a solution of the problem. For this, we first provide in Section 3.2 a computational description of the boundary of $S^{\ominus\varepsilon}$ when S is a union of disks. This result is then used in Section 3.3 to design a polynomial time algorithm to check that a given collection of disks covers the ε -erosion $S^{\ominus\varepsilon}$ and is included in the δ -dilation $S^{\oplus\delta}$ using arrangements of circles.
- The rational (δ, ε) -ball approximation problem in \mathbb{R}^2 is NP-hard. We prove this through a reduction from the well-known NP-complete vertex cover problem in Section 3.4.

3.2 Boundary of the erosion

Let \mathcal{S} be a finite collection of disks in \mathbb{R}^2 and ε a non-negative real number. The purpose of this section is to provide a computational description of the boundary of the erosion $(\bigcup \mathcal{S})^{\ominus\varepsilon}$. For this, we introduce in Section 3.2.1 a superset of the boundary called the *wavefront* and characterize which part of the wavefront belongs to the boundary. In Section 3.2.2, we decompose the wavefront into elements that are easy to compute, based on the decomposition of $\partial(\bigcup \mathcal{S})$ into vertices and edges that we derive from Section 2.4.

Note that Section 3.2.1 is valid in any dimension, whereas Section 3.2.2 is specific to \mathbb{R}^2 .

3.2.1 Boundary of the erosion and wavefront

Recall that \mathcal{S} is a finite collection of balls in \mathbb{R}^d and let $S = \bigcup \mathcal{S}$ be the union of those balls. In this section, we relate the boundary of the erosion $\partial S^{\ominus\varepsilon}$ to a (possibly self-intersecting) surface of dimension $d - 1$ that we call the *wavefront* of ∂S at time ε . Roughly speaking, the wavefront at time ε is obtained by moving with speed 1 some of the points on the boundary of S a given distance ε down the inward-pointing normals. Formally, the contribution of a point $u \in \partial S$ to the wavefront at time ε is:

$$w_\varepsilon(u) = \{x \in \partial b(u, \varepsilon) \mid \exists r > 0, b(x, \varepsilon) \cap b(u, r) \subseteq S\}.$$

The wavefront of ∂S is obtained by taking the union of all contributions $w_\varepsilon(\partial S) = \bigcup_{u \in \partial S} w_\varepsilon(u)$ (see Figure 3.2.b). It is not difficult to see that the wavefront contains the boundary of the erosion (see Figure 3.2.a). Indeed, if a point x belongs to the boundary of the erosion $\partial S^{\ominus\varepsilon}$, then it is the center of a ball $b(x, \varepsilon) \subseteq S$ whose boundary meets the boundary of S at some point u . It follows that $x \in \partial b(u, \varepsilon)$ and thus belongs to the contribution $w_\varepsilon(u)$ of point u . The goal of this section is to establish that the boundary of the erosion $\partial S^{\ominus\varepsilon}$ is the set of points on the wavefront $w_\varepsilon(\partial S)$ at distance ε or more to vertices of ∂S (see Figure 3.2.c):




Figure 3.2 – (a) A shape defined as the union of three (blue) disks and its erosion (in orange); (b) wavefront of the boundary (red curve); (c) arrangement of circles used to decompose the erosion into cells (see Section 3.3.1). Three of the circles are centered on the (red) vertices of the boundary. The vertices of the arrangement are the black dots.

Proposition 3.1. *Let \mathcal{V} be the vertex set of ∂S . For $\varepsilon \geq 0$, we have*

$$\partial S^{\ominus \varepsilon} = w_\varepsilon(\partial S) \setminus (\mathcal{V}^{\oplus \varepsilon})^\circ$$

Proof. To prove this result, we relate the join of a point to its contribution to the wavefront. Note that by definition of the wavefront, $w_\varepsilon(u)$ is the collection of centers of balls of radius ε whose boundary goes through u , and that satisfy local inclusion in S around u . Hence using Proposition 2.40 we have the identity below. Refer to Figure 3.3.

$$w_\varepsilon(u) = \text{coni}^*(u, S) \cap \partial b(u, \varepsilon) \tag{3.1}$$

First, we show that $w_\varepsilon(u) \setminus (\mathcal{V}^{\oplus \varepsilon})^\circ = \text{vor}_S(u, \partial S) \cap \partial b(u, \varepsilon)$. When u is not a simple vertex, we know from Property 2.53 that $\text{coni}^*(u, S) = \text{join}(u, S) = \text{vor}_S(u, \partial S)$ and thus $w_\varepsilon(u) = \text{vor}_S(u, \partial S) \cap \partial b(u, \varepsilon)$. It follows that every point of the wavefront contribution of u admits u as a closest point in ∂S , and in particular that they are at distance at least ε from any other vertices of ∂S . Hence $w_\varepsilon(u) \setminus (\mathcal{V}^{\oplus \varepsilon})^\circ = w_\varepsilon(u) = \text{vor}_S(u, \partial S) \cap \partial b(u, \varepsilon)$, and this set of points is at distance precisely ε from ∂S .

If u is a simple vertex, we have $\text{vor}_S(u, \partial S) \cap \partial b(u, \varepsilon) = \text{join}(u, S) \cap \partial b(u, \varepsilon) = \text{coni}^*(u, S) \cap \text{vor}(u, \mathcal{V}) \cap \partial b(u, \varepsilon) = w_\varepsilon(u) \cap \text{vor}(u, \mathcal{V})$ using Property 2.53, and equality (3.1). Now, for $x \in w_\varepsilon(u)$,

$$\begin{aligned} x \in \text{vor}(u, \mathcal{V}) &\iff \forall v \in \mathcal{V}, \|x - u\| \leq \|x - v\| \\ &\iff \forall v \in \mathcal{V}, \|x - v\| \geq \varepsilon \\ &\iff x \notin (\mathcal{V}^{\oplus \varepsilon})^\circ \end{aligned}$$

Thus the equality $w_\varepsilon(u) \setminus (\mathcal{V}^{\oplus\varepsilon})^\circ = \text{vor}_S(u, \partial S) \cap \partial b(u, \varepsilon)$ holds. Again, this set of points is at distance ε from ∂S .

Finally, because $\bigcup_{u \in \partial S} \text{vor}_S(u, \partial S) = S$, we deduce

$$w_\varepsilon(\partial S) \setminus (\mathcal{V}^{\oplus\varepsilon})^\circ = \{x \in S \mid d(x, \partial S) = \varepsilon\} = \partial S^{\ominus\varepsilon}$$

which concludes the proof. □




Figure 3.3 – A shape represented locally around u by two (blue) disks b_1 and b_2 . The contribution of u to the wavefront is the (black) circular arc, closed at x_1 and open at x_2 . By Equality (3.1), it is the intersection of the (orange) affine conical hull and $\partial b(u, \varepsilon)$. The endpoint x_1 belongs to the wavefront (unlike the endpoint x_2) because the disk $b(x_1, \varepsilon)$ is locally included in the shape around u (unlike the disk $b(x_2, \varepsilon)$).

3.2.2 Computational description of the wavefront in \mathbb{R}^2

The goal of this section is to give a computational description of the wavefront in dimension 2. In particular, we shall see that the wavefront is composed of circular arcs; convex circular arcs are contributed by edges (1-faces) of ∂S while concave circular arcs are contributed by vertices of ∂S . To describe those arcs, recall from identity (3.1) that the contribution of a point $u \in \partial S$ to the wavefront at time ε can be expressed as the intersection of $\partial b(u, \varepsilon)$ and the inclusion normal cone of S at point u . We thus inspect for each element on the boundary of S its contribution to the wavefront.

§ Boundary points of S Based on Section 2.4, we list the various types of boundary points that exist in dimension 2. Consider $u \in \partial S$, it is either part of an open 1-face, usually

referred to as an edge, or a non-simple vertex, or a simple vertex. Refer to Figure 2.16 for an illustration.

- In \mathbb{R}^2 , edges stem from elliptic 0-pencils, that is pencils that only contain one disk of positive radius. It follows by Property 2.46 that the whole edge is supported by the boundary of that unique disk.
- Non-simple vertices correspond to parabolic pencils of dimension 0 or 1. In the former case, it is an isolated singleton in the shape, and in the latter it is the point of tangency between two (non-singleton) disks.
- Finally, any other configuration yields a simple vertex. Note that for simple vertices, any collection of disks that locally describes its neighbourhood in the shape necessarily contains at least two distinct (non-singleton) disks.

§ Contribution of edges Let e be an edge of ∂S . As observed above, there is a disk $b \in \mathcal{S}$ whose boundary supports e and which represents S locally around each point $u \in e$. Letting c be the center of b and plugging in identity (3.1) the equality $\text{coni}^*(u, S) = [uc]$ yields $w_\varepsilon(u) = [uc] \cap \partial b(u, \varepsilon) = [uc] \cap \partial b^{\ominus \varepsilon}$ for all $u \in e$. It follows that the contribution $w_\varepsilon(e) = \bigcup_{u \in e} w_\varepsilon(u)$ of an edge e is either empty if b has a radius smaller than ε or is the convex circular arc obtained by projecting e orthogonally onto $b^{\ominus \varepsilon}$ otherwise.

§ Contribution of vertices Assume that v is a vertex of ∂S . If v is a simple vertex, then $\text{coni}^*(v, S)$ has non empty interior, and $w_\varepsilon(v)$ is a concave circular arc. If v is an isolated singleton, then $\text{coni}^*(v, S) = \{v\}$, and its contribution to the wavefront is empty for $\varepsilon > 0$. Finally in the special case when \mathcal{S} has two disks b_1 and b_2 meeting tangentially at v and whose pair $\mathcal{B} = \{b_1, b_2\}$ represents locally S around v , then $w_\varepsilon(v) = [c_1 c_2] \cap \partial b(v, \varepsilon)$, where c_i is the center of b_i for $i \in \{1, 2\}$. It follows that $w_\varepsilon(v)$ is composed of zero, one or two degenerate circular arcs (each reduced to a point), depending on whether zero, one or two disks among b_1 and b_2 have radii larger than or equal to ε .

§ Boundary of the erosion The boundary of the erosion $\partial S^{\ominus \varepsilon}$ being a subset of the wavefront, it is also composed of circular arcs supported by the set of circles (see Figure 3.2.c):

$$\mathcal{C}_\varepsilon(\mathcal{S}) = \{\partial b(v, \varepsilon) \mid v \in \mathcal{V}\} \cup \{\partial b^{\ominus \varepsilon} \mid b \in \mathcal{S} \text{ and } \text{radius}(b) \geq \varepsilon\}.$$

3.3 The rational (δ, ε) -ball approximation is in NP

To check whether a finite collection of disks \mathcal{B} is a (δ, ε) -ball approximation of a finite union of disks $\bigcup \mathcal{S}$, we have to test the two inclusions $(\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B} \subseteq (\bigcup \mathcal{S})^{\oplus \delta}$ in polynomial

time. To do so, we build on results of Section 3.2 and describe a method that given the triplet $(\mathcal{S}, \varepsilon, \mathcal{B})$ returns true if and only if the first inclusion $(\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B}$ holds. To test the second inclusion, it suffices to apply the same method to the triplet $(\mathcal{B}, 0, \{s^{\oplus \delta} \mid s \in \mathcal{S}\})$.

Before describing our method in Section 3.3.2, we explain in Section 3.3.1 how to encode the erosion (or union) of a finite union of disks, using cells of an arrangement of circles.

3.3.1 Capturing the erosion through an arrangement

Recall that an arrangement of circles is the decomposition of \mathbb{R}^2 into open cells induced by the circles. Cells in the arrangement are either vertices, edges or faces. Such arrangements have a size quadratic in the number of circles and can be computed in polynomial time [Pac12]. In this section, we let $\mathcal{A}_\varepsilon(\mathcal{S})$ be the arrangement formed by circles in $\mathcal{C}_\varepsilon(\mathcal{S})$ and show that the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$ can be captured by cells in $\mathcal{A}_\varepsilon(\mathcal{S})$. This means that we can label cells in $\mathcal{A}_\varepsilon(\mathcal{S})$ as IN or OUT in such a way that cells with label IN partition the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$. We shall see that labeling the arrangement can be done in polynomial time. Furthermore, the arrangement can be computed in an exact manner [WZ06] because all circles involved in the construction are rational (center coordinates and radii are rational numbers).

We adopt the convention that each time a circle in $\mathcal{C}_\varepsilon(\mathcal{S})$ has radius zero, it counts as a vertex in the arrangement $\mathcal{A}_\varepsilon(\mathcal{S})$. We already know from the previous section that the boundary of the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$ is supported by circles in the collection $\mathcal{C}_\varepsilon(\mathcal{S})$. Hence, the edges and vertices of $\mathcal{A}_\varepsilon(\mathcal{S})$ capture the boundary of the erosion. It follows that every cell of $\mathcal{A}_\varepsilon(\mathcal{S})$ is either inside or outside of the erosion. We establish the following result:

Lemma 3.2. *We can label in polynomial time $\mathcal{A}_\varepsilon(\mathcal{S})$ in such a way that cells with label IN partition the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$.*

To label cells in the arrangement, we need the following technical lemma:

Lemma 3.3. *Let e be an edge of $\mathcal{A}_\varepsilon(\mathcal{S})$ that supports the boundary of the erosion. Then e is adjacent to a face in the erosion and a face in its complement.*

Proof. Let $S = \bigcup \mathcal{S}$ and $x \in e$. By Proposition 3.1, there is $u \in \partial S$ such that $x \in w_\varepsilon(u) \setminus (\mathcal{V}^{\oplus \varepsilon})^\circ$. Let z_t be the point on half-line $[ux)$ such that $\|z_t - u\| = t$ (see Figure 3.4). To establish the lemma, we proceed in three steps:

(1) First, we prove that $z_t \notin (\mathcal{V}^{\oplus t})^\circ$ for all t close enough to ε . Recall that $x \notin (\mathcal{V}^{\oplus \varepsilon})^\circ$ or equivalently $d(x, \mathcal{V}) \geq \varepsilon$. We claim that if we exclude u from the vertex set \mathcal{V} , then the distance from x to the remaining vertices of \mathcal{V} can only be larger, that is, $d(x, \mathcal{V} \setminus \{u\}) > \varepsilon$. Indeed, if this is not the case, then x would be equidistant to u and a vertex $v \in \mathcal{V} \setminus \{u\}$. This would imply that x lies on at least two circles of the arrangement (the one supporting e and another bounding $b(v, \varepsilon)$). Hence, x would be a vertex of the arrangement, contradicting our assumption that x lies on edge e . It follows that we can find $T > \varepsilon$ such that for all $0 \leq t \leq T$,

we have $d(z_t, \mathcal{V} \setminus \{u\}) \geq d(z_T, \mathcal{V} \setminus \{u\}) - \|z_t - z_T\| > T - (T - t) = t$. Hence, $d(z_t, \mathcal{V}) \geq t$ for all $0 \leq t \leq T$ or equivalently $z_t \notin (\mathcal{V}^{\oplus t})^\circ$ for all $0 \leq t \leq T$.

(2) Second, we prove that $z_t \in w_t(u)$ for all t close enough to ε . Let \mathcal{B} a collection of disks that represents S locally around u and $C = c(\mathcal{B})$ the collection of their centers. Recall that by Equation (3.1), $w_t(u) = \text{con}_i^*(u, C) \cap \partial b(u, t)$, hence $x \in \text{con}_i^*(u, C)$. We want to prove that we also have $z_t \in \text{con}_i^*(u, C)$, for t in an open neighbourhood of ε . If $x \in \text{con}_i^*(u, C)$, then for all $t > 0$, $z_t \in \text{con}_i^*(u, C)$. If $x \notin \text{con}_i^*(u, C)$, then there is $c \in C$ such that $x \in [uc]$. Without loss of generality we can assume that $\text{con}_i^*(u, C) \cap [ux] = [uc]$, where $[ux]$ denotes the half-line originating from u that goes through x . We claim that $x \neq c$. Indeed let $b = b(c, r)$ be the disk in \mathcal{S} centered at c . Assume for a contradiction that $x = c$. Then, $r = \varepsilon$ and $b^{\ominus \varepsilon} = \{c\}$, showing that $c = x$ is a vertex of the arrangement and contradicting our assumption $x \in e$. Thus, $x \neq c$ and $r > \varepsilon$. Also for all $0 \leq t \leq r$, we have $z_t \in [uc] \subseteq \text{con}_i^*(u, C)$, hence $z_t \in w_t(u)$.

(3) Steps (1) and (2) imply that for t close enough to ε , we have $z_t \in w_t(u) \setminus (\mathcal{V}^{\oplus t})^\circ$. By Proposition 3.1, this implies that $z_t \in \partial(S^{\ominus t})$ and thus that $d(z_t, \partial S) = t$. In particular, for $t < \varepsilon$, $d(z_t, \partial S) < \varepsilon$ and $z_t \in (S^{\ominus \varepsilon})^c$ while for $t > \varepsilon$, $d(z_t, \partial S) > \varepsilon$ and $z_t \in S^{\ominus \varepsilon}$, which suffices to conclude. \square




Figure 3.4 – Notations for the proof of Lemma 3.3.

We now prove Lemma 3.2.

Proof of Lemma 3.2. We label cells in $\mathcal{A}_\varepsilon(\mathcal{S})$ as follows:

(a) We label as BOUNDARY all cells on the boundary of the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$, using results in Section 3.2. More precisely, we compute first vertices and edges on the boundary of $\bigcup \mathcal{S}$ (building for instance the arrangement $\mathcal{A}_0(\mathcal{S})$). Then, for each element on the boundary of $\bigcup \mathcal{S}$, we compute its contribution to the wavefront. Each contribution is composed of cells and we label the cells belonging to the boundary of the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$ as BOUNDARY.

(b) We label as IN each face covered by a disk $b^{\ominus\varepsilon}$ for $b \in \mathcal{S}$ and label as OUT each face covered by a disk $b(v, \varepsilon)$ for $v \in \mathcal{V}$. If no faces have received a label so far, this means that $S^{\ominus\varepsilon}$ is either empty or a finite collection of singletons. In this case, each singleton is labeled IN, all other cells are labeled OUT and we are done. Otherwise, we continue.

(c) Starting from a face either with label IN or label OUT, we do a traversal of the arrangement $\mathcal{A}_\varepsilon(\mathcal{S})$, stepping from one face to an adjacent face. If we cross an edge labeled BOUNDARY, the face currently visited receives a label different from the face on the other side of the crossed edge. Otherwise, the face currently visited receives the same label as the face on the other side. Lemma 3.3 guarantees that after this step, every face is properly labeled.

(d) To label remaining vertices and edges, we use the fact that the erosion is a closed set. We thus give label IN to every vertex or edge bounding a face labeled IN. Other vertices and edges receive the label OUT.

Clearly, all steps can be done in time polynomial in the size of the arrangement. \square

3.3.2 Checking set inclusion

We explain in this section how to check the inclusion $(\bigcup \mathcal{S})^{\ominus\varepsilon} \subseteq \bigcup \mathcal{B}$ in polynomial time. We compute the arrangement $\mathcal{A}_\varepsilon(\mathcal{S})$ and label its cells as IN or OUT depending on whether they belong to the erosion $(\bigcup \mathcal{S})^{\ominus\varepsilon}$ or not as explained in the previous section. Likewise, we compute a second arrangement $\mathcal{A}_0(\mathcal{B})$ and label its cells as IN or OUT depending on whether they belong to the union $\bigcup \mathcal{B}$ or not. We then overlay the two arrangements. This operation consists in computing the arrangement of the joint collection of circles, and it is possible to do so while inheriting the labels from both input arrangements [Ber+02]; [Wei+]. By inspecting the labels of the overlay, we can thus determine if every cell in the erosion is also in the union and conclude on the inclusion test. All operations can be done in polynomial time. In Appendix A.3, we provide an alternative method that does not use any overlay.

3.4 Rational (δ, ε) -ball approximation is NP-hard

In this section, we establish that the rational (δ, ε) -ball approximation problem in \mathbb{R}^d is already NP-hard for $d = 2$. We prove NP-hardness through a reduction from a variant of the vertex cover problem. Recall that for a graph $G = (V, E)$, a subset $C \subseteq V$ is a vertex cover of G if every edge of E is incident to a vertex of C . Finding a minimum vertex cover is NP-hard, even when restricted to cubic planar graphs, that is, planar graphs of degree at most 3 [GJ77]. We shall perform the reduction from this particular variant. Given a cubic planar graph $G = (V, E)$ and real non-negative numbers δ and ε with $\delta + \varepsilon > 0$, we show how to build a finite collection $\mathcal{S}(G, \delta, \varepsilon)$ of disks in the plane such that G has a vertex cover of cardinality k if and only if $\bigcup \mathcal{S}(G, \delta, \varepsilon)$ admits a (δ, ε) -ball approximation of cardinality $k + L$, where L is an integer depending only on $\mathcal{S}(G, \delta, \varepsilon)$.

3.4.1 Reduction from vertex cover

§ **Gadgets** We build the set of disks $\mathcal{S}(G, \delta, \varepsilon)$ from the triple (G, δ, ε) by converting each vertex of G into a *vertex* gadget and each edge of G into an *edge* gadget. Given a vertex $v \in V$, the vertex gadget associated to v is a singleton containing only one disk of radius $\lambda\varepsilon$: any value of λ in $[0, 1[$ works for our construction, meaning in particular that these disks are not mandatory, but including them make the explanations below clearer. In Figures 3.5, 3.6, 3.8 and 3.9 we choose $\lambda = 0.8$. The center of that disk is called a *ghost* and is denoted g_v . The reason for this is that ghosts won't belong to the ε -erosion of $\bigcup \mathcal{S}(G, \delta, \varepsilon)$. Given an edge $\{x, y\} \in E$, the edge gadget associated to $\{x, y\}$ is a finite sequence of disks of odd length at least three; see Figure 3.5.b. The disks in the sequence all share the same radius ε . The centers of those disks are called *rotulae*. We distinguish two types of rotulae: *Extreme* rotulae are the centers of the two extreme disks in the sequence and are denoted as c_{xy} and c_{yx} . *Regular* rotulae are the centers of other disks. The rotulae of two consecutive disks in the sequence are said to be neighbours of one another. Note that extreme rotulae have only one neighbour and regular rotulae have two neighbours. We say that an extreme rotula c_{xy} and a ghost g_v are *linked* to one another whenever $x = v$. Hence, each extreme rotula is linked to a unique ghost, whereas a ghost may be linked to up to three extreme rotulae depending on the degree of the vertex that generated the ghost. Finally, we define $\mathcal{S}(G, \delta, \varepsilon)$ as the collection of disks resulting from the conversion of vertices and edges into vertex gadgets and edge gadgets respectively.

At this point, we haven't yet specified the centers of disks, nor the number of rotulae per edge that we need (only that it should be an odd number). We postpone to Section 3.4.2 the description of how disks composing gadgets are chosen. This will be done from an orthogonal grid drawing of G ; see Figure 3.5 for an example. At this stage it suffices to know that we can always choose disks composing gadgets so that they fulfill the properties (i)-(viii) listed below. To simplify notations, we shall refer to $\mathcal{S}(G, \delta, \varepsilon)$ and $S(G, \delta, \varepsilon) = \bigcup \mathcal{S}(G, \delta, \varepsilon)$ simply as \mathcal{S} and S .

- (i) The ε -erosion $S^{\ominus\varepsilon}$ is exactly the collection of rotulae.
- (ii) Any disk $b \subseteq S^{\oplus\delta}$ contains at most 2 regular rotulae.
- (iii) If a disk $b \subseteq S^{\oplus\delta}$ contains 2 regular rotulae, then they are neighbours of each other and b does not contain any other rotula, neither regular nor extreme.
- (iv) Let c be a regular rotula, c_1 and c_2 its two neighbours. There exist two disks $b_1, b_2 \subseteq S^{\oplus\delta}$ such that $\{c_1, c\} \subseteq b_1$ and $\{c_2, c\} \subseteq b_2$.
- (v) Any disk $b \subseteq S^{\oplus\delta}$ contains at most 3 extreme rotulae.
- (vi) If a disk $b \subseteq S^{\oplus\delta}$ contains 2 or 3 extreme rotulae, then these extreme rotulae are linked to the same ghost and b only contains extreme rotulae linked to that ghost.
- (vii) Let g be a ghost. There exists a disk $b \subseteq S^{\oplus\delta}$ that contains all extreme rotulae linked to g .

- (viii) If a disk $b \subseteq S^{\oplus\delta}$ contains both a regular rotula and an extreme rotula, then these are neighbours and b does not contain any other rotula, neither regular nor extreme.




Figure 3.5 – Conversion from (a) an orthogonal grid drawing of a graph G with 5 vertices and 5 edges to (b) the collection of disks $\mathcal{S}(G, \delta, \varepsilon)$ with 5 (yellow) ghosts disks, 10 (purple) extreme rotulae and 31 (blue) regular rotulae. Ghost disks have radii $\lambda\varepsilon$ with $\lambda = 0.8$. The δ -dilation of the union of disks is bounded by the dashed blue lines.

§ Canonical coverings Recall that each edge of G is associated to an odd number of rotulae (which are the centers of disks in the associated edge gadget). Given an edge $e \in E$, we ask for the smallest number of disks contained in $S^{\oplus\delta}$ that suffices to cover the rotulae of e . Let $2\ell(e) + 1$ be the number of rotulae in e , where $\ell(e) \geq 1$ is an integer. Since e has two extreme rotulae, the number of regular rotulae is $2\ell(e) - 1$. From property (ii) we need at least $\lceil (2\ell(e) - 1)/2 \rceil = \ell(e)$ disks in $S^{\oplus\delta}$ in order to cover these regular rotulae. By (iv) there always exists a collection of $\ell(e)$ disks in $S^{\oplus\delta}$ which covers these rotulae plus one of the two extreme rotulae of e , and that extreme rotula can be chosen arbitrarily. Indeed, it suffices to

cover pairs of neighbouring rotulae of e with disks in $S^{\oplus\delta}$, while making sure that the extreme rotula to cover and its neighbour are one of these pairs (see Figure 3.6). This gives two possible coverings of the regular rotulae of e (one for each extreme rotulae) which we shall refer to as *canonical* for e . Furthermore, properties (iii) and (viii) guarantee that any disk containing a regular rotula will only contain rotulae belonging to the same edge gadget. Therefore it is necessary and sufficient to use $\ell(e)$ disks in $S^{\oplus\delta}$ to cover the regular rotulae of an edge gadget e , and these $\ell(e)$ disks exclusively cover rotulae of e . However, one of the two extreme rotulae of e will not be covered by these $\ell(e)$ disks and one extra disk is required to cover all rotulae of e for a total of $\ell(e) + 1$ disks. Unlike the previous $\ell(e)$ disks, by (v) this extra disk may be used to cover extreme rotulae of several edges assuming these extreme rotulae are all linked to the same ghost. We define $L = \sum_{e \in E} \ell(e)$ which is the smallest number of disks needed to cover all regular rotulae and which can be achieved by taking for each edge, one of its two possible canonical coverings.




Figure 3.6 – A canonical covering of an edge gadget with two green disks. Same color convention as in Figure 3.5.b.

Property 3.4. *If G has a vertex cover $C \subseteq V$, then S has a (δ, ε) -ball approximation \mathcal{B} with $|\mathcal{B}| = L + |C|$.*

Proof. By (i), we know that any (δ, ε) -ball approximation of S is a collection of disks contained in $S^{\oplus\delta}$ whose union covers the rotulae of all edges of G . To build such a (δ, ε) -ball approximation, we use property (vii) and start by selecting for each vertex $v \in C$ a disk that covers all extreme rotulae linked to the ghost g_v . This gives a set \mathcal{C} of $|C|$ disks that cover at least one extreme rotula per edge. Using the appropriate canonical covering of each edge, that is, the one excluding the extreme rotula already covered by \mathcal{C} , we complete the set \mathcal{C} into a (δ, ε) -ball approximation adding L more disks. \square

Property 3.5. *If S has a (δ, ε) -ball approximation \mathcal{B} , then G has a vertex cover C with $|C| \leq |\mathcal{B}| - L$.*

Proof. Without loss of generality, we may assume that all disks in \mathcal{B} cover at least one rotula. Indeed, if a disk b does not cover any rotula, it can be removed from \mathcal{B} while keeping the property that \mathcal{B} is a (δ, ε) -ball approximation. Starting from \mathcal{B} , we first deduce a (δ, ε) -ball approximation \mathcal{R} of S having the property that it contains one of the two canonical coverings of each edge $e \in E$. For $e \in E$, let $\mathcal{B}_e = \{b \in \mathcal{B} \mid b \text{ contains a regular rotula of } e\}$. Note that for different edges the \mathcal{B}_e are disjoint and that $|\mathcal{B}_e| \geq \ell(e)$. Given an extreme rotula c_{xy} ,

we denote by \mathcal{C}_{xy} the canonical covering that contains c_{xy} and all the regular rotulae of edge $\{x, y\}$. Initializing \mathcal{R} to \mathcal{B} , we modify \mathcal{R} as follows, applying the following transformation to each edge e of G :

- If disks in \mathcal{B}_e cover no extreme rotula, we replace \mathcal{B}_e by one of the two canonical coverings of e .
- If disks in \mathcal{B}_e cover 1 extreme rotula, we replace \mathcal{B}_e by the canonical covering of e that covers this extreme rotula.
- If disks in \mathcal{B}_e contains both extreme rotulae of e , then $|\mathcal{B}_e| \geq \ell(e) + 1$. We choose arbitrarily one of the two extreme rotulae, say c_{xy} , and let b be a disk in $S^{\oplus\varepsilon}$ containing the other extreme rotula c_{yx} but no regular rotula. We replace \mathcal{B}_e by $\mathcal{C}_{xy} \cup \{b\}$.

Each of these substitutions preserves the ε -covering property and does not increase the cardinality of the resulting collection of disks. Consider the disks of \mathcal{R} that do not contain any regular rotula, $\mathcal{C} = \mathcal{R} \setminus (\cup_{e \in E} \mathcal{R}_e)$. By construction, $|\mathcal{C}| = |\mathcal{R}| - L \leq |\mathcal{B}| - L$ and \mathcal{C} covers at least one extreme rotula of each edge. Let C be the set of vertices $v \in V$ whose ghost g_v is linked to an extreme rotula covered by a disk in \mathcal{C} . We claim that C is a vertex cover of G and that its cardinality satisfies $|C| \leq |\mathcal{C}|$. All $b \in \mathcal{C}$ must contain at least one extreme rotula, thus C is empty if and only if \mathcal{C} is empty. In this particular case, G has no edges because otherwise \mathcal{R} would only cover half of the extreme rotulae. It follows that in this case the empty set is a vertex cover of G . Assume now that \mathcal{C} is not empty. By construction, any $b \in \mathcal{C}$ covers only extreme rotulae, which by (vi) are all linked to the same ghost g_v . It is thus possible to map each $b \in \mathcal{C}$ to one vertex $v \in C$, showing that $|C| \leq |\mathcal{C}|$. To see that C is a vertex cover, recall that \mathcal{C} covers at least one of the two extreme rotulae of each edge. The definition of C thus ensures it contains at least one endpoint of each edge. \square

3.4.2 Practical construction

Given as input a graph G and real non-negative numbers δ and ε such that $\delta + \varepsilon > 0$, we build a collection \mathcal{S} fulfilling properties (i) to (viii). To do so, we rely on the following result.

Theorem 6 ([CP95]). There is a linear time and space algorithm to draw a connected at most cubic graph on an orthogonal grid.

Given a drawing of G on an orthogonal grid, we now describe a way to convert it into an appropriate collection of disks; see Figure 3.5.a for an example of such a drawing. To perform the conversion, we fix the size of the grid to $8(\delta + \varepsilon)$ so that we can fit square blocks of size $4(\delta + \varepsilon) \times 4(\delta + \varepsilon)$ as in Figure 3.7. Note that in Figures 3.5 to 3.9, we have $\delta = \varepsilon$.²

²See http://www.gipsa-lab.grenoble-inp.fr/~isabelle.sivignon/recherches_en.html#ballapproximation for Geogebra files with parameterized generic constructions for any values of δ and ε .




Figure 3.7 – Drawing (in blue) of a cubic graph on an orthogonal grid and grid division into blocks. Gray lines represent the grid and dashed red lines are the blocks.

There are two different ways in which the blocks meet the graph drawing: the block either contains one vertex or only a portion of one edge. Blocks containing a vertex only vary depending on the number and layout of incident edges. Similarly, blocks containing a portion of an edge vary depending on whether the edge bends within the block or not. In each case, we convert the graph drawing covered by the block into a set of disks that satisfies properties (i) to (viii). For blocks containing a vertex, see Figure 3.8 for the four subcases. Similarly for edges, we have two subcases. However, recall that edge gadgets must have an odd number of rotulae. To achieve this, we use the fact that every edge necessarily crosses at least one block in a straight line and provide an odd and an even conversion for this type of block. The three block conversions are presented in Figure 3.9.




Figure 3.8 – Block conversions for a vertex of degree (a) 3, (b) 1, (c) 2 in a bend and (d) 2 in a straight line. The red dashed square delimits the block. Same color convention as in Figure 3.5.b.

3.5 Conclusion

In our proof of NP-hardness, we have chosen carefully the set of disks composing the shape, so that its erosion is composed of a finite set of points (the rotulae). It follows that the reduction




Figure 3.9 – Block conversions for (a) a bent edge, (b) an even and (c) odd straight edge. Same color convention as in Figure 3.8.

from vertex cover induces a set cover problem of a finite set of points (the rotulae), using an infinite range of balls (the medial axis of the dilation).

We notice however that it is possible to define a similar construction where the erosion of the shape is an infinite set of points too. Furthermore, a construction where the shape, its erosion and its dilation all have the same homotopy type is also possible. To build such variants, it suffices to tune the radii of disks composing the gadgets, and distances between consecutive disks so that the eight properties (i)-(viii) are fulfilled. This does not change the overall result.

Because the rational (δ, ε) -ball approximation is NP-complete, there is little hope of finding a general polynomial algorithm. However for shapes simple enough, we argue that it is still possible to achieve polynomial time and space complexity. We present in Chapter 4 an algorithm that works with cycle-free finite union of disks, and that can compute $(0, \varepsilon)$ -ball approximations. Building upon the notions introduced in that first algorithm, we then extend it to a wider class of shapes in Chapter 5.

Greedy optimal algorithm for $(0, \varepsilon)$ -ball approximation

Contents

4.1 Preliminary notions	104
4.2 Projection on the medial axis	107
4.3 Testing for set inclusion in finite unions of balls	109
4.3.1 Technical results	110
4.3.2 Proof of Lemma 4.3	114
4.4 Critical balls	115
4.4.1 Fundamental properties of T -small/large components	115
4.4.2 Critical ball for a point	118
4.4.3 Critical ball for a set	121
4.5 Critical balls and T-maximal candidates	127
4.5.1 Critical balls are candidates	129
4.5.2 Critical balls are T -maximal candidates	133
4.5.3 T -maximal candidates are critical for $A(S, \varepsilon, \mathcal{B})$	133
4.6 Computing critical balls	135
4.6.1 Reduction to elliptic pencils	136
4.6.2 Simplifying the constraint set $A(S, \varepsilon, \mathcal{B})$	139
4.6.3 Reducing to primitive constraints	142
4.7 Algorithm	142
4.7.1 Specification	143
4.7.2 Complexity	143
4.7.3 Optimality	145
4.8 Conclusion	147

In this chapter, we describe an algorithm for computing a cardinal minimum solution to certain instances of the (δ, ε) -ball approximation in \mathbb{R}^2 . As the problem is NP-complete, we make some simplifying assumptions. We suppose that $\delta = 0$ and $\varepsilon > 0$. We also restrain the class of shapes we consider. Specifically, we investigate finite unions of disks that do not contain any cycle in their medial axis. As per [Lie04], it is well-known that open bounded

subsets of \mathbb{R}^d have the same homotopy type as their medial axis, and it follows that the interior of these shapes are also cycle-free. However we here consider unions of closed disks, hence situation as in Figure 4.6 can arise, and the shape itself may in fact contain a cycle despite its medial axis being cycle-free.

Our main result is as follow.

Theorem 7. Under the Real-RAM model, there is a polynomial-time algorithm to compute $(0, \varepsilon)$ -ball approximations of minimum cardinality for finite unions of disks whose medial axes are cycle-free.

Remark 46. Most algorithms and data structures in computational geometry are described in the Real-RAM model [Tó4]; [BY95]. This model assumes that any real number can be represented by a single memory cell, and that the usual arithmetic operations also take constant time.

Throughout this chapter, S will denote a finite union of disks whose medial axis is cycle-free. In Section 4.1, we present the intuition behind our proposed algorithm. This intuition however fails to provide a suitable algorithmic solution. We thus introduce several tools and properties in Sections 4.2 and 4.3 that we use to define and study the properties of critical balls in Sections 4.4 through 4.6. We finally describe our algorithm in Section 4.7.

4.1 Preliminary notions

Because S is compact, any ball $b \subseteq S$ is always contained within some (inclusion-wise) maximal ball of S . It follows that S always admits a cardinal-minimum $(0, \varepsilon)$ -ball approximation that only contains medial balls. Our strategy thus consists in exploring the family of medial balls of S in order to find good medial balls to include in the approximation. In order to efficiently search the family of medial balls, we introduce a partial order over the medial axis and balls of S . From this partial order, we introduce several definitions that will be central to our algorithm.

§ **Partial ordering on medial balls** We already know from Chapter 2 that the medial axis of a finite union of balls can be seen as a finite collection of convex subsets of elliptic pencils. In \mathbb{R}^2 , it is thus a collection of line segments. $\text{MA}(S)$ being a collection of segments, it can be viewed as the embedding of a graph in \mathbb{R}^2 . By assumption on the class of shapes considered, $\text{MA}(S)$ is cycle-free, hence it is a forest. Since we can process each tree of the forest independently, we assume without loss of generality that $\text{MA}(S)$ is a tree. By picking any point of $\text{MA}(S)$ as a root, we obtain an orientation of $\text{MA}(S)$ which induces a partial order on $\text{MA}(S)$. Indeed, we simply have to orient all the edges of $\text{MA}(S)$ from the leaves to the root. We denote by T the resulting oriented tree. The structure represented by T is at times called anti-arborescence or in-tree, and can also be viewed as a directed acyclic graph with a unique sink. For any $x, y \in \text{MA}(S)$, we say that x is T -smaller than or equal to y ,

and note $x \leq_T y$, if y belongs to the unique path from x to the root of T . We also use the usual order symbols and notions such as being T -larger or equal to, \geq_T , or also being strictly T -smaller, $<_T$.

Note that this T -order is valid for all points of $\text{MA}(S)$, and not simply vertices of T . Because points of $\text{MA}(S)$ are centers of medial balls of S , this T -order extends to medial balls. Specifically, we can T -compare two medial balls of S , but also a medial ball of S with a point of $\text{MA}(S)$.

Since T only induces a partial order, we say that two balls that cannot be ordered by T are T -incomparable. An additional useful notion is that of T -maximal ball for a collection: given a collection of balls \mathcal{B} , $b \in \mathcal{B}$ is T -maximal in \mathcal{B} if for all $b' \in \mathcal{B}$, either $b' \leq_T b$ or b and b' are T -incomparable. Likewise, b is T -minimal in \mathcal{B} if for all $b' \in \mathcal{B}$, either $b' \geq_T b$ or b and b' are T -incomparable.

Finally, we introduce the notion of degree for any point $c \in \text{MA}(S)$. If $c \in \text{MA}(S)$ is a node of T , we denote by $\deg c$ the usual degree of c from graph theory, that is the number of neighbours of c in T . If $c \in \text{MA}(S)$ is not a node of T but an inner edge point, we adopt the convention that c has degree $\deg c = 2$.

Remark 47. Let $c \in \text{MA}(S)$, then $\text{MA}(S) \setminus \{c\}$ has $\deg c$ (path-)connected components.

As a closing remark, note that the existence of a T -order as introduced above only requires the medial axis to be acyclic. The notion can thus be adapted to shapes other than finite unions of balls. Note also that whenever we consider a shape to have a T -order on its medial balls, we implicitly assume its medial axis to be acyclic.

§ Principle of the algorithm For our proposed algorithm, the partial ordering we introduced allows the definition of clear start and end points, as well as a measure of progress. Indeed, let $b = b(c, r)$ be a medial ball of S . Its center c splits $\text{MA}(S)$ into $\deg c$ connected components. We denote these components by $\text{branch}(S, c, i)$, $1 \leq i \leq \deg c$. For our purpose, we want to express the domain covered by balls centered at points of these components of $\text{MA}(S)$, but not covered by b . First we define the related collection of medial balls, $\mathcal{C}(S, b, i) = \{b' \text{ medial in } S \mid c(b') \in \text{branch}(S, c, i)\}$. Then the domain of each $\mathcal{C}(S, b, i)$ is $C(S, b, i) = \bigcup \mathcal{C}(S, b, i) \setminus b$. With these notations, b also splits S into the different $C(S, b, i)$'s, and $S \setminus b = \bigcup_{i=1}^{\deg c} C(S, b, i)$. Unless c is the root of T , one of these domains corresponds to balls T -larger than or T -incomparable to b . The other $(\deg c - 1)$ domains correspond to balls T -smaller than b . To promote clarity, we refer to the former domain simply as the T -large component for b in S , and denote it by $C(S, b, +)$. As for the later domains, we refer to their union as the T -small component for b in S and note $C(S, b, -)$. Hence, we have $S \setminus b = C(S, b, +) \cup C(S, b, -)$. See Figure 4.1. From the definition, we also deduce the following:

Property 4.1. *If $b_1 \leq_T b_2$, then $C(S, b_1, -) \cup b_1 \subseteq C(S, b_2, -) \cup b_2$ and $b_1 \cup C(S, b_1, +) \supseteq b_2 \cup C(S, b_2, +)$.*




Figure 4.1 – T -large and T -small components of a medial ball. Green segments are points T -smaller than c , blue ones are points T -larger, black ones are T -incomparable. In (a), the T -small component is the domain of S that was swept from the leaf to b . In (b), that interpretation requires sweeping two branches simultaneously.

Now assume we want to traverse and sweep S with a medial ball, starting from a leaf of T , toward its root. When we reach a medial ball b , and assuming there was no branching from the starting leaf until b (inclusive), then at that moment, $C(S, b, -) \cup b$ corresponds to the domain of S that was swept by our medial ball, and $C(S, b, +)$ to the domain of S that was not swept by it. Our approach is based on this particular decomposition of S . We want to use a greedy approach to iteratively compute a $(0, \varepsilon)$ -ball approximation of the T -small component $C(S, b, -)$. Because T may have several leaves and thus several branches, it is necessary to extend the above definitions of T -small and T -large components to collections \mathcal{B} of medial balls, while preserving the interpretation that $C(S, \mathcal{B}, -) \cup (\bigcup \mathcal{B})$ is the domain of S already processed, and $C(S, \mathcal{B}, +)$ is the domain of S that has yet to be processed. Refer to Figure 4.1.b for an illustration of a potential issue when sweeping with a single ball. The domain already processed for a collection of balls should thus be the union of the domains already processed by some $b \in \mathcal{B}$. Hence the T -small component for \mathcal{B} in S is $C(S, \mathcal{B}, -) = \bigcup_{b \in \mathcal{B}} C(S, b, -)$. Likewise, the domain that still needs to be processed for \mathcal{B} should be the intersection, over all $b \in \mathcal{B}$, of the domains to be processed for b . Hence the T -large component for \mathcal{B} in S is $C(S, \mathcal{B}, +) = \bigcap_{b \in \mathcal{B}} C(S, b, +)$. Owing to Property 4.1, these definitions emphasize the importance of the T -maximal balls of \mathcal{B} . Let $\mathcal{B}_{T\text{-max}}$ be the collection of these T -maximal balls. Then $C(S, \mathcal{B}, +) = C(S, \mathcal{B}_{T\text{-max}}, +)$ and likewise $C(S, \mathcal{B}, -) = C(S, \mathcal{B}_{T\text{-max}}, -)$.

To formalize the procedure presented above, we require two more definitions.

Definition 15. Let \mathcal{B} be a collection of medial balls in S . We say that \mathcal{B} is a T -small $(0, \varepsilon)$ -ball approximation of S if it covers $S^{\ominus \varepsilon}$ in its T -small component $C(S, \mathcal{B}, -)$, that is if $C(S, \mathcal{B}, -) \cap S^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B}$.

Note that every $(0, \varepsilon)$ -ball approximation is also a T -small $(0, \varepsilon)$ -ball approximation of S . As such, we employ the term *partial T -small $(0, \varepsilon)$ -ball approximation* if we need to distinguish from complete $(0, \varepsilon)$ -ball approximations.

Definition 16. Let \mathcal{B} be a partial T -small $(0, \varepsilon)$ -ball approximation of S , and b be medial in S . We say that b is a *candidate* ball with respect to \mathcal{B} , if $\mathcal{B}' = \mathcal{B} \cup \{b\}$ is also a T -small ε -covering of S , and $S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}' \subsetneq S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}$.

The strict inclusion $S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}' \subsetneq S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}$ ensures that \mathcal{B}' is closer to being a complete $(0, \varepsilon)$ -ball approximation than \mathcal{B} . Hence, for any partial T -small $(0, \varepsilon)$ -ball approximation, iteratively adding a candidate to the collection ensure that at some point we will obtain a complete $(0, \varepsilon)$ -ball approximation. Because any partial T -small $(0, \varepsilon)$ -ball approximation always has an infinity of candidates, we choose to add only T -maximal candidates, that is T -maximal among the collection of candidates.

§ **Chapter overview** The approach presented above relies on the fairly intuitive idea of T -maximal candidates to progressively build up an approximation. The main difficulty here lies in the ability to compute such T -maximal candidates. From Definition 16, it is indeed difficult to reach a computational description of these T -maximal candidates.

To circumvent this problem, we introduce and study the notion of critical ball in Section 4.4. We show that they coincide with T -maximal candidates in Section 4.5, and that their definition is suited to an algorithmic computation in Section 4.6, which enables us to state our algorithm in Section 4.7. In order to properly define the notion of critical balls and prove their properties, we need several technical tools and intermediary results. One major difficulty in directly adapting Definition 16 into an algorithm is to ensure that the collection of balls satisfies the desired inclusion properties. We provide in Section 4.3 conditions that guarantee inclusion in finite unions of balls. Because we focus on medial balls, we introduce in Section 4.2 a natural way to map balls contained in S to medial balls of S . This mapping will prove useful to explore the family of medial balls.

4.2 Projection on the medial axis

One easy manipulation when dealing with collection of balls that are contained in a compact shape, is to substitute every ball in the collection by an (inclusion-wise) maximal ball of that shape, so that the maximal ball contains the original ball. This is a form of projection onto the family of maximal balls.

For a finite union of balls S , we present here a continuous projection for any ball $b \subseteq S$ such that $c(b) \in \mathring{S}$. This mapping is based on the face join decomposition of S (Figure 4.2), and inspired by [Ede93]. Indeed, after defining his joins in [Ede93], Edelsbrunner goes on to define an explicit, (continuous,) deformation retraction from the interior of a finite union of balls to its dual shape. By directly adapting his approach, we also obtain a deformation retraction from \mathring{S} to $\text{MA}(S)$, and thus a continuous mapping $\hat{\pi} : \mathring{S} \rightarrow \text{MA}(S)$. We then extend this mapping into a continuous function π that maps balls of S onto medial balls of S .

Formally, consider $x \in \mathring{S}$. If $x \in \text{MA}(S)$, we let $\hat{\pi}(x) = x$ and we are done. Otherwise x




Figure 4.2 – Example of decomposition into face joins. The medial axis is represented by the bold black line segments and points. The face joins only overlap at points of the medial axis. The non-medial boundary edges in the interior only belong to the join of the corresponding colour.

admits one unique closest point on the boundary ∂S . Let u_x be that point. Because $x \in \overset{\circ}{S}$, consider \mathbf{e}_x the unit vector from u_x to x . Recall that J_{u_x} is the collection of points of $\text{MA}(S)$ that admits u_x as a closest point in ∂S . The properties of joins guarantee that the half-line $u_x + \mathbb{R}^+ \mathbf{e}_x$ will intersect the medial axis of S at some point in J_{u_x} . We let $\hat{\pi}(x)$ be that point of intersection. See Figure 4.3 for an illustration. The function $\hat{\pi}$ is well-defined for every interior points of S . It is easy to see that it is continuous, because it is continuous on every face join.

We can now define the mapping π from $\hat{\pi}$ as follows. Consider $b \subseteq S$ with $c(b) \in \overset{\circ}{S}$. We then let $\pi(b)$ be the medial ball centered at $\hat{\pi}(c(b))$.

$$\pi : \left\{ b \subseteq S \mid c(b) \in \overset{\circ}{S} \right\} \rightarrow \{b \text{ medial in } S\}$$

$$b \mapsto \pi(b) \text{ medial in } S \text{ such that } c(\pi(b)) = \hat{\pi}(c(b))$$

If we consider balls of \mathbb{R}^d to be $(d+1)$ -dimensional points, then the mapping π is continuous because $\hat{\pi}$ is. Note that the only balls in S for which π is undefined, are the singleton balls that lie on the boundary of S . Because singleton balls are in essence points, we will often consider the image of interior points of S by π to be well-defined. From now on, we will thus only use π and will not mention $\hat{\pi}$ anymore.

One key property of π is that for any ball b , its image by π will contain b .

Property 4.2. For all $b \subseteq S$ with $c(b) \in \overset{\circ}{S}$, we have $b \subseteq \pi(b)$.

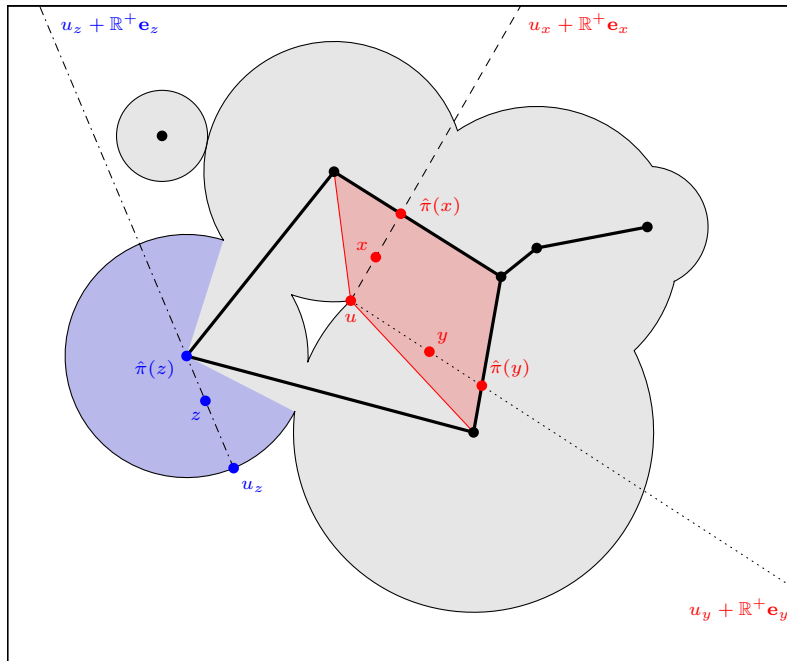


Figure 4.3 – Illustration of the definition of π for various interior points. Note that $u = u_x = u_y$.

Proof. For this proof, refer to Figure 4.4 for an illustration. If b is a medial ball of S , the result is trivial. Suppose that this is not the case. Let $c = c(b)$ be the center of b , and b_c the ball centered at c whose bounding sphere goes through $u_c \in \partial S$. By definition, c belongs to the Voronoi cell of u_c , in the Voronoi diagram of ∂S . Thus we have $b \subseteq b_c$. From the property of joins, we also know that the medial ball centered at $\pi(c)$ contains b_c . Finally, $b \subseteq b_c \subseteq \pi(b)$. \square

4.3 Testing for set inclusion in finite unions of balls

We give here a specialized version of Lemma 1.1 in which one of the sets is a finite union of balls. This result holds in \mathbb{R}^d .

Lemma 4.3. *Let A be a bounded subset of \mathbb{R}^d and consider $X = \bigcup \mathcal{X}$ a finite union of balls. Let \mathcal{V} be the collection of vertices of ∂X . If*

$$(i) \partial A \subseteq X$$

$$(ii) \forall v \in \mathcal{V}, \exists N_v \text{ an open neighborhood of } v \text{ such that } N_v \cap A \subseteq X$$

then $A \subseteq X$.




Figure 4.4 – Illustration for Property 4.2. The blue ball b is contained in the green ball b_c which is contained in the red ball $\pi(b)$.

In order to prove this, we rely on several technical results that we present in Section 4.3.1, while the proof of Lemma 4.3 is deferred to Section 4.3.2.

4.3.1 Technical results

There are two technical results that we will need for the proof of Lemma 4.3. The first one relates the subset $A \setminus X$ to the holes of X given the assumption $\partial A \subseteq X$, and the second investigate the structure of these holes in finite unions of balls.

§ **Holes of a subset X** We start by specifying the definition of a hole.

Definition 17. Let $X \subseteq \mathbb{R}^d$. We call *hole* of X a bounded connected component of X^c .

In general, a subset X may very well have no hole. When X is bounded, then X^c has exactly one unbounded connected component. If that unbounded connected component is not X^c itself, each remaining connected component of X^c is a hole of X .

Lemma 4.4. Let X be a subset and A a bounded subset of \mathbb{R}^d . Assume that $\partial A \subseteq X$, and consider $H = A \setminus X$. Then H is either empty, or a union of holes of X .

Proof. Assume that H is non-empty and let $x \in H$. Let X_x^c be the connected component of X^c that contains x , and $H_x = A \cap X_x^c$. By definition we have $H_x \subseteq X_x^c$. We claim here and

prove in the next paragraph that we also have the reverse inclusion and thus that $H_x = X_x^c$. Hence H_x is a connected component of X^c . Because in addition, $H_x \subseteq A$ is bounded, H_x is a bounded connected component of X^c and is thus a hole of X . It follows that H is either empty, or an union of holes of X .

Let us now prove our claim, that $X_x^c \subseteq H_x$. Let $G = X_x^c \setminus H_x = X_x^c \cap A^c$ and assume by contradiction that $G \neq \emptyset$. We have $X_x^c = G \cup H_x$ (because $H_x \subseteq X_x^c$). Because X_x^c is connected, G and H_x cannot be separated. Thus either $\overline{G} \cap H_x \neq \emptyset$ or $G \cap \overline{H_x} \neq \emptyset$. In both cases, there exists a point $y \in \overline{G} \cap \overline{H_x} \cap X^c$. Because $H_x \subseteq A$, we know that $y \in \overline{A}$. However, $\partial A \subseteq X$ and $y \in X^c$, thus $y \in \overline{A} \setminus \partial A = \overset{\circ}{A}$. It follows that $y \in \overline{G} \cap \overset{\circ}{A} \subseteq \overline{A^c} \cap \overset{\circ}{A} = \emptyset$. This is impossible, hence $G = \emptyset$ and $X_x^c \subseteq H_x$. Thus H_x is indeed a hole of X . \square

§ Holes in finite unions of balls We now provide some insight on the holes of a finite union of balls. In this paragraph, we thus consider $X = S = \bigcup \mathcal{S}$ a finite union of balls. Eventually, we want to show that any hole of S always contains some vertex of ∂S in its boundary. To do so, we prove a series of two lemmas, from which we deduce two corollaries.

Lemma 4.5. *Let $S = \bigcup \mathcal{S} \subseteq \mathbb{R}^d$ be a finite union of balls and $u \in \partial S$. If every open neighbourhood of u meets two locally distinct connected components of S^c , then u is a non-simple vertex of ∂S .*

Proof. Let $\mathcal{B} \subseteq \mathcal{S}$ be a (cardinal-minimum) collection of balls that locally represents S at u and $C = c(\mathcal{B})$ be the collection of their centers. We prove that $\text{coni}(u, C)$ is a hyperplane. This, in turn, entails that $u \in \text{aff}(C)$ and $\dim \text{aff}(C) = d - 1$. \mathcal{B} thus defines a parabolic $(d - 1)$ -pencil, with $\text{coni}(u, C) = \emptyset$. It follows that u is a non-simple vertex. We first specify this hyperplane then show that we have double inclusion.

Let X be a connected component of S^c such that $u \in \overline{X}$. Consider $(x_n) \subseteq X$ a sequence that converges to u and let \mathbf{e}_n be the unit vector such that $x_n = u + \|x_n - u\|\mathbf{e}_n$. The sequence (\mathbf{e}_n) admits some limit point \mathbf{e}_X . We claim that there is some $\rho > 0$, such that for all $0 < h < \rho$, $u + h\mathbf{e}_X \in X$. Indeed let $b \in \mathcal{B}$, we denote its center by c . By applying Lemma 2.27 to b and $b_0 = b(u, 0) = \{u\}$, we know that $b = \{u + \rho\mathbf{e} \mid 0 \leq \rho \leq \rho_{\mathbf{e}}, \|\mathbf{e}\| = 1\}$, with $\rho_{\mathbf{e}} = \langle \mathbf{e}, c - u \rangle + |\langle \mathbf{e}, c - u \rangle|$. Because $x_n = u + \|x_n - u\|\mathbf{e}_n \notin b$, we deduce that $\|x_n - u\| > \rho_{\mathbf{e}_n} \geq 2\langle \mathbf{e}_n, c - u \rangle$. Thus $0 \geq \langle \mathbf{e}_X, c - u \rangle$, and it follows that $u + h\mathbf{e}_X \notin b$. This holds for every $b \in \mathcal{B}$, therefore $u + h\mathbf{e}_X \notin \bigcup \mathcal{B}$. Because \mathcal{B} locally describes S around u , there exists $\rho > 0$ such that $S \cap b(u, \rho) \subseteq \bigcup \mathcal{B}$. If $h \leq \rho$ we thus have $u + h\mathbf{e}_X \in S^c$. In addition, notice that if $\|x_n - u\| \leq h$, then the line segment $\{u + h\mathbf{e}_n \mid \|x_n - u\| \leq h \leq \rho\}$ is path-connected to $x_n \in X$ and thus lies in X . We deduce that $u + h\mathbf{e}_X$ is in the closure of X , since it is the limit point to the sequence $(u + h\mathbf{e}_n)_n$. Because S is closed, S^c is open. Necessarily, $u + h\mathbf{e}_X \in S^c \cap \overline{X}$ implies that $u + h\mathbf{e}_X \in X$.

By assumption, there exists $Y \neq X$, another connected component of S^c with $u \in \overline{Y}$. Likewise we can build \mathbf{e}_Y such that for some $\rho > 0$, we have for all $0 < h \leq \rho$, $u + h\mathbf{e}_Y \in Y$. We claim that $\mathbf{e}_X = -\mathbf{e}_Y$. Because $X \neq Y$, we have $\mathbf{e}_X \neq \mathbf{e}_Y$. Assume by contradiction that \mathbf{e}_X and \mathbf{e}_Y are linearly independent. For $\lambda \in [0, 1]$, let $\mathbf{e}_\lambda = \lambda\mathbf{e}_X + (1 - \lambda)\mathbf{e}_Y$. We have $0 < \|\mathbf{e}_\lambda\| \leq$

1. Let $b \in \mathcal{B}$ and $c = c(b)$ its center, we have $\langle \mathbf{e}_\lambda, c - u \rangle = \lambda \langle \mathbf{e}_X, c - u \rangle + (1 - \lambda) \langle \mathbf{e}_Y, c - u \rangle \leq 0$. Therefore $u + \rho \mathbf{e}_\lambda \in S^c$. This is however absurd because $\{u + \rho \mathbf{e}_\lambda \mid \lambda \in [0, 1]\}$ is then a path in S^c connecting X and Y . Hence we must have $\mathbf{e}_X = -\mathbf{e}_Y$.

Let H be the affine hyperplane through u with normal \mathbf{e}_X . Clearly we have $C \subseteq H$. Indeed for all $c \in C$, $\langle \mathbf{e}_X, c - u \rangle \leq 0$ and $\langle \mathbf{e}_Y, c - u \rangle \leq 0$, which implies $\langle \mathbf{e}_X, c - u \rangle = 0$. Hence $\text{coni}(u, C) \subseteq H$. We now prove the reverse inclusion. Consider $b \in \mathcal{B}$ with its corresponding center c . Because $c \in H$, the projection of b onto H is the $(d - 1)$ -ball $b \cap H$, it has the same center and radius as b . Let $\mathcal{B}' = \{b \cap H \mid b \in \mathcal{B}\}$. The projection of S on H is thus locally equal to $(\bigcup \mathcal{B}) \cap H = \bigcup \mathcal{B}'$. We show that in a small enough neighbourhood of u , the projection of S onto H is H itself. This implies that every $(d - 1)$ -ball supported by H , and whose boundary goes through u , is locally contained in S around u . By Lemma 2.42 we thus have that $H \subseteq \text{coni}(u, C)$ which concludes the proof.

Consider $x = u + \rho \mathbf{e}_X$ and $y = u - \rho \mathbf{e}_X$, $x, y \in S^c$. Because S^c is open, there is $h_0 > 0$ such that $b(x, h_0) \cup b(y, h_0) \subseteq S^c$. Let \mathbf{e} be any unit vector orthogonal to \mathbf{e}_X and h , $0 \leq h \leq h_0$. Consider $\gamma = [xy] + h\mathbf{e}$ the translation of segment $[xy]$ by $h\mathbf{e}$. The endpoints of γ belong to X and Y , hence γ is a path connecting two distinct connected components of S^c and must intersect S . By definition of \mathbf{e} , the projection of γ onto H is the singleton $\{u + h\mathbf{e}\}$. Because $\gamma \cap S \neq \emptyset$ we deduce that $u + h\mathbf{e} \in S \cap H$. This holds for every \mathbf{e} orthogonal to \mathbf{e}_X and $0 \leq h \leq h_0$, therefore $H \cap b(u, h_0) \subseteq H \cap S$.

For any $z \in H$, denote by b_z the ball centered at z whose bounding $(d - 1)$ -sphere goes through u . We have $b_z \cap H \cap b(u, h_0) \subseteq H \cap S$. Because this holds for all $z \in H$, Lemma 2.42 implies $H \subseteq \text{coni}(u, C)$ and thus $H = \text{coni}(u, C)$. Point u necessarily belongs to $\text{CH}(C)$ and it follows that $\text{aff}(C) = \text{coni}(u, C) = H$. The pencil defined by \mathcal{B} thus has dimension $d - 1$. In addition, u belongs to its radical sphere and also lies in the centers' space $\text{aff}(C)$, therefore the pencil is parabolic. Hence, $\text{coni}(u, C) = \emptyset$ and u is a non-simple vertex of ∂S . \square

Consider H , a hole of S . Note that by Lemma A.2 and A.3 (see Appendix) we know that $\partial H \subseteq \partial S$. We can thus use the structure on ∂S to describe that of ∂H .

Lemma 4.6. *Let $x \in \mathbb{R}^d$ and the function $f : y \in \mathbb{R}^d \mapsto \|x - y\|$. Let S be a finite union of balls with a hole H , and consider $f|_{\overline{H}}$ the restriction of f to \overline{H} . Given these assumptions, then $f|_{\overline{H}}$ only achieves its maximum at vertices of ∂S .*

Proof. We proceed by contradiction. Let $u \in \overline{H}$ such that $f|_{\overline{H}}(u)$ is a maximum of $f|_{\overline{H}}$, and assume (by contradiction) that u is not a vertex of ∂S . First, note that we necessarily have $u \in \partial H$. Let \mathcal{S} be a finite collection of balls such that $S = \bigcup \mathcal{S}$ and consider $\mathcal{B} = \{b \in \mathcal{S} \mid u \in b\}$. \mathcal{B} locally describes S around u . We can derive a pencil $\mathcal{P} = \text{aff}(\mathcal{B})$ from \mathcal{B} , its radical sphere $s_0 = s_0(\mathcal{P})$ contains u . By assumption u cannot be a vertex, hence s_0 must have positive radius and have dimensionality at least 1. In other words, s_0 is at least a circle. If we restrict f to $s_0 \cap \partial H$, u still achieves a maximum. It follows that either u is the global maximum of $f|_{s_0}$, or belongs to the boundary of $s_0 \cap \partial H$. We claim that the latter case is impossible. Indeed if u is on the boundary of $s_0 \cap \partial H$, then there must exist a ball

$b \in \mathcal{S}$ such that $s_0 \not\subseteq \partial b$ and $u \in s_0 \cap \partial b$. By definition of \mathcal{B} , we must however have $b \in \mathcal{B}$ and hence $s_0 \subseteq b$, which contradicts the above.

$f(u)$ is a maximum of f over the restricted domain s_0 . Let c_0 be the center of s_0 and $r_0 > 0$ its radius. Consider \mathbf{e} the unit vector from c_0 towards u , we have $u = c_0 + r_0\mathbf{e}$. For $h > 0$ let $y_h = c_0 + (r_0 + h)\mathbf{e}$. Refer to Figure 4.5 for an illustration in \mathbb{R}^2 . We claim that $\|x - y_h\| > \|x - u\|$, and that if h is small enough then $y_h \in S^c$. By Lemma 4.5, such points y_h must belong to H . Indeed, if y_h belonged to another connected component of S^c for any arbitrary small h , this lemma would imply that u is a vertex, which contradicts our initial assumption. There is thus a neighbourhood of u in which $y_h \in H$, but this is impossible because it contradicts the definition of u as a maximum, which concludes the proof.

Remains to prove our two claims, that is $\|x - y_h\| > \|x - u\|$ and $y_h \in S^c$ for h small enough. Let $\mathcal{N} = \{\mathbf{n} \mid c_0 + r_0\mathbf{n} \in s_0\}$ be the collection of unit directions from the center c_0 to points of the sphere s_0 , and note that $\mathbf{e} \in \mathcal{N}$. For $\mathbf{n} \in \mathcal{N}$ we have $\|c_0 + r_0\mathbf{n} - x\|^2 = \|c_0 - x\|^2 + r_0^2 + 2r_0\langle c_0 - x, \mathbf{n} \rangle$. Because u achieves the maximum of $f|_{s_0}$, necessarily $\max_{\mathbf{n} \in \mathcal{N}} \langle c_0 - x, \mathbf{n} \rangle = \langle c_0 - x, \mathbf{e} \rangle$. In particular, if $\mathbf{n} \in \mathcal{N}$ then $-\mathbf{n} \in \mathcal{N}$ and it follows that $\langle c_0 - x, \mathbf{e} \rangle \geq 0$. Because $y_h = u + h\mathbf{e}$, we have

$$\begin{aligned} \|y_h - x\|^2 &= \|u - x\|^2 + h^2 + 2h\langle u - x, \mathbf{e} \rangle \\ &= \|u - x\|^2 + h^2 + 2hr_0 + 2h\langle c_0 - x, \mathbf{e} \rangle \\ &\geq \|u - x\|^2 + h(h + 2r_0) \\ &> \|u - x\|^2 \end{aligned}$$

Hence $\|x - y_h\| > \|x - u\|$. By definition, \mathbf{e} is a direction in the radical space of the pencil $\mathcal{P} = \text{aff}(\mathcal{B})$, therefore $y_h \notin \bigcup \mathcal{B}$. Since \mathcal{B} locally represents S in the neighbourhood of u , for h small enough we obtain $y_h \in S^c$ and our two claims are thus true. \square

We can finally state the property we needed for the proof of Lemma 4.3 as a corollary.

Corollary 4.7. *Let H be a hole of S . Then ∂H contains a vertex of ∂S .*

We also state another corollary of Lemma 4.6 that will be useful later on in Section 4.5.

Corollary 4.8. *Let H be a hole of $S = \bigcup \mathcal{S}$. Let $b \in \mathcal{S}$ a ball such that it contributes to the boundary of H , $\partial b \cap \partial H \neq \emptyset$. Then, there exists a vertex v of ∂S on the boundary of H that does not lie on ∂b , $v \in \partial H \setminus \partial b$.*

Proof. Denote by \mathcal{V} the vertex set of ∂S . By contradiction, suppose that every vertices in ∂H lie on ∂b , $\mathcal{V} \cap \partial H \subseteq \partial b$. Let c be the center of b and r its radius. Consider the function $f : y \in \mathbb{R}^d \mapsto \|c - y\|$. As per Lemma 4.6, $f|_{\overline{H}}$ reaches its maximum in \mathcal{V} . Let $u \in \overline{H}$ be such a maximum point, and recall that necessarily $u \in \partial H$. It follows that $u \in \mathcal{V} \cap \partial H \subseteq \partial b$, and hence $\|c - u\| = r$. Because u is a maximum of $f|_{\overline{H}}$, thus for all $y \in \overline{H}$, $\|c - y\| \leq r$ and $y \in b$. This implies $H \subseteq b \subseteq S$, which is absurd. \square




Figure 4.5 – Illustration for the proof of Lemma 4.6 in \mathbb{R}^2 when s_0 is (a) a 1-sphere or (b) a 0-sphere. For the various (red) points x_i , their furthest point in s_0 is u . For $h > 0$, the (blue) point y_h is always further away from x_i than u .

4.3.2 Proof of Lemma 4.3

We can finally prove Lemma 4.3.

Proof of Lemma 4.3. Let $H = A \cap X^c$. The scheme of the proof is similar to the proof of Lemma 1.1, with only one extra stage in the middle:

- (a) First, we prove that $\partial H \subseteq \partial X$. For this, we proceed as in the first stage of the proof of Lemma 1.1.
- (b) Second, we prove that H is either empty or an union of holes of X . This is Lemma 4.4.
- (c) Last, we prove that $H = \emptyset$. For this, we proceed as in the last stage of the proof of Lemma 1.1, with the only difference being that we pick a point x on ∂H among vertices of ∂X . This is always possible thanks to Corollary 4.7. \square

In essence, the proof above exploits the structural properties of holes in finite unions of balls, to reduce the set of boundary points in ∂X that have to be examined to check whether the inclusion holds. Indeed, steps (a) and (b) only assume that A is bounded, and X is non-empty. Extending this result (and proof) to other classes of containing shapes X can thus be done by investigating the nature of the holes for those classes of shapes.

4.4 Critical balls

For our algorithm, we intuitively want to compute the T -maximal candidates of a partial T -small $(0, \varepsilon)$ -ball approximation \mathcal{B} . In order to reach a computational characterization of these T -maximal candidates, we introduce here the notion of critical ball (Definitions 18 and 19). A critical ball in S is defined by a constraint subset that lies in \mathring{S} . The location of critical balls thus vary depending on the specific constraint used to define them. We show in the next Section 4.5 that it is possible to define a constraint subset from the partial T -small $(0, \varepsilon)$ -ball approximation \mathcal{B} , so that critical balls and T -maximal candidates coincide.

Before giving the formal definition of critical balls, we first have a look at T -small and T -large components of a medial ball. The properties of these components will indeed be useful to prove the properties of critical balls.

4.4.1 Fundamental properties of T -small/large components

In order to prove several properties of critical balls, we will rely heavily on the decomposition of S into $C(S, b, i)$ components. Indeed, we will define these critical balls as T -extremal balls within a specific family of medial balls. As such, the collection of balls T -smaller than the critical ball, and likewise that of balls T -larger or T -incomparable, will naturally arise in the discussion. Knowing some properties of the domain of points they cover, which can be expressed with the $C(S, b, i)$'s, will thus be helpful.

Specifically, we want to establish Proposition 4.10 which states that the $C(S, b, i)$'s are disjoint within \mathring{S} . In general, the $C(S, b, i)$ may not be disjoint themselves because some degeneracies as in Figure 4.6 may occur. The restriction to \mathring{S} enables us to disregard these degeneracies, and simplifies the manipulation of sets involving the $C(S, b, i)$'s.

In order to prove our Proposition 4.10, we first establish a technical lemma. We conclude this section with two corollaries.

Lemma 4.9. *Let S be a finite union of balls whose medial axis is cycle-free. Consider b_- and b_+ two medial balls of S . If $\mathring{b}_- \cap \mathring{b}_+$ is non-empty, then it is included in \mathring{b} , for any medial ball b such that $c(b)$ lies on the path (in $\text{MA}(S)$) from $c(b_-)$ to $c(b_+)$.*

Proof. For this proof, refer to Figure 4.7 for an illustration.

Consider the pencil segment $[b_- b_+]$, it constitutes a path between b_- and b_+ in \mathbb{R}^{d+1} . Because b_- and b_+ are not interior disjoint, the path $c([b_- b_+]) = [c(b_-) c(b_+)]$ lies in the interior of S . As per Section 4.2, that path can be projected onto a path in the medial axis of S , $\pi([b_- b_+])$. Because $\text{MA}(S)$ is cycle-free, any path in $\text{MA}(S)$ connecting $c(b_-)$ to $c(b_+)$ must contain the unique shortest path in $\text{MA}(S)$ from $c(b_-)$ to $c(b_+)$. Therefore for any ball b on that shortest path, there exists a ball $b' \in [b_- b_+]$ such that $\pi(b') = b$.

Let $x \in \mathring{b}_- \cap \mathring{b}_+$. By Property 2.1, we know that for all $b' \in [b_- b_+]$, we have $\text{pow}(x, b') \leq$




Figure 4.6 – Point x lies on the boundary of S , and belongs to both b_- (green) and b_+ (blue). It follows that $C(S, b, -) \cap C(S, b, +)$ is non-empty.

$\max\{\text{pow}(x, b_-); \text{pow}(x, b_+)\} < 0$. Hence $x \in \overset{\circ}{b}'$. Additionally, by the properties of the projection π we know $b' \subseteq \pi(b') = b$. We thus deduce that $x \in \overset{\circ}{b}$. Because this holds for every ball b between b_- and b_+ in $\text{MA}(S)$, and every $x \in \overset{\circ}{b}_- \cap \overset{\circ}{b}_+$, this concludes the proof. \square

In essence, Lemma 4.9 affirms that when we split the shape S by removing a medial ball b , the intersections between balls T -smaller and balls T -greater than b is removed too, meaning that the remaining $C(S, b, i)$ components are nearly disjoint. When restricted to the interior of S , these components are indeed disjoint. However some degeneracies can still occur, as in Figure 4.6, and the $C(S, b, i)$ components may not be disjoint themselves. Formally, we have the proposition below.

Proposition 4.10 (Disjoint interior components). *Let S be a finite union of balls, and b a medial ball in S . If $x \in \overset{\circ}{S}$, then x belongs to exactly one of the three subsets $C(S, b, -)$, b , and $C(S, b, +)$.*

Proof. We already know from the definition of $C(S, b, -)$ and $C(S, b, +)$ that they are both disjoint from b , $C(S, b, -) \cap b = \emptyset = b \cap C(S, b, +)$. Remains to show that x cannot belong to $C(S, b, -) \cap C(S, b, +)$. Assume by contradiction that $x \in C(S, b, -) \cap C(S, b, +)$. First, we argue that for any medial ball of S that contains x , x must lie on the boundary of that medial




Figure 4.7 – The two balls b_1 and b_2 (green) have a non-empty intersection. This intersection is covered by any ball b_3 (blue) on the pencil $[b_1b_2]$, and also by its projection $\pi(b_3)$ (red) on $MA(S)$.

ball. Then we argue that there exists a medial ball that must contain x in its interior, thus leading to a contradiction.

Denote by \mathcal{B}_- the medial balls of S that contain x which are strictly T -smaller than b , and similarly \mathcal{B}_+ the medial balls of S containing x that are strictly T -larger than or T -incomparable to b . Because b and $C(S, b, -) \cup C(S, b, +)$ are disjoint, necessarily $x \notin b$. Thus \mathcal{B}_+ and \mathcal{B}_- are disjoint and their union is exactly the medial balls of S that contain x . In addition because $x \in C(S, b, -)$, there exists a medial ball $b_- \leq_T b$ such that $x \in b_- \setminus b$. We have $b_- \in \mathcal{B}_-$ and thus \mathcal{B}_- is non-empty. Similarly, \mathcal{B}_+ is also non-empty.

Let $b_- \in \mathcal{B}_-$ and $b_+ \in \mathcal{B}_+$. By Lemma 4.9, we have $\overset{\circ}{b}_- \cap \overset{\circ}{b}_+ \subseteq \overset{\circ}{b}$. If $\overset{\circ}{b}_- \cap \overset{\circ}{b}_+$ is non-empty, this implies $b_- \cap b_+ \subseteq b$. Because $x \in b_- \cap b_+$ and $x \notin b$, this is impossible. Therefore, b_- and b_+ are interior disjoint, they are tangent at point x . In particular, we deduce that for all medial balls b_m of S , we must have $x \in \partial b_m$.

Because $x \in \overset{\circ}{S}$, let $r > 0$ such that $b(x, r) \subseteq S$. Let b_m be a medial ball of S such that $b(x, r) \subseteq b_m$. Such a ball always exists because medial balls and inclusion-wise maximal balls coincide for S . We have $x \in \overset{\circ}{b}_m$, which contradicts the property above. Therefore, x cannot

belong to both $C(S, b, -)$ and $C(, S, b, +)$, this concludes the proof. \square

Remark 48. Although we stated the lemma above in terms of $C(S, b, -)$ and $C(S, b, +)$, note that it actually remains valid when considering the (potentially) more numerous $C(S, b, i)$'s.

Proposition 4.10 has several interesting consequences. The most straightforward one is that when restricted to $\overset{\circ}{S}$, (or any subset of $\overset{\circ}{S}$, like $S^{\ominus\varepsilon}$ for instance), the complement of $C(S, b, -)$, b , and $C(S, b, +)$ is the union of the other two subsets. Because situations as in Figure 4.6 may arise in general, we shall make extensive use of this result and refer to Proposition 4.10 when we perform such a complement manipulation.

Proposition 4.10 also enables us to state a refinement of Lemma 4.9 that is not restricted to the interior of balls:

Corollary 4.11. *Let b_- , b , and b_+ be medial balls in S such that $c(b)$ lies on the path from $c(b_-)$ to $c(b_+)$ in $\text{MA}(S)$. If $b_- \cap b_+ \neq \emptyset$, then either $b_- \cap b_+ \subseteq b$, or b_- and b_+ are tangent at a vertex of ∂S . In both cases we have $b_- \cap b_+ \subseteq b \cup \mathcal{V}$, where \mathcal{V} denotes the vertex set of ∂S .*

Remark 49. When b_- and b_+ are tangent at a vertex of ∂S , the tangency point may or may not belong to b .

Proof. If $\overset{\circ}{b}_- \cap \overset{\circ}{b}_+$ is non-empty, we easily deduce $b_- \cap b_+ \subseteq b$ by relying on Lemma 4.9. Otherwise b_- and b_+ are interior disjoint. If these two balls intersect, they must thus be tangent. Denote by u the tangency point of b_- and b_+ . If $u \in b$ we are done, otherwise assume that $u \notin b$. Because $u \in b_- \setminus b$, we have $u \in C(S, b, -)$. Likewise $u \in b_+ \setminus b$ implies $u \in C(S, b, +)$. Hence $u \in C(S, b, -) \cap C(S, b, +)$, by Proposition 4.10 we thus deduce that $u \in \partial S$. Any local representation of S around u must define a pencil that contains both b_- and b_+ . These two balls are however tangent at u , the pencil must therefore be parabolic and u is necessarily a vertex of ∂S . \square

Conjecture 1. Using notations of the last proof, we conjecture that when b_- and b_+ are tangent at a point u , with $u \notin b$, then u is a non-simple vertex.

We proved above that u was necessarily a vertex, and that any local representation of S around u would define a parabolic pencil. In order to obtain the stronger statement that u is a non-simple vertex of S , we would have to prove that these pencils have dimension exactly 1. Intuitively, we would deduce that $\{b_-, b_+\}$ is a local representation of S around their tangency point.

4.4.2 Critical ball for a point

Definition 18 (Critical ball for a point). Consider a point $x \in \overset{\circ}{S}$. Let b be a medial ball in S such that $x \in b$. We say b is $(T-)$ critical for point x in shape S if equivalently:

- (i) b is the T -maximum in the set of balls $\{b' \text{ medial in } S \mid x \in b'\}$

- (ii) if $x \in b'$, b' medial in S , then $b' \leq_T b$
- (iii) if $b' >_T b$, or b' and b are T -incomparable, then $x \notin b'$
- (iv) if $b' >_T b$, then $x \notin b' \cup C(S, b', +)$
- (v) if $x \in b' \cup C(S, b', +)$, then $b' \leq_T b$ or b' and b are T -incomparable

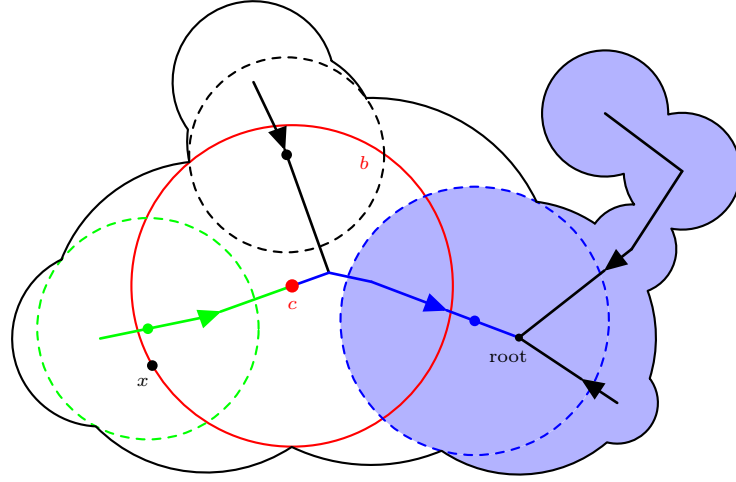


Figure 4.8 – The red ball b is critical for x in S . The green ball contains x but is strictly T -smaller than b , and is not x -critical. The black and blue balls are strictly T -larger than or T -incomparable to b , and do not contain x .

Said otherwise, a ball b is critical for $x \in S$ if it is the furthest maximal ball along the orientation of T that still covers x . See Figure 4.8. The following property is pretty straightforward from the definition of T -larger and T -smaller components.

Property 4.12. *Definition 18 is consistent, and for any $x \in \mathring{S}$, x always admits one unique critical ball.*

Proof. For the proof, we rely on the properties of the set $F_x = \{c(b) \mid b \text{ medial in } S, x \in \mathring{b}\}$. Using the projection π from Section 4.2, we argue that F_x is path connected. Indeed consider b_1, b_2 such that $c(b_1), c(b_2) \in F_x$. Because $x \in \mathring{b}_1 \cap \mathring{b}_2$, we know that b_1 and b_2 are not interior disjoint. Similarly as in the proof of Lemma 4.9, we can project the pencil segment $[b_1 b_2]$ onto medial balls of S and obtain a path $\pi([b_1 b_2])$. From the properties of π , we easily deduce that the centers of this path of medial balls are contained in F_x . Thus F_x is path-connected. Therefore, F_x admits one unique T -maximal element, because $F_x \subseteq \text{MA}(S)$ and $\text{MA}(S)$ is a tree. This property suffices to prove the consistency of Definition 18.

Indeed, the series of implication (i) \implies (ii) \implies (iii) \implies (iv) \implies (v) are fairly obvious. In addition, if (v) holds, we easily deduce that b must be T -maximal in $\overline{F_x} = \{b' \text{ medial in } S \mid x \in b'\}$. Due to the remark above, b is actually the T -maximum of $\overline{F_x}$ and therefore we reach the last implication (v) \implies (i).

By statement (i) of Definition 18, we know that $x \in \overset{\circ}{S}$ admits at most one critical ball. Because $x \in \overset{\circ}{S}$, there always exists some medial ball b_m of S that covers x . Thus $\overline{F_x}$ is non-empty, and x always admits a critical ball. This concludes on the existence and unicity of critical balls. \square

Following Property 4.12, we can define the function pcrit that maps a point of the interior to the medial ball critical for that point.

$$\begin{aligned} \text{pcrit}: \overset{\circ}{S} &\rightarrow \{b \text{ medial in } S\} \\ x &\mapsto b \text{ critical for } x \end{aligned}$$

Because the notion of critical ball for a point is the basis of the later notion of critical ball for a set (Definition 19), we dedicate the next few paragraphs to the study of the properties of balls critical for a point. Eventually, we want some insight on the location of x within the ball $\text{pcrit}(x)$. In other words, we will study pcrit^{-1} and more precisely, the preimage of a ball b . This location property will let us establish some equivalence between the various notions of criticality: criticality for a point, the later criticality for a set, and even criticality across distinct shapes. This in turn will allow us to simplify the computation of critical balls.

We now investigate the inverse function pcrit^{-1} : given a medial ball b , to what points $x \in b$ will it be critical for, if any? Said otherwise, what is $\text{pcrit}^{-1}(b)$? Notice by statement (iii) of Definition 18 that we have

$$\text{pcrit}^{-1}(b) = \left(\overset{\circ}{S} \cap b \right) \setminus \bigcup_{\substack{b' >_T b \text{ or} \\ b', b \text{ } T\text{-incomparable}}} b' \quad (4.1)$$

When b is centered at the root of T , we easily deduce $\text{pcrit}^{-1}(b) = \overset{\circ}{S} \cap b$. Assume that b isn't at the root of T . We argue that then, $\text{pcrit}^{-1}(b)$ is a (possibly empty) union of open connected circular arcs supported by ∂b . Specifically, because b is not at the root of T , there exists a medial ball b_+ of S such that $b_+ >_T b$, and $[bb_+]$ is a pencil segment of $\text{MA}(S)$. We claim that $\text{pcrit}^{-1}(b) = \left(\overset{\circ}{S} \cap \partial b \right) \setminus b_+$.

Property 4.13. *Let S be a finite union of balls whose medial axis is cycle-free. Let b and b^+ be two medial balls of S such that $b <_T b^+$ and $[bb^+] \subseteq \text{MA}(S)$ is a pencil segment of the medial axis with no branching. Then*

$$\text{pcrit}^{-1}(b) = \left(\overset{\circ}{S} \cap \partial b \right) \setminus b_+$$

Proof. Let \mathcal{B}_+ be the collection $\mathcal{B}_+ = \{b' \text{ medial in } S \mid b' >_T b, \text{ or } b' \text{ and } b \text{ } T\text{-incomparable}\}$. Note that then, Equation (4.1) can be rewritten as $\text{pcrit}^{-1}(b) = \left(\overset{\circ}{S} \cap b \right) \setminus \bigcup \mathcal{B}_+$. We start by showing the following equality:

$$\text{pcrit}^{-1}(b) = \left(\overset{\circ}{S} \cap b \right) \setminus \bigcup_{b <_T b' \leq_T b_+} b'$$

To prove this, consider $b' \in \mathcal{B}_+$. Assume that b' is not T -smaller than b_+ , hence it is either T -larger than or T -incomparable to b_+ . We can thus apply Corollary 4.11 and deduce that $b \cap b' \subseteq b_+ \cup \mathcal{V}$, where \mathcal{V} denotes the vertex set of ∂S . This yields $b \setminus (b_+ \cup \mathcal{V}) \subseteq b \setminus b'$. Since $\mathcal{V} \cap \mathring{S} = \emptyset$, we can intersect with \mathring{S} which then yields $(\mathring{S} \cap b) \setminus b_+ \subseteq (\mathring{S} \cap b) \setminus b'$. Because $b_+ \in \mathcal{B}_+$, we can thus ignore any $b' \in \mathcal{B}_+$ that is T -larger than or T -incomparable to b_+ , since they do not restrict $(\mathring{S} \cap b) \setminus \bigcup \mathcal{B}_+$ anymore than b_+ already does. This directly implies the equality above.

To complete the proof, we now show that $b \setminus \bigcup_{b <_T b' \leq_T b_+} b' = \partial b \setminus b_+$. Recall that as per Property 2.1, the power of an affine combination of balls, is the affine combination of the powers. Thus for any $x \in \mathring{b}$, there exists a ball $b' \in]bb_+[$ such that $\text{pow}(x, b') < 0$, and thus $x \in \mathring{b}'$. It follows that $b \setminus \bigcup_{b <_T b' \leq_T b_+} b' \subseteq \partial b$. Obviously $\partial b \cap b_+$ cannot contribute to the left hand side of the inclusion, thus $b \setminus \bigcup_{b <_T b' \leq_T b_+} b' \subseteq \partial b \setminus b_+$. Conversely, every point in $\partial b \setminus b_+$ are exclusively covered by b in $b \cup (\bigcup \mathcal{B}_+)$, as per Section 2.2.2. Indeed we can consider the power cell $\text{pcell}(b)$ of b , it is an affine half-space whose interior covers $\partial b \setminus b_+$. It follows that for $x \in \partial b \setminus b_+$ and $b' \in]bb_+[$, $\text{pow}(x, b') > \text{pow}(x, b) = 0$, and $x \notin b'$. Therefore $\partial b \setminus b_+ \subseteq b \setminus \bigcup_{b <_T b' \leq_T b_+} b'$ and we actually have equality. Finally, $\text{pcrit}^{-1}(b) = (\mathring{S} \cap \partial b) \setminus b_+$. \square

See Figure 4.9 for examples of preimages.

Corollary 4.14. *If b is critical for x , then $x \in (C(S, b, -) \cup b)^\circ$.*

Proof. Consider $x \in \text{pcrit}(b)$. By definition of critical balls, we must have $x \in \mathring{S}$, and there is $r > 0$ such that $b(x, r) \subseteq S$. Because inclusion-wise maximal balls and medial balls coincide in finite unions of balls, there is some medial ball b_m of S such that $b(x, r) \subseteq b_m$, and we have $x \in \mathring{b}_m$. Because x cannot belong to any ball T -larger than or T -incomparable to b , it follows that $b_m \leq_T b$ and thus $b_m \subseteq C(S, b, -) \cup b$. Therefore, x must be in the interior of $C(S, b, -) \cup b$, which concludes the proof. \square

As a closing remark, note that Property 4.13 implies that any medial ball that is not a leaf of $\text{MA}(S)$ will be critical for some points of S . Indeed, when b is not a leaf, then there exists a ball $b_- <_T b$, such that $[b_- b]$ is a pencil segment of $\text{MA}(S)$. Then, $\mathring{b}_- \cap (\partial b \setminus b_+)$ is a non-empty open circular that is included in $\text{pcrit}^{-1}(b)$. Conversely, if b is a leaf of $\text{MA}(S)$, it cannot be critical for any point $x \in \mathring{S}$. Indeed, we then have $\partial b \setminus b_+ \subseteq \partial S$ and thus $\text{pcrit}^{-1}(b) = \emptyset$.

Remark 50. $\text{pcrit}^{-1}(b) \neq \emptyset \iff b$ is not a leaf of $\text{MA}(S)$

4.4.3 Critical ball for a set

With the basics of point criticality now established, we can now move on to the notion of set criticality. For a subset $X \subseteq \mathbb{R}^d$ that admits a T -order, and a non-empty subset $Y \subseteq \mathring{X}$,




Figure 4.9 – For medial ball b centered at c , we can apply Property 4.13 with ball b_+ (blue). The purple arc is the portion of ∂b covered by b_+ . The orange arcs are portions of ∂b that contribute to the boundary of S . Hence $\text{pcrit}^{-1}(b)$ is the red arc (excluding its endpoints). The green-blue arcs are other examples of preimages for the medial balls centered at the center of the corresponding colour.

consider the following two collections of balls:

$$F_{\text{cover}}(X, Y) = \{b' \text{ medial in } X \mid Y \subseteq b' \cup C(X, b', +)\}$$

and

$$F_{\text{pcrit}}(X, Y) = \{b' \text{ medial in } X \mid \exists y \in Y, b' \text{ critical for } y \text{ (in shape } X)\}$$

Definition 19 (Critical ball for a set). Let A be a non-empty subset of \hat{S} and consider a ball b medial in S . We say b is (T) -critical for A in shape S if equivalently:

- (i) b is T -maximal in $F_{\text{cover}}(S, A)$
- (ii) b is T -minimal in $\overline{F_{\text{pcrit}}(S, A)}$

Remark 51. A can be closed, open, or neither.

Intuitively, we want to preserve the idea that a ball critical for a set A is the furthest ball along orientation T that still allows for a covering of A . As such, assume we are sweeping S

with a medial ball b . Ball b becomes critical when we cannot advance further along T without having some part of A in $C(S, b, -)$. See Figure 4.10 for an illustration of balls in $F_{\text{cover}}(S, A)$, and Figure 4.11 for balls in $F_{\text{pcrit}}(S, A)$. Figure 4.12 showcases the corresponding critical balls.




Figure 4.10 – The two green dashed balls belong to $F_{\text{cover}}(S, A)$, because the set A (orange) entirely lies within $b \cup C(S, b, +)$ (green). The centers of balls in $F_{\text{cover}}(S, A)$ are the green line segments, including their endpoint, except for the root.

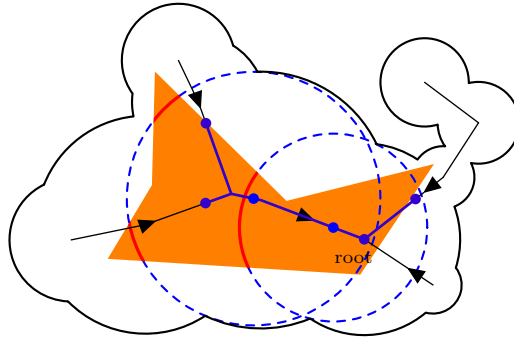


Figure 4.11 – The two blue dashed balls belong to $F_{\text{pcrit}}(S, A)$, because they are critical for some points of A (the red circular arcs). The centers of balls in $F_{\text{pcrit}}(S, A)$ are represented by the blue line segments.

Following Definitions 18 and 19, a ball may be critical for a point or a set. We refer to that point or set as the *constraint* of the critical ball. Note that when a set reduces to a singleton, the two definitions are equivalent.

Remark 52. Consider a constraint $A \subseteq \overset{\circ}{S}$ and b critical for A . In many cases, if b is not centered at the root of $\text{MA}(S)$, then b and A are not in “general position” with respect to one another. That is, an “extremity” of A lies on the boundary of b , or b and A are “tangent” to one another. When A is a singleton, this is obvious from Property 4.13. In fact, we can show through the later Property 4.16 that a degeneracy always occurs when the set constraint satisfies $\overline{A} \subseteq \overset{\circ}{S}$.

§ **Outline** In the remainder of this section, we first address the consistency of Definition 19, and prove the equivalence between the two statements (i) and (ii). Following that, we give




Figure 4.12 – Balls b_1, b_2, b_3 are critical for set A . Note that their centers c_1, c_2, c_3 lie on the boundary between the centers of $F_{\text{cover}}(S, A)$ (green) and $F_{\text{pcrit}}(S, A)$ (blue). The ball centered at the root of $\text{MA}(S)$ belongs to $F_{\text{pcrit}}(S, A)$ but is clearly not T -minimal, and it does not belong to $F_{\text{cover}}(S, A)$, it thus is not critical. Note also that $x_i \in \text{pcrit}^{-1}(b_i) \cap A$ and thus b_i satisfies Property 4.16.

again some remarks about the existence and unicity of critical balls for a set. We then conclude with a third characterization of balls critical for a set, which only holds given some additional assumption over the subset A . In practice, this characterization is easier to manipulate than T -extremal balls of $F_{\text{cover}}(S, A)$ or $F_{\text{pcrit}}(S, A)$, and will thus simplify the discussion later on.

§ **Consistency of set criticality** Before going any further, we first address why statements (i) and (ii) in Definition 19 are indeed equivalent.

Property 4.15. *Definition 19 is consistent.*

Proof. Consider $b_{\text{cover}} \in F_{\text{cover}}(S, A)$ and $b_{\text{pcrit}} \in F_{\text{pcrit}}(S, A)$, we first argue that (a) either $b_{\text{cover}} \leq_T b_{\text{pcrit}}$, or b_{cover} and b_{pcrit} are T -incomparable. Because this holds for every balls in $F_{\text{cover}}(S, A)$ and $F_{\text{pcrit}}(S, A)$, this also extends to T -maximal balls of $F_{\text{cover}}(S, A)$ and T -minimal balls of $F_{\text{pcrit}}(S, A)$. We conclude the equivalence by showing that (b) T -minimal balls of $F_{\text{pcrit}}(S, A)$ belong to $F_{\text{cover}}(S, A)$, while (c) T -maximal balls of $F_{\text{cover}}(S, A)$ also belong to $F_{\text{pcrit}}(S, A)$.

(a) By definition of $F_{\text{pcrit}}(S, A)$, let $x \in A$ such that b_{pcrit} is critical for x . By statement (v) of Definition 18, a medial ball b' may only satisfy $x \in b' \cup C(S, b', +)$ if it is T -smaller than or T -incomparable to b_{pcrit} . By definition of $F_{\text{cover}}(S, A)$, we know that $x \in A \subseteq b_{\text{cover}} \cup C(S, b_{\text{cover}}, +)$. It follows that either $b_{\text{cover}} \leq_T b_{\text{pcrit}}$, or b_{cover} and b_{pcrit} are T -incomparable.

(b) Assume now that b_{pcrit} is a T -minimal ball for $F_{\text{pcrit}}(S, A)$. We argue that $A \subseteq b_{\text{pcrit}} \cup C(S, b_{\text{pcrit}}, +)$. Indeed let $x \in A$, and consider b' a ball critical for x . We have $b' \in F_{\text{pcrit}}(S, A)$.

Because b_{pcrit} is T -minimal in $F_{\text{pcrit}}(S, A)$, then either $b_{\text{pcrit}} \leq_T b'$, or b' and b_{pcrit} are T -incomparable. Either way, we have $x \in b' \subseteq b_{\text{pcrit}} \cup C(S, b_{\text{pcrit}}, +)$. Because this holds for every $x \in A$, we deduce $A \subseteq b_{\text{pcrit}} \cup C(S, b_{\text{pcrit}}, +)$, and $b_{\text{pcrit}} \in F_{\text{cover}}(S, A)$.

(c) Now, assume that b_{cover} is a T -maximal ball for $F_{\text{cover}}(S, A)$. If b_{cover} is at the root of T , we then have $A \subseteq b_{\text{cover}}$ and b_{cover} is critical for every $x \in A$. Hence, $b_{\text{cover}} \in F_{\text{pcrit}}(S, A)$. Otherwise, there exists $b_+ >_T b_{\text{cover}}$ such that $[b_{\text{cover}}b_+] \subseteq \text{MA}(S)$. Consider a medial ball b' such that $b_{\text{cover}} <_T b' \leq_T b_+$. Because $b' >_T b_{\text{cover}}$, the T -maximality of b_{cover} in $F_{\text{cover}}(S, A)$ implies that $A \not\subseteq b' \cup C(S, b', +)$. Let $x \in A \setminus (b' \cup C(S, b', +))$ and b_x a medial ball critical for x , we have $b_x \in F_{\text{pcrit}}(S, A)$. Because $x \in b_x$ and $x \notin b' \cup C(S, b', +)$, we necessarily have $b_x \leq_T b'$. Because $x \in b_{\text{cover}} \cup C(S, b_{\text{cover}}, +)$, and b_x is critical for x , either $b_{\text{cover}} \leq_T b_x$ or b_{cover} and b_x are T -incomparable (by statement (v) of Definition 18). We claim that necessarily, $b_{\text{cover}} \leq_T b_x$. Indeed, b_x cannot be T -smaller than b_{cover} , yet it is T -smaller than b' . The only medial balls T -smaller than b' but not b_{cover} must lie on the pencil segment $[b_{\text{cover}}b_+]$, all of which can be T -ordered with respect to b_{cover} . It follows that $b_{\text{cover}} \leq_T b_x \leq_T b'$. Thus for every b' between b_{cover} and b_+ , there exists a ball b_x between b_{cover} and b' in $F_{\text{pcrit}}(S, A)$. We can thus build a sequence of balls in $F_{\text{pcrit}}(S, A)$ that converges to b_{cover} , thus $b_{\text{cover}} \in \overline{F_{\text{pcrit}}(S, A)}$.

Finally the equivalence (i) \iff (ii) holds. \square

§ General remarks (existence, unicity/multiplicity, closure) Because a point always admits a critical ball, and A is non-empty, the two collections of balls $F_{\text{cover}}(S, A)$ and $\overline{F_{\text{pcrit}}(S, A)}$ are non-empty, they thus admit T -extremal balls. It follows that any non-empty set $A \subseteq \mathring{S}$ always has a critical ball. However contrary to a single point x , a set A can very easily have several distinct critical balls, see Figure 4.12. Indeed recall that as per Property 4.12, any point x admits one unique critical ball. But for a set constraint A , it suffices for $F_{\text{cover}}(S, A)$ or $\overline{F_{\text{pcrit}}(S, A)}$ to span several T -incomparable branches of the medial axis, and we immediately obtain several distinct critical balls.

Remark 53. If a set A admits several critical balls, they are all pairwise T -incomparable.

As a remark, note that a non-empty subset A has the same critical balls as its closure \overline{A} , assuming that $\overline{A} \subseteq \mathring{S}$. Indeed we consider medial balls to be closed and thus for any medial ball b , the subset $b \cup C(S, b, +)$ is closed. It follows that the two inclusions $A \subseteq b \cup C(S, b, +)$ and $\overline{A} \subseteq b \cup C(S, b, +)$ are equivalent. Hence the two subsets A and \overline{A} define the same collection of balls $F_{\text{cover}}(S, A) = F_{\text{cover}}(S, \overline{A})$, they thus have the same critical balls.

§ Weak(er) characterization of set criticality Given some mild assumption on the set A , we argue that a ball b will be critical for the set constraint A if there exists a specific point $x \in \overline{A}$ such that b is critical for the point constraint x . The existence of that specific point x implies that in a sense, it is equivalent to be critical for either A or x . This equivalence is stronger than taking any random point within $y \in \text{pcrit}^{-1}(b)$. Indeed, as y may not belong to \overline{A} , there may be some interferences when A admits several critical balls. With the additional

property $x \in \bar{A}$, we can effectively substitute A by a finite collection of points, one per critical ball. This relationship enables us to reduce set constraints to point constraints, which greatly simplifies several discussions and considerations later on. Specifically we have the property below.

Property 4.16. *Consider A a non-empty subset of S , with $\bar{A} \subseteq \mathring{S}$. Then, b is critical for A if and only if b satisfies both*

- (1) $A \subseteq b \cup C(S, b, +)$, and
- (2) there exists $x \in \bar{A}$ such that b is critical for x .

Proof. We first show that if a medial ball b satisfies both (1) and (2), then it is critical for A . Specifically we show that then, b must be T -maximal in $F_{\text{cover}}(S, A)$. Because $A \subseteq b \cup C(S, b, +)$ by assumption (1), we have $b \in F_{\text{cover}}(S, A)$ and b is thus critical for A . Recall that we have

$$F_{\text{cover}}(S, A) = F_{\text{cover}}(S, \bar{A}) = \{b' \text{ medial in } S \mid \bar{A} \subseteq b' \cup C(S, b', +)\}$$

As per (2), let $x \in \bar{A}$ such that b is critical for x , and consider $b_{\text{cover}} \in F_{\text{cover}}(S, A)$. We have $x \in b_{\text{cover}} \cup C(S, b_{\text{cover}}, +)$. Since b is critical for A , we deduce that either $b_{\text{cover}} \leq_T b$, or b and b_{cover} are T -incomparable. Because this holds for every ball of $F_{\text{cover}}(S, A)$, hence b is T -maximal in $F_{\text{cover}}(S, A)$.

Consider now a ball b critical for A . Then, b necessarily belongs to $F_{\text{cover}}(S, A)$ and it thus satisfies (1), $A \subseteq b \cup C(S, b, +)$. We show that b also fulfills (2), there is $x \in \bar{A}$ such that b is critical for x . If b is centered at the root of T , it is then critical for every point of A and we are done. Otherwise, there exists a medial ball $b_+ >_T b$ of S such that $[bb_+] \subseteq \text{MA}(S)$. We know that b is T -minimal in $\overline{F_{\text{pcrit}}(S, A)}$. If $b \in F_{\text{pcrit}}(S, A)$, we immediately deduce that b satisfies (2). Otherwise, there is a sequence of balls (b_n) in $F_{\text{pcrit}}(S, A)$ that converges to b , such that $\forall n, b_{n+1} <_T b_n$. Without loss of generality, we can additionally assume that $\forall n, b_n \leq_T b_+$. Let $x_n \in A$ such that b_n is critical for x_n , we have $x_n \in \text{pcrit}^{-1}(b_n)$. Let x be a limit point of (x_n) , we argue that necessarily, $x \in \text{pcrit}^{-1}(b)$, and thus that b is critical for $x \in \bar{A}$. Indeed, note that every b_n belongs to the pencil segment $[bb_+]$, and we have $\text{pcrit}^{-1}(b_n) = (\mathring{S} \cap \partial b_n) \setminus b_+$ as per Property 4.13. In addition, we have $\bar{A} \subseteq \mathring{S}$, hence there is a positive real $r > 0$ such that $d(\bar{A}, \partial S) \geq r$. In particular, consider $\{u, v\}$ the radical sphere of pencil $[bb_+]$, then for all n we have $\|u - x_n\| \geq r$ and $\|v - x_n\| \geq r$. It follows that $x_n \in \text{pcrit}^{-1}(b_n) \setminus (b(u, r) \cup b(v, r))^\circ \subseteq \partial b_n \setminus (b(u, r) \cup b(v, r) \cup b_+)^\circ$. As n goes to infinity, these converge to the closed set $\partial b \setminus (b(u, r) \cup b(v, r) \cup b_+)^\circ$. Thus $x \in \mathring{S} \cap \partial b \setminus (b(u, r) \cup b(v, r) \cup b_+)^\circ \subseteq (\mathring{S} \cap \partial b) \setminus b_+ = \text{pcrit}^{-1}(b)$. Hence, $x \in \text{pcrit}^{-1}(b)$, and b is critical for x . Therefore b satisfies (2), which concludes the proof. \square

See Figure 4.13 for a counter-example when the assumption $\bar{A} \subseteq \mathring{S}$ is not verified. Note that in the proof above, the assumption $\bar{A} \subseteq \mathring{S}$ was only used to prove one of the two implications, we thus also obtain this corollary.

Corollary 4.17. *Let A be a non-empty subset of S and let b be a medial ball of S . If*

- (1) $A \subseteq b \cup C(S, b, +)$, and
- (2) there exists $x \in \bar{A}$ such that b is critical for x ,

then b is critical for A .




Figure 4.13 – Consider the set constraint A to be the red circle punctured at point u . Every medial ball except b_- (green) is critical for some point of A . In fact b_- is the unique T -minimal ball in $F_{\text{perit}}(S, A)$, thus it is critical for A . However, b_- isn't critical for any point in \bar{A} . Note that as we approach the center of b_- , the point for which the corresponding medial ball is critical converges to u .

Hence when the constraint set A satisfies $\bar{A} \subseteq \mathring{S}$, Property 4.16 yields a third equivalent definition of criticality for A . See Figure 4.12. Every constraint set we consider henceforth will fulfill the assumption of Property 4.16. As such the underlying strategy we adopt in Section 4.6 and in Chapter 5 when computing these critical balls revolves around finding this particular point $x \in \bar{A}$.

4.5 Critical balls and T -maximal candidates

Using various results of the previous sections, we now explicit the relationship between critical balls and T -maximal candidates. Recall that in order to compute a $(0, \varepsilon)$ -ball approximation of S , we want to build T -small approximations by computing T -maximal candidates. We claim and prove in this section that for well chosen constraints, critical balls coincide with T -maximal candidates.

§ **Intuition and computation** Let \mathcal{B} be a finite collection of medial balls, and assume it is a T -small $(0, \varepsilon)$ -ball approximation of S . Then, let A be the set $A = S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}$, and consider a medial ball b of S . We can fairly easily deduce the equivalence between b being critical for A , and b being a T -maximal candidate for \mathcal{B} . Indeed, it suffices to exploit the properties of balls T -maximal in $F_{\text{cover}}(S, A)$. Consider $b \in F_{\text{cover}}(S, A)$. Intuitively, the condition $A \subseteq b \cup C(S, b, +)$ in conjunction with Proposition 4.10 will yield $S^{\ominus\varepsilon} \cap C(S, b, -) \subseteq \bigcup \mathcal{B}$. In other words, the collection $\mathcal{B} \cup \{b\}$ properly covers the erosion of S in its T -small component. Adding in the T -maximality of critical balls in $F_{\text{cover}}(S, A)$, one could then prove the equivalence between criticality for A , and T -maximal candidacy for \mathcal{B} .

One major caveat of using the set constraint $S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}$ directly, is that computation of a critical ball for such a set is difficult. Instead, we build on Lemma 4.3 to define a smaller set, for which computation is more easily tractable, and that still fulfills the equivalence between critical balls and T -maximal candidates. Let $A(S, \varepsilon, \mathcal{B})$ be

$$A(S, \varepsilon, \mathcal{B}) = \left((\partial S^{\ominus\varepsilon}) \setminus \bigcup \mathcal{B} \right) \cup \left\{ v \text{ vertex of } \bigcup \mathcal{B} \mid \forall N_v \text{ neighbourhood of } v, S^{\ominus\varepsilon} \cap N_v \not\subseteq \bigcup \mathcal{B} \right\}$$

In other words, $A(S, \varepsilon, \mathcal{B})$ is the part of $\partial S^{\ominus\varepsilon}$ not yet covered by $\bigcup \mathcal{B}$, plus vertices of $\bigcup \mathcal{B}$ that do not satisfy proper coverage of $S^{\ominus\varepsilon}$ in their neighbourhood. Refer to Figure 4.14 for an illustration of various features in $A(S, \varepsilon, \mathcal{B})$. Observe also in Figure 4.14.c that it is mandatory to include some vertices of $\bigcup \mathcal{B}$ in the set $A(S, \varepsilon, \mathcal{B})$.

Notice that Lemma 4.3 ensures that:

$$A(S, \varepsilon, \mathcal{B}) = \emptyset \iff \mathcal{B} \text{ is a } (0, \varepsilon)\text{-ball approximation}$$

When $A(S, \varepsilon, \mathcal{B})$ is non-empty, observe that the closure of $A(S, \varepsilon, \mathcal{B})$ lies in $S^{\ominus\varepsilon}$. Thus this set can be used as a constraint. It also fulfills the conditions of Property 4.16.

The main result of this section is as follows.

Theorem 8. Let \mathcal{B} be a partial T -small $(0, \varepsilon)$ -ball approximation of S . The T -maximal candidates to \mathcal{B} are precisely the balls critical for $A(S, \varepsilon, \mathcal{B})$.

Remark 54. Because \mathcal{B} is a partial T -small $(0, \varepsilon)$ -ball approximation, we know that $A(S, \varepsilon, \mathcal{B}) \neq \emptyset$, thus the notion of ball critical for $A(S, \varepsilon, \mathcal{B})$ is always well-defined.

§ **Outline** The remainder of this section is dedicated to the proof of this result. We start by proving that a ball critical for $A(S, \varepsilon, \mathcal{B})$ is a T -maximal candidate, then argue that any candidate is always T -smaller than some of these critical balls. It follows that the collections of critical balls and T -maximal candidates thus coincide.

We defer to the later Section 4.6 the explanation as to why computing balls critical for the set constraint $A(S, \varepsilon, \mathcal{B})$ is computationally feasible.




Figure 4.14 – Illustration of $A(S, \varepsilon, \mathcal{B})$ (red arcs and point). The set \mathcal{B} is composed of the two green balls. The dark grey area is the part of $S^{\ominus \varepsilon}$ covered by $\bigcup \mathcal{B}$, the light grey is the non-covered part. In (b), ball b_* (blue) contains all of $A(S, \varepsilon, \mathcal{B})$ within $b_* \cup C(S, b_*, +)$, and it is a candidate. The blue area is the additional parts of $S^{\ominus \varepsilon}$ that would be covered by including b_* in \mathcal{B} . Observe in (c) that if vertex v is not properly contained in ball b_* , then b_* cannot be a candidate to \mathcal{B} .

4.5.1 Critical balls are candidates

Let \mathcal{B} be a partial T -small $(0, \varepsilon)$ -ball approximation, and consider b_* a medial ball of S such that for all $b' \in \mathcal{B}$, either $b' <_T b_*$, or b' and b_* are T -incomparable. We have a closer look at what it means for b_* to be a candidate ball to \mathcal{B} . By definition, b_* is a candidate with respect to \mathcal{B} if and only if the addition of b_* to the collection \mathcal{B} preserves the covering property of $S^{\ominus \varepsilon}$ in the T -small component, and also strictly increases coverage of $S^{\ominus \varepsilon}$. Let $\mathcal{B}_* = \mathcal{B} \cup \{b_*\}$.

Because we know that \mathcal{B} is a T -small $(0, \varepsilon)$ -ball approximation, any change to the part of $S^{\ominus \varepsilon}$ covered in $C(S, \mathcal{B}_*, -)$ must occur within $C(S, b_*, -)$. Indeed, we recall that by definition we have $C(S, \mathcal{B}_*, -) = C(S, \mathcal{B}, -) \cup C(S, b_*, -)$. We can thus disregard everything that happens in branches T -incomparable to b_* , and can assume without loss of generality that b_* is the unique T -maximum ball in \mathcal{B}_* . In particular, we obtain $C(S, \mathcal{B}_*, -) = C(S, b_*, -)$.

For now we focus on the first property to be verified, $S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -) \subseteq \bigcup \mathcal{B}_*$.

Proposition 4.18. *Consider b_* critical for $A(S, \varepsilon, \mathcal{B})$. Then, $S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -) \subseteq \bigcup \mathcal{B}_*$.*

In order to prove this proposition, we shall rely on results from Section 4.3. Specifically, we use a variation of Lemma 4.3. First, we argue that the boundary of $S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -)$ belongs to $\bigcup \mathcal{B}_*$. From there, we consider $H = (S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -)) \setminus \bigcup \mathcal{B}_*$, that is the part of $S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -)$ not covered by $\bigcup \mathcal{B}_*$. Because $S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -)$ is bounded, we can apply Lemma 4.4 and deduce that H is either empty, or a union of holes of $\bigcup \mathcal{B}_*$. Finally, we will show that H cannot contain any hole of $\bigcup \mathcal{B}_*$, thus implying $H = \emptyset$ and thus that $S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -) \subseteq \bigcup \mathcal{B}_*$.

Lemma 4.19. *Consider b_* critical for $A(S, \varepsilon, \mathcal{B})$. Then, $\partial(S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -)) \subseteq \bigcup \mathcal{B}_*$.*

Proof. Without loss of generality, it suffices to prove that $\partial(S^{\ominus\varepsilon} \cap C(S, b_*, -)) \subseteq \bigcup \mathcal{B}_*$. By Lemma A.1, we have

$$\partial(S^{\ominus\varepsilon} \cap C(S, b_*, -)) \subseteq \left(\partial S^{\ominus\varepsilon} \cap \overline{C(S, b_*, -)} \right) \cup \left(\overline{S^{\ominus\varepsilon}} \cap \partial C(S, b_*, -) \right)$$

Recall that both $C(S, b_*, -) \cup b_*$ and $S^{\ominus\varepsilon}$ are closed. In particular, $\overline{C(S, b_*, -)} \subseteq C(S, b_*, -) \cup b_*$ and $\overline{S^{\ominus\varepsilon}} = S^{\ominus\varepsilon}$. Let $\gamma = \partial S^{\ominus\varepsilon} \cap (C(S, b_*, -) \cup b_*)$ and $\delta = S^{\ominus\varepsilon} \cap \partial C(S, b_*, -)$, obviously $\gamma \cup \delta$ is a superset of the boundary of interest. See Figure 4.15. We successively show that (a) $\gamma \subseteq \bigcup \mathcal{B}_*$ and (b) $\delta \subseteq \bigcup \mathcal{B}_*$, thus proving the result.

(a) Recall that $\partial S^{\ominus\varepsilon} \subseteq S = C(S, \mathcal{B}, -) \cup (\bigcup \mathcal{B}) \cup C(S, \mathcal{B}, +)$, It follows that $\partial S^{\ominus\varepsilon}$, and thus γ as well, can be split without loss between $C(S, \mathcal{B}, -) \cup (\bigcup \mathcal{B})$ and $C(S, \mathcal{B}, +)$. Because \mathcal{B} is a T -small $(0, \varepsilon)$ -ball approximation, hence $\partial S^{\ominus\varepsilon} \cap C(S, \mathcal{B}, -) \subseteq \bigcup \mathcal{B}$. We thus easily obtain that

$$\gamma \cap \left(C(S, \mathcal{B}, -) \cup \left(\bigcup \mathcal{B} \right) \right) \subseteq \bigcup \mathcal{B}$$

Consider the remaining parts of γ , that is $\gamma \setminus (C(S, \mathcal{B}, -) \cup (\bigcup \mathcal{B}))$. By Proposition 4.10, this coincides with $\gamma \cap C(S, \mathcal{B}, +)$. Note that this is a subset of $A(S, \varepsilon, \mathcal{B})$, since it is included in $\partial S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}$. Thus $\gamma \cap C(S, \mathcal{B}, +) \subseteq A(S, \varepsilon, \mathcal{B}) \cap (C(S, b_*, -) \cup b_*)$. Because $A(S, \varepsilon, \mathcal{B}) \subseteq \mathring{S}$, it follows by Proposition 4.10 that $A(S, \varepsilon, \mathcal{B}) \cap (C(S, b_*, -) \cup b_*) = A(S, \varepsilon, \mathcal{B}) \setminus C(S, b_*, +)$. Since b_* is critical for $A(S, \varepsilon, \mathcal{B})$, we have $A(S, \varepsilon, \mathcal{B}) \subseteq b_* \cup C(S, b_*, +)$, which implies $A(S, \varepsilon, \mathcal{B}) \setminus C(S, b_*, +) \subseteq b_*$. Therefore,

$$\gamma \cap C(S, \mathcal{B}, +) \subseteq A(S, \varepsilon, \mathcal{B}) \cap (C(S, b_*, -) \cup b_*) \subseteq b_*$$

In conclusion,

$$\gamma = \left(\gamma \cap \left(C(S, \mathcal{B}, -) \cup \left(\bigcup \mathcal{B} \right) \right) \right) \cup \left(\gamma \cap C(S, \mathcal{B}, +) \right) \subseteq \left(\bigcup \mathcal{B} \right) \cup b_* = \bigcup \mathcal{B}_*$$

and thus γ satisfies the desired inclusion.

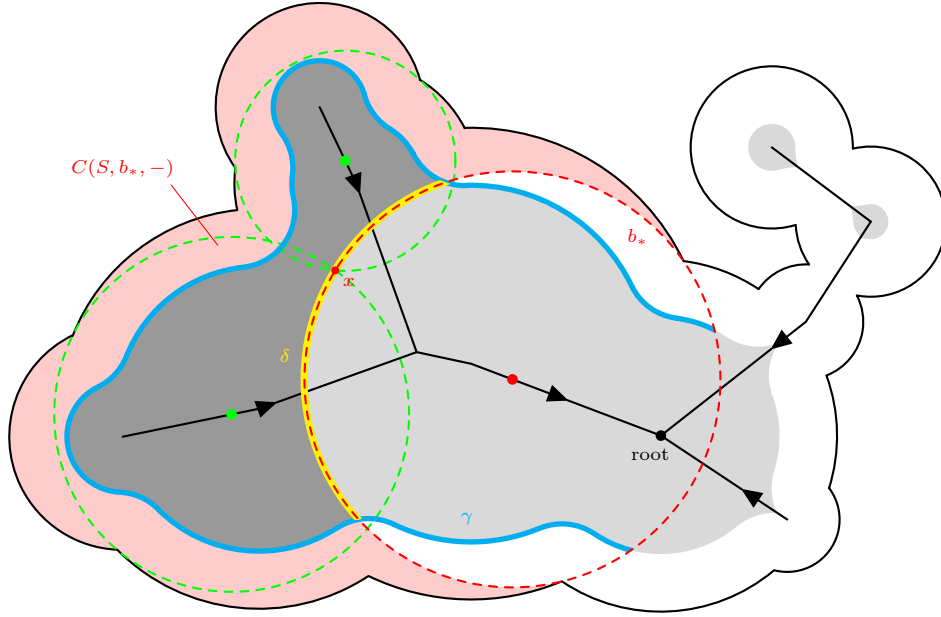


Figure 4.15 – Illustration for the proof of Lemma 4.19. The erosion of S is the (light and dark) grey areas. Consider \mathcal{B} to be the two green balls. Then b_* (red) is critical for $A(S, \varepsilon, \mathcal{B})$ because it is critical for x . Observe that the domain $S^{\ominus\varepsilon} \cap C(S, b_*, -)$ is the dark grey area, and that its boundary is indeed covered by the two curves γ (blue) and δ (yellow).

(b) Since $S^{\ominus\varepsilon} \subseteq \overset{\circ}{S}$ we have $\delta \subseteq \overset{\circ}{S} \cap \partial C(S, b_*, -)$. We argue that $\overset{\circ}{S} \cap \partial C(S, b_*, -) \subseteq b_*$. Indeed let $x \in \overset{\circ}{S} \cap \partial C(S, b_*, -)$. If $x \in b_*$, we are done. Assume that $x \notin b_*$, and let $r > 0$ such that $b(x, r) \subseteq S \setminus b_*$. There exists some ball b_m medial in S such that $b(x, r) \subseteq b_m$. We claim that b_m is either strictly T -larger than, or T -incomparable to b_* . Indeed, if we had $b_m <_T b_*$, then $x \in \overset{\circ}{b_m} \setminus b_* \subseteq \overset{\circ}{C}(S, b_*, -)$, which contradicts the definition of x . Because $x \in \partial C(S, b_*, -) \subseteq C(S, b_*, -) \cup b_*$, necessarily there exists some ball $b_- \leq_T b_*$ such that $x \in b_-$. Since we know that x is in the interior of b_m , thus b_m and b_- cannot be interior disjoint. We deduce from Corollary 4.11 that $b_- \cap b_m \subseteq b_*$. Specifically this implies $x \in b_*$. Therefore, $\delta \subseteq \overset{\circ}{S} \cap \partial C(S, b_*, -) \subseteq b_* \subseteq \bigcup \mathcal{B}_*$, which concludes the proof. \square

From the above, $S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -)$ has its boundary in $\bigcup \mathcal{B}_*$. Because it is also bounded, it thus satisfies the conditions of lemma 4.4, from which we deduce that $H = (S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -)) \setminus \bigcup \mathcal{B}_*$ is either empty or a union of holes of $\bigcup \mathcal{B}_*$. We now argue that H must be empty.

Proof of Proposition 4.18. Recall that without loss of generality, we can assume $C(S, \mathcal{B}_*, -) = C(S, b_*, -)$. Suppose for a contradiction that H is non-empty, it is then a union of holes of $\bigcup \mathcal{B}_*$. As per Corollary 4.8, the boundary of H must contain a vertex v that does not lie on ∂b_* . This vertex v is also a vertex in $\partial \bigcup \mathcal{B}$. Because $v \in \partial H$, and by definition of v , the subset $S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -)$ is not properly included in $\bigcup \mathcal{B}_*$ in the neighbourhood of v . Hence, $v \in A(S, \varepsilon, \mathcal{B})$. Since b_* is critical for $A(S, \varepsilon, \mathcal{B})$, it follows that $v \in b_* \cup C(S, b_*, +)$. Since

$v \in S^{\ominus \varepsilon} \subseteq \overset{\circ}{S}$, we deduce by Proposition 4.10 that $v \notin C(S, b_*, -)$. By definition of v , we also know that $v \notin b_*$. Thus, $v \notin C(S, b_*, -) \cup b_* \supseteq \partial H$, which is absurd. Therefore, H must be empty. \square

We now know as per Proposition 4.18 that a ball b_* critical for $A(S, \varepsilon, \mathcal{B})$ will properly preserve the covering property of $S^{\ominus \varepsilon}$ within the T -small component of the collection of balls. We now address the second property of candidate balls, namely that we must have the strict inclusion $S^{\ominus \varepsilon} \setminus \bigcup \mathcal{B}_* \subsetneq S^{\ominus \varepsilon} \setminus \bigcup \mathcal{B}$. To do so, consider b_* critical for $A(S, \varepsilon, \mathcal{B})$. We show that, around a point x that lies in both $\text{pcrit}^{-1}(b_*)$ and $\overline{A(S, \varepsilon, \mathcal{B})}$, there exists a part of $S^{\ominus \varepsilon}$ not covered by $\bigcup \mathcal{B}$ but covered by $\bigcup \mathcal{B}_*$.

Lemma 4.20. *Consider b_* critical for $A(S, \varepsilon, \mathcal{B})$. Then for all x in $\overline{A(S, \varepsilon, \mathcal{B})}$ such that b_* is critical for x , and for all N_x open neighbourhood of x , we have $S^{\ominus \varepsilon} \cap N_x \not\subseteq \bigcup \mathcal{B}$.*

Proof. Let $x \in \overline{A(S, \varepsilon, \mathcal{B})} \cap \text{pcrit}^{-1}(b_*)$. We distinguish three main cases depending on the nature of x , see Figure 4.17.

- (a) $x \in A(S, \varepsilon, \mathcal{B})$ is a vertex of $\bigcup \mathcal{B}$. By definition of $A(S, \varepsilon, \mathcal{B})$, we immediately deduce that for every neighbourhood N_x of x , $S^{\ominus \varepsilon}$ does not satisfy local inclusion in $\bigcup \mathcal{B}$ around x .
- (b) $x \in A(S, \varepsilon, \mathcal{B})$ and $x \in \partial S^{\ominus \varepsilon} \setminus \bigcup \mathcal{B}$. Because $S^{\ominus \varepsilon}$ is closed, it contains its own boundary, and the existence of x itself suffices to conclude.
- (c) $x \notin A(S, \varepsilon, \mathcal{B})$, there exists an infinite sequence of points $(x_n) \subseteq A(S, \varepsilon, \mathcal{B})$ that converges to x . Without loss of generality, we can assume that (x_n) does not meet the vertices of $\bigcup \mathcal{B}$, hence $(x_n) \subseteq \partial S^{\ominus \varepsilon} \setminus \bigcup \mathcal{B}$. For any open neighbourhood N_x of x , we have by definition of (x_n) that $(x_n) \cap N_x \neq \emptyset$. Therefore, $(x_n) \cap N_x \subseteq S^{\ominus \varepsilon} \cap N_x \neq \emptyset$. \square




Figure 4.17 – In the proof of Lemma 4.20, point v is in case (a), x in case (b) and y in case (c).

Corollary 4.21. *Consider b_* critical for $A(S, \varepsilon, \mathcal{B})$. Then $S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}_* \subsetneq S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}$.*

Proof. Recall that $S^{\ominus\varepsilon}$ is closed and included within \mathring{S} . Therefore by Property 4.16, there exists some $x \in \overline{A(S, \varepsilon, \mathcal{B})}$ such that b_* is critical for x . Because $x \in \mathring{S}$, we deduce by Corollary 4.14 that $x \in (C(S, b_*, -) \cup b_*)^\circ$, and thus there exists an open neighbourhood N_x of x such that $N_x \subseteq C(S, b_*, -) \cup b_*$. In addition, we know by Lemma 4.20 that $S^{\ominus\varepsilon} \cap N_x \not\subseteq \bigcup \mathcal{B}$. By Proposition 4.18 however, we have $S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -) \subseteq \bigcup \mathcal{B}_*$. In particular,

$$S^{\ominus\varepsilon} \cap N_x \subseteq S^{\ominus\varepsilon} \cap (C(S, b_*, -) \cup b_*) \subseteq S^{\ominus\varepsilon} \cap (C(S, \mathcal{B}_*, -) \cup b_*) \subseteq \bigcup \mathcal{B}_*$$

To conclude, $\bigcup \mathcal{B}_*$ covers strictly more points of $S^{\ominus\varepsilon}$ than $\bigcup \mathcal{B}$ does, therefore we have the desired strict inclusion. \square

4.5.2 Critical balls are T -maximal candidates

Following Proposition 4.18 and Corollary 4.21, it follows that if b_* is critical for $A(S, \varepsilon, \mathcal{B})$, then b_* is a candidate with respect to \mathcal{B} . To conclude on the first implication of Theorem 8, it remains to show that b_* is T -maximal among candidates to \mathcal{B} . This can easily be achieved as a corollary to previous lemmas.

Corollary 4.22. *Consider b_* critical for $A(S, \varepsilon, \mathcal{B})$. Then, b_* is a T -maximal candidate for \mathcal{B} .*

Proof. Recall that Property 4.16 holds and let $x \in \overline{A(S, \varepsilon, \mathcal{B})}$ such that b_* is critical for x . Assume for a contradiction that there exists a candidate ball b_+ to \mathcal{B} that satisfies $b_+ >_T b_*$. By criticality of b_* to point constraint x , we have $x \notin b_+$. Hence there is a neighbourhood N_x of x , such that $N_x \cap b_+ = \emptyset$. Using Corollary 4.14 we know that $x \in (C(S, b_*, -) \cup b_*)^\circ$, thus we can adjust N_x so that $N_x \subseteq C(S, b_*, -) \cup b_* \subseteq C(S, b_+, -) \cup b_+$, which here simplifies to $N_x \subseteq C(S, b_+, -)$. Indeed, we already have by definition of N_x that $N_x \cap b_+ = \emptyset$.

By Lemma 4.20, we know that $S^{\ominus\varepsilon} \cap N_x \not\subseteq \bigcup \mathcal{B}$. As such, let $y \in S^{\ominus\varepsilon} \cap N_x \setminus \bigcup \mathcal{B}$. Because $y \in N_x \subseteq C(S, b_+, -)$, we have $y \in S^{\ominus\varepsilon} \cap C(S, b_+, -)$. Since b_+ is a candidate to \mathcal{B} , this implies $y \in (\bigcup \mathcal{B}) \cup b_+$. However, $y \notin \bigcup \mathcal{B}$, therefore y must belong to b_+ . Thus $y \in N_x \cap b_+ = \emptyset$, which is impossible. It follows that b_+ cannot exist, and thus that b_* is a T -maximal candidate. \square

4.5.3 T -maximal candidates are critical for $A(S, \varepsilon, \mathcal{B})$

We now know that balls critical to $A(S, \varepsilon, \mathcal{B})$ are necessarily T -maximal candidates to \mathcal{B} . To conclude the proof of Theorem 8, we must still show the reverse property, that is T -maximal candidates must be critical to $A(S, \varepsilon, \mathcal{B})$.

Lemma 4.23. *Consider b_* a T -maximal candidate to \mathcal{B} . Then b_* is critical for $A(S, \varepsilon, \mathcal{B})$.*

Proof. We proceed in two steps. We start by showing that b_* belongs to $F_{\text{cover}}(S, A(S, \varepsilon, \mathcal{B}))$, then that it also belongs to $F_{\text{pcrit}}(S, \overline{A(S, \varepsilon, \mathcal{B})})$. Since a set and its closure have the same critical balls, this suffices to conclude.

Note that $A(S, \varepsilon, \mathcal{B}) \subseteq \overline{S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}}$. Let $\mathcal{B}_* = \mathcal{B} \cup \{b_*\}$. Because b_* is a candidate to \mathcal{B} , we have $S^{\ominus\varepsilon} \cap C(S, b_*, -) \subseteq S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_*, -) \subseteq \bigcup \mathcal{B}_*$, that is the portion of $S^{\ominus\varepsilon}$ within $C(S, b_*, -)$ is entirely covered by $\bigcup \mathcal{B}_*$. This implies $S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}_* \subseteq S \setminus C(S, b_*, -) \subseteq b_* \cup C(S, b_*, +)$. From the definition of \mathcal{B}_* , we can rewrite this last inclusion as $S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B} \subseteq b_* \cup C(S, b_*, +)$. Because $b_* \cup C(S, b_*, +)$ is closed, we finally obtain $A(S, \varepsilon, \mathcal{B}) \subseteq b_* \cup C(S, b_*, +)$. Therefore $b_* \in F_{\text{cover}}(S, A(S, \varepsilon, \mathcal{B}))$, where $F_{\text{cover}}(S, A(S, \varepsilon, \mathcal{B})) = \{b' \text{ medial in } S \mid A(S, \varepsilon, \mathcal{B}) \subseteq b' \cup C(S, b', +)\}$.

Because b_* is candidate to \mathcal{B} , b_* must contains points of $S^{\ominus\varepsilon}$ that are not covered by $\bigcup \mathcal{B}$. Consider a ball b' such that $b' >_T b_*$. By T -maximality of b_* among candidates, b' cannot be a candidate, and $b_* \setminus b'$ must still contain some points of $S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}$. Because this holds for every $b' >_T b_*$, we deduce the existence of a point $x \in \overline{S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}} \cap \text{pcrit}^{-1}(b_*)$. We claim that $x \in \overline{A(S, \varepsilon, \mathcal{B})}$. By Corollary 4.14 we know that x lies in $(C(S, b_*, -) \cup b_*)^\circ$. Let N_x be a neighbourhood of x such that $N_x \subseteq C(S, b_*, -) \cup b_*$. We distinguish three cases depending on the nature of x .

(a) $x \in (S^{\ominus\varepsilon})^\circ$. First, we will show that x must be a vertex of $\bigcup \mathcal{B}$. From here, by definition of x , $\bigcup \mathcal{B}$ does not locally cover $S^{\ominus\varepsilon}$ in the neighbourhood of x , which enables to conclude that $x \in A(S, \varepsilon, \mathcal{B})$.

Because x is in the interior of $S^{\ominus\varepsilon}$, we can choose the neighbourhood N_x such that $N_x \subseteq S^{\ominus\varepsilon}$. Since $N_x \subseteq C(S, b_*, -) \cup b_*$ and b_* is a candidate to \mathcal{B} , we must have $N_x \subseteq \bigcup \mathcal{B}_*$. Let \mathcal{B}_x the collection of balls in \mathcal{B}_* that contain x . By definition of x , we know that $x \in \partial b_*$ and $x \in (\bigcup \mathcal{B})^c$. Hence for all $b \in \mathcal{B}_x$ we have $x \in \partial b$, and \mathcal{B}_x defines a pencil. In addition, \mathcal{B}_x locally represents $\bigcup \mathcal{B}_*$ in the neighbourhood of x , and we can enforce $N_x \subseteq \bigcup \mathcal{B}_x$. Any ball through x must thus be locally contained in $\bigcup \mathcal{B}_x$, and we deduce from Property 2.40 that $\text{coni}(x, c(\mathcal{B}_x)) = \mathbb{R}^2$. This implies that $\mathcal{B}_x \setminus \{b_*\}$ is non-empty, and defines either an elliptic pencil of dimension 1, or a parabolic pencil. In both cases, the degree of x in $\partial \bigcup \mathcal{B}$ is null, and x is a vertex of $\bigcup \mathcal{B}$.

(b) $x \in \partial S^{\ominus\varepsilon}$ and there is a neighbourhood N_x of x such that $N_x \cap \partial S^{\ominus\varepsilon} \subseteq \bigcup \mathcal{B}$. We show again that x must be a vertex of $\bigcup \mathcal{B}$ that belongs to $A(S, \varepsilon, \mathcal{B})$. We have $x \in N_x \cap \partial S^{\ominus\varepsilon}$, thus x belongs to $\bigcup \mathcal{B}$. From the definition of x , we also know that $x \in (\bigcup \mathcal{B})^c$, therefore $x \in \partial \bigcup \mathcal{B}$. Let \mathcal{B}_x be a (cardinal minimum) collection of balls in \mathcal{B} that locally represents $\bigcup \mathcal{B}$ around x . We must show that \mathcal{B}_x defines either an elliptic pencil of dimension 1, or a parabolic pencil. Recall that as per Lemma 4.5, if every neighbourhood of x meets two locally distinct connected components of $(\bigcup \mathcal{B})^c$, then we know that x is a vertex of $\bigcup \mathcal{B}$. See Figure 4.18.a. We address the case where $(\bigcup \mathcal{B})^c$ only has one unique connected component in N_x . Refer to Figure 4.18.b for an illustration. Because $\bigcup \mathcal{B}$ covers all of $\partial S^{\ominus\varepsilon}$ in N_x , that connected component is either entirely in the interior of $S^{\ominus\varepsilon}$, or entirely in the interior of $(S^{\ominus\varepsilon})^c$. In the latter case, it would imply that $\bigcup \mathcal{B}$ entirely covers $S^{\ominus\varepsilon}$ in N_x , which would contradict the definition of x . Thus necessarily, that connected component is in the interior

of $S^{\ominus\varepsilon}$. This in turn implies that $(S^{\ominus\varepsilon})^c \cap N_x \subseteq \bigcup \mathcal{B}_x$. Now since b_* is a candidate, and by definition of N_x , we know that $S^{\ominus\varepsilon} \cap N_x \subseteq b_* \cup (\bigcup \mathcal{B}_x)$. It follows that $N_x \subseteq b_* \cup (\bigcup \mathcal{B}_x)$ and that $\text{coni}^*(x, \{c(b_*)\} \cup c(\mathcal{B}_x)) = \mathbb{R}^2$. As in the previous case, we deduce that the degree of x in $\partial \bigcup \mathcal{B}$ is null, and that x is a vertex.

(c) $x \in \partial S^{\ominus\varepsilon}$ and in every neighbourhood N_x of x , we have $N_x \cap \partial S^{\ominus\varepsilon} \not\subseteq \bigcup \mathcal{B}$. Thus x belongs to the closure of $(\partial S^{\ominus\varepsilon}) \setminus \bigcup \mathcal{B}$, which is a subset of $\overline{A(S, \varepsilon, \mathcal{B})}$.

In all three cases, we have $x \in \overline{A(S, \varepsilon, \mathcal{B})}$. Since b_* is critical for x , we finally obtain $b_* \in F_{\text{pcrit}}(S, \overline{A(S, \varepsilon, \mathcal{B})})$, which concludes the proof \square

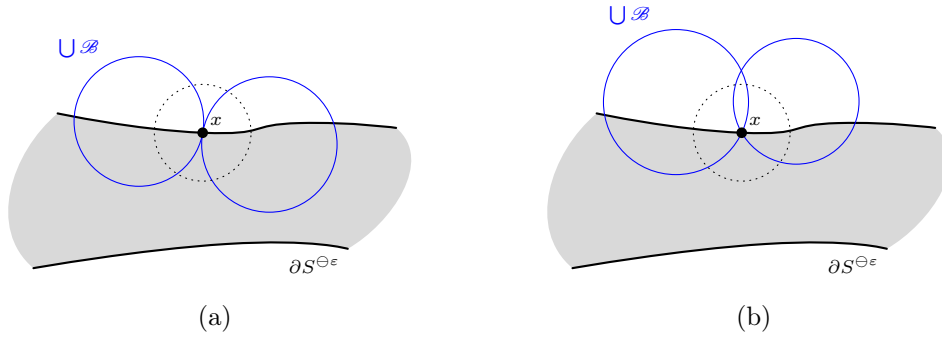


Figure 4.18 – Schematic illustration for the proof of Lemma 4.23. In the neighbourhood of x , $\bigcup \mathcal{B}$ is locally represented by the blue balls, and its complement admits (a) several connected components, or (b) one unique connected component.

In conclusion, we have just proved Theorem 8. As a consequence, we now have a characterization of T -maximal candidates that is much more suited for an algorithmic computation.

Remark 55. Note that if b_1 and b_2 both are critical for $A(S, \varepsilon, \mathcal{B})$, then b_2 remains critical for $A(S, \varepsilon, \mathcal{B} \cup \{b_1\})$. Indeed, recall first that when two balls are critical for the same set, they are T -incomparable, as per Remark 53. Thus b_2 can potentially remain critical. Then, notice that there must exist $x_2 \in \text{pcrit}^{-1}(b_2) \cap \overline{A(S, \varepsilon, \mathcal{B})}$. Because b_1 and b_2 are T -incomparable, we have by definition of pcrit^{-1} that $x_2 \notin b_1$. In addition, we have $b_1 \subseteq b_2 \cup C(S, b_2, +)$. It follows that $x_2 \in \overline{A(S, \varepsilon, \mathcal{B})} \setminus b_1 \subseteq \overline{A(S, \varepsilon, \mathcal{B} \cup \{b_1\})} \subseteq \overline{A(S, \varepsilon, \mathcal{B})} \cup \partial b_1 \subseteq b_2 \cup C(S, b_2, +)$. Thus b_2 remains critical for $A(S, \varepsilon, \mathcal{B} \cup \{b_1\})$.

4.6 Computing critical balls

Our proposed algorithm (that we detail in Section 4.7) requires the ability to compute critical balls for various constraints. In this section we explain how this computation can always be done with a few simple predicates. We show in the later Chapter 5, and in a broader context, that those predicates can be implemented, and focus here on the reduction to those predicates. First we show that we can reduce to elliptic 1-pencils in section 4.6.1. We then

detail in sections 4.6.2 and 4.6.3 how we can break down the set constraint of interest, that is $A(S, \varepsilon, \mathcal{B})$, into simpler primitive constraints, and how to process these primitives.

4.6.1 Reduction to elliptic pencils

We prove in this section that given a non-empty subset A in \mathring{S} , finding a medial ball b of S that is critical for A can be accomplished by looking for critical balls in specific pencil segments. This is due to the fact that the medial axis is a collection of pencil segments, and that looking for a critical ball can be performed locally. Indeed, assume we are looking for a critical ball to some constraint set A within $[b_-b_+]$, a specific pencil segment of the medial axis. One first necessary condition to verify, would be that there are no critical balls to A within the balls T -smaller than the pencil segment. Assume that this is the case, and denote by \mathcal{P} the (elliptic) pencil generated by b_- and b_+ . Then, we claim that it is equivalent to search the range $[b_-b_+[$ (b_- inclusive, b_+ exclusive) for a ball critical to A in the whole shape S , or to search that range for a ball critical to a certain restriction of A , in the shape $\bigcup \mathcal{P}$. Note that the balls in $[b_-b_+]$ are part of the medial axis of both shapes, and we make two assumptions in our previous claim: (1) segment $[b_-b_+]$ does not contain any branching, and (2) the orientation of the medial axes agree on segment $[b_-b_+]$.

More formally, by no branching we mean that for all balls b in the open range $]b_-b_+[$, the center of b must have degree 2. Said otherwise, segment $[b_-b_+]$ must not contain any vertex of T in its interior. This immediately implies that b_- and b_+ can be ordered with respect to T . Assume for now that we have $b_- \leq_T b_+$. The segment $[b_-b_+]$ uniquely defines a 1-pencil \mathcal{P} on which we can define a total order $T^{\mathcal{P}}$ derived from T , simply imposing that $b_- \leq_{T^{\mathcal{P}}} b_+$. Thus, any two balls of \mathcal{P} can be ordered according to $T^{\mathcal{P}}$, and if the two balls are on the segment $[b_-b_+]$ their order according to T is the same as their order according to $T^{\mathcal{P}}$. See Figure 4.19 for an illustration.




Figure 4.19 – The order $T^{\mathcal{P}}$ on \mathcal{P} is inherited from the order T defined on the medial axis of the three balls (in blue). We have $b_1 \leq_{T^{\mathcal{P}}} b_- \leq_{T^{\mathcal{P}}} b_2$. Additionally, we have both $b_- \leq_T b_+$ and $b_- \leq_{T^{\mathcal{P}}} b_+$.

Then, we have the following.

Lemma 4.24. *Let A be a non-empty subset of \mathring{S} . Let $[b_-b_+]$ be a pencil segment of $\text{MA}(S)$ with no branching and denote by \mathcal{P} the 1-pencil spanned by b_- and b_+ . Assume that there are no critical balls for A among the balls strictly T -smaller than b_- . Then a ball b on $[b_-b_+[$ (b_- inclusive, b_+ exclusive) is critical for A in S if and only if it is critical for $A \cap (b_- \cup b_+)$ in $\bigcup \mathcal{P}$.*

Proof. Recall that for a shape X that admits a partial T -order, and a subset $Y \subseteq X$, we denote by $F_{\text{pcrit}}(X, Y)$ the collection of balls

$$F_{\text{pcrit}}(X, Y) = \{b' \text{ medial in } X \mid \exists y \in Y, b' \text{ critical for } y\}$$

Let $A' = A \cap (b_- \cup b_+)$. We prove the series of equalities

$$\begin{aligned} F_{\text{pcrit}}(S, A) \cap \{b' \text{ medial in } S \mid b' <_T b_+\} \\ = F_{\text{pcrit}}(S, A) \cap [b_-b_+[\end{aligned} \quad (4.2)$$

$$= F_{\text{pcrit}}\left(\bigcup \mathcal{P}, A'\right) \cap [b_-b_+[\quad (4.3)$$

$$= F_{\text{pcrit}}\left(\bigcup \mathcal{P}, A'\right) \cap \{b' \text{ medial in } \bigcup \mathcal{P} \mid b' <_{T\mathcal{P}} b_+\} \quad (4.4)$$

Hence, if one of the collection of balls $\overline{F_{\text{pcrit}}(S, A)}$ or $\overline{F_{\text{pcrit}}(\bigcup \mathcal{P}, A')}$ admits a T -minimal ball in the range $[b_-b_+[$, then that ball is also T -minimal for the other collection, which concludes the proof.

For the first equality (4.2), recall that we assumed there were no ball critical for A that are T -smaller than b_- , and no branching of $\text{MA}(S)$ along the pencil segment $[b_-b_+]$. Hence any potential critical ball for A strictly T -smaller than b_+ must lie in $[b_-b_+[$.

For the second equality (4.3), consider $b \in [b_-b_+[$. Note that b is medial in both S and $\bigcup \mathcal{P}$. We proceed by double inclusion. For the first inclusion, suppose that additionally $b \in F_{\text{pcrit}}(S, A)$. Then let $x \in A$ such that b is critical for x . We have $x \in b \subseteq b_- \cup b_+$, thus $x \in A'$. We show that b must be critical for x in $\bigcup \mathcal{P}$, which thus implies that $b \in F_{\text{pcrit}}(\bigcup \mathcal{P}, A')$. Let $b_{\mathcal{P}}$ medial in $\bigcup \mathcal{P}$ such that $b_{\mathcal{P}} >_{T\mathcal{P}} b$. Let $b' \in [b_-b_+]$ such that $b <_{T\mathcal{P}} b' \leq_{T\mathcal{P}} b_{\mathcal{P}}$. Note that we still have $b <_T b' \leq_T b_+$ using the T -order over S . By criticality of b for x in S , we know by statement (iv) of Definition 18, that $x \notin b' \cup C(S, b', +)$. Because $x \in b$ we thus obtain $x \in b \setminus b'$. By definition of b' , we have $b <_{T\mathcal{P}} b'$. By definition of T -small components in $\bigcup \mathcal{P}$ and Property 4.1, we have $b \setminus b' \subseteq C(\bigcup \mathcal{P}, b', -) \subseteq C(\bigcup \mathcal{P}, b_{\mathcal{P}}, -)$. Therefore $x \in C(\bigcup \mathcal{P}, b_{\mathcal{P}}, -)$. Although Proposition 4.10 does not include shapes defined by a complete pencils such as $\bigcup \mathcal{P}$, we actually have the stronger property that $C(\bigcup \mathcal{P}, b_{\mathcal{P}}, -)$, $b_{\mathcal{P}}$, and $C(\bigcup \mathcal{P}, b_{\mathcal{P}}, +)$ are pairwise disjoint for any $b_{\mathcal{P}} \in \mathcal{P}$. This can easily be proved with Property 2.1. Thus $x \notin b_{\mathcal{P}} \cup C(\bigcup \mathcal{P}, b_{\mathcal{P}}, +)$. Because this holds for every ball $b_{\mathcal{P}} >_{T\mathcal{P}} b$, we deduce by statement (iv) of Definition 18 that b is critical for x in $\bigcup \mathcal{P}$. In conclusion, $b \in F_{\text{pcrit}}(\bigcup \mathcal{P}, A')$.

Conversely, suppose now that $b \in F_{\text{pcrit}}(\bigcup \mathcal{P}, A')$. Let $x \in A'$ such that b is critical for x in $\bigcup \mathcal{P}$. Since $A' \subseteq A$, $x \in A$. We show that b is critical for x in S , which implies that $b \in F_{\text{pcrit}}(S, A)$. Let b_S a medial ball of S such that $b_S >_T b$. As above, consider $b' \in [b_-b_+[$

such that $b <_T b' \leq_T b_S$. A similar reasoning yields $x \in b \setminus b' \subseteq C(S, b', -) \subseteq C(S, b_S, -)$. Because we assumed $A \subseteq \mathring{S}$, thus $x \in \mathring{S}$ and Proposition 4.10 implies that $x \notin b_S \cup C(S, b_S, +)$. Because this is true for every $b_S >_T b$, we deduce that b is critical for x in S and therefore that $b \in F_{\text{pcrit}}(S, A)$. Equality (4.3) immediately follows.

For the third and last equality (4.4), recall that the medial axis of $\bigcup \mathcal{P}$ is the centers' space of pencil \mathcal{P} . Denote by c_+ (resp. c_-) the center of b_+ (resp. b_-). Equivalently, the medial axis is the line through c_- and c_+ . Taking into account the orientation and partial order $T^{\mathcal{P}}$, said line is oriented from c_- towards c_+ . Hence the collection of balls $T^{\mathcal{P}}$ -smaller than b_+ corresponds to balls centered on the half-line originating from c_+ and going through c_- . Excluding the point c_+ yields the collection of balls strictly $T^{\mathcal{P}}$ -smaller than b_+ . Consider now a ball $b \in \mathcal{P}$ strictly $T^{\mathcal{P}}$ -smaller than b_- . By contradiction, suppose that $b \in F_{\text{pcrit}}(\bigcup \mathcal{P}, A')$. Then, let $x \in A'$ such that b is critical for x in $\bigcup \mathcal{P}$. We have $x \in \text{pcrit}^{-1}(b) = b \setminus \bigcup_{b' >_T b} b' \subseteq b \setminus (b_- \cup b_+) \subseteq b \setminus A'$. Hence $x \notin A'$ which is absurd. Therefore, $F_{\text{pcrit}}(\bigcup \mathcal{P}, A') \subseteq \{b' \mid b' \geq_T b_-\}$, which easily yields equality (4.4). \square

Owing to the lemma above, looking for a ball critical for A within $[b_- b_+[$ can be reduced to looking for a ball critical for $A \cap (b_- \cup b_+)$ in the complete pencil \mathcal{P} , and then verifying that the critical ball lies within the range $[b_- b_+[$. The two assumptions required to apply the lemma and perform this reduction, are that the medial axis must locally be a path without branching, and that there must be no critical ball in the T -smaller parts of the medial axis. In practice, both of these assumptions can easily be fulfilled. Because the medial axis of a finite union of balls is a finite collection of pencil segments, it is sufficient to process each of these pencil segments to avoid any branching within the interior of the segment. Because the pencil segments incident to the leaves of the medial axis necessarily satisfy the “no T -smaller critical ball” assumption, there always exist some pencil segment we can apply the reduction to. Because the lemma states an equivalence, we then know whether those pencil segments contain a critical ball or not. If they contain one such critical ball, then any other pencil segment that is T -larger cannot contain a critical ball. If they do not contain a critical ball, then we know, up to medial axis branching, that the incident pencil segment of T -larger balls thus satisfies the conditions of Lemma 4.24. Hence, by processing the pencil segments of the medial axis in an appropriate order and carefully applying the reduction, it is possible to compute every critical ball of a given constraint.

§ Topological ordering. Because branching of the medial axis must still be taken into account, this motivates sweeping the medial axis in a topological order. With our partial ordering on $\text{MA}(S)$, a topological order of its vertices (v_1, \dots, v_{n+1}) is such that for $i \leq j$, then either $v_i \leq_T v_j$, or v_i and v_j are T -incomparable. Besides the root, each vertex is incident to exactly one pencil segment composed of T -larger points of $\text{MA}(S)$, hence any topological ordering of vertices induces an ordering of pencil segments $([v_1 \hat{v}_1], \dots, [v_n \hat{v}_n])$, where $\hat{v}_i \in \{v_1, \dots, v_{n+1}\}$. This ordering of pencil segments thus satisfies that for $i < j$, then either $v_i <_T \hat{v}_i \leq_T v_j <_T \hat{v}_j$, or \hat{v}_i and v_j are T -incomparable.

This property of topological orders is in particular suited to the computation of critical

balls. Indeed, when looping over pencil segments in a topological order, we are guaranteed to have processed every pencil segments T -smaller than the current one. Using a topological order thus naturally helps in satisfying and verifying the conditions of Lemma 4.24.

Remark 56. Note that this lemma states an equivalence when looking for a critical ball in a range that exclude the T -large endpoint. After searching every pencil in that manner, we will have inspected every medial ball, except for the root of T itself.

4.6.2 Simplifying the constraint set $A(S, \varepsilon, \mathcal{B})$

Given a T -small $(0, \varepsilon)$ -ball approximation \mathcal{B} , we essentially want to investigate each pencil segment of the medial axis of S in a topological order, and determine whether or not they contain a T -maximal candidate to \mathcal{B} , that is a ball critical for $A(S, \varepsilon, \mathcal{B})$. We show in this section that when processing a given pencil segment, we can ignore a large part of $A(S, \varepsilon, \mathcal{B}) \cap \partial S^{\ominus \varepsilon}$ and only focus on points that are, in a certain sense, close to the pencil segment.

Note that it is very natural to be able to ignore most of $A(S, \varepsilon, \mathcal{B})$. Because the assumption $\overline{A(S, \varepsilon, \mathcal{B})} \subseteq \mathring{S}$ holds, we already know by Property 4.16 that there exists a point $x \in A(S, \varepsilon, \mathcal{B})$ such that criticality for $A(S, \varepsilon, \mathcal{B})$ and criticality for x will be equivalent within the local branch. To take into account the branching of the medial axis, it is sufficient to take one point for each ball that is critical for the complete set constraint, which yields a finite number of singletons. Thus we can theoretically reduce a set with infinitely many points to a discrete, finite set. In practice, one major difficulty is to identify, a priori, to what points can a given set be reduced to. Precisely pinpointing these points prior to computation seems difficult, it is however easy to narrow down their location. Indeed, assume b is critical for $A(S, \varepsilon, \mathcal{B})$. Because Property 4.16 holds, we know there exists $x \in \overline{A(S, \varepsilon, \mathcal{B})} \cap \text{pcrit}^{-1}(b)$. In particular, if $b_+ >_T b$, then $x \in C(S, b_+, -) \cup b_+$. By criticality of b for x , we have $x \notin b_+$ and we actually have $x \in C(S, b_+, -)$. It follows that when looking for a critical ball within the range $[b_-, b_+]$, we can consider only $A(S, \varepsilon, \mathcal{B}) \cap C(S, b_+, -)$, or equivalently by Proposition 4.10 $A(S, \varepsilon, \mathcal{B}) \setminus (b_+ \cup C(S, b_+, +))$. See Figure 4.20. In this section, we describe how we can simply obtain a superset of $A(S, \varepsilon, \mathcal{B}) \cap C(S, b_+, -)$.

§ Overview We previously investigated the boundary of $S^{\ominus \varepsilon}$ in Section 3.2 and showed that it can be described as a finite collection of circular arcs. Our strategy to simplify the constraint set $A(S, \varepsilon, \mathcal{B})$ is to ignore specific arcs of $\partial S^{\ominus \varepsilon}$. Indeed the constraint set is composed of portions of arcs of $\partial S^{\ominus \varepsilon}$, as well as vertices of $\bigcup \mathcal{B}$. To determine which arcs can be ignored and which arcs should be taken into account, we map these arcs to pencil segments and vertices of $\text{MA}(S)$. We argue that the mapping we introduce enables us to extend the partial T -order to these arcs, and that we can still obtain balls critical for $A(S, \varepsilon, \mathcal{B})$ after simply filtering out arcs that are T -larger or T -incomparable than a specific threshold.




Figure 4.20 – If b is critical for $A(S, \varepsilon, \mathcal{B})$, the set of constraints can be reduced to $A(S, \varepsilon, \mathcal{B}) \cap C(S, b_+, -)$ (blue arcs and point).

§ Projection mapping and arc ordering Relying on the mapping π introduced in Section 4.2, we can project any arc e of $\partial S^{\ominus \varepsilon}$ onto medial balls of S . Indeed we have $e \subseteq \mathring{S}$ so its projection $\pi(e)$ is well defined, and is a continuous path. We immediately obtain that $e \subseteq \bigcup \pi(e)$. Because the centers of $\pi(e)$ satisfy $c(\pi(e)) \subseteq \text{MA}(S)$, we can subdivide $\pi(e)$ according to what pencil segments and vertices of $\text{MA}(S)$ it meets, which induces a subdivision of e . Those subdivisions can be T -ordered.

One possibility is to compute the subdivision of every arc e of the boundary $\partial S^{\ominus \varepsilon}$. We can then directly work with those subdivided arcs, and directly use the extended T -order of these subdivided arcs. Due to the properties of the projection π , this computation is nearly equivalent to verifying the non-empty intersections between arc e and the face joins. Indeed, a face join may encompass several pencil segments and vertices of $\text{MA}(S)$, it is thus necessary to subdivide face joins to retrieve the subdivided arcs. See for instance the face join represented in Figure 4.4, it covers three vertices and two pencil segments of the medial axis. In fact, our description of the boundary of $S^{\ominus \varepsilon}$ implies that one arc of the boundary is always contained within a single face join, because each arc originates from a single face of ∂S .

Alternatively, it is sufficient to simply remember which (subdivided) face joins are crossed over by arc e . This lets us work with the original arcs of $\partial S^{\ominus \varepsilon}$, without the need to explicitly compute each subdivision of those arcs. Specifically, it is sufficient to remember for each pencil segment and vertex of $\text{MA}(S)$, which circular arcs of $\partial S^{\ominus \varepsilon}$ are projected onto those pencil segments and vertices.

Then, consider b a medial ball of S centered at a vertex of the medial axis $\text{MA}(S)$. Let e be a circular arc of $\partial S^{\ominus \varepsilon}$. We say that e is T -smaller than b , and denote $e \leq_T b$, if there is some ball $b_m \in \pi(e)$ such that $c(b_m) \leq_T c(b)$.

Remark 57. The definition above can extend to any ball b medial in S , instead of being restricted to simply balls centered at vertices of $\text{MA}(S)$. That level of detail will not be

necessary for what follows.

§ Filtering (arc) constraints Consider a pencil segment $[b_-b_+]$ of $\text{MA}(S)$. Suppose there exists b_{crit} , critical for $A(S, \varepsilon, \mathcal{B})$, with $b_- \leq_T b_{\text{crit}} <_T b_+$. Then b_{crit} must be critical for $A(S, \varepsilon, \mathcal{B}) \cap C(S, b_+, -)$. In order to simplify the computation of critical balls, we now show that there is a simple way to find a superset of $A(S, \varepsilon, \mathcal{B}) \cap C(S, b_+, -)$.

Let $\mathcal{V}(S, \varepsilon, \mathcal{B}, b_+)$ be the collection of vertices of $\bigcup \mathcal{B}$ that belong to both $A(S, \varepsilon, \mathcal{B})$ and $C(S, b_+, -)$, and let $\mathcal{E}(S, \varepsilon, \mathcal{B}, b_+)$ the collection of circular arcs of $\partial S^{\ominus \varepsilon} \setminus \bigcup \mathcal{B}$ that are T -smaller than b_+ ,

$$\mathcal{E}(S, \varepsilon, \mathcal{B}, b_+) = \left\{ e \subseteq (\partial S^{\ominus \varepsilon}) \setminus \bigcup \mathcal{B} \mid e \text{ circular arc, } e \leq_T b_+ \right\}$$

Assuming that $\partial S^{\ominus \varepsilon}$ has been pre-processed in advance, it is easy to determine $\mathcal{E}(S, \varepsilon, \mathcal{B}, b_+)$ as we loop over the list of pencils in topological order. We let $A(S, \varepsilon, \mathcal{B}, b_+)$ be the union of every arcs and vertices in $\mathcal{E}(S, \varepsilon, \mathcal{B}, b_+)$ and $\mathcal{V}(S, \varepsilon, \mathcal{B}, b_+)$,

$$A(S, \varepsilon, \mathcal{B}, b_+) = \left(\bigcup \mathcal{E}(S, \varepsilon, \mathcal{B}, b_+) \right) \cup \left(\bigcup \mathcal{V}(S, \varepsilon, \mathcal{B}, b_+) \right)$$

Lemma 4.25. *Let \mathcal{B} be a T -small $(0, \varepsilon)$ -ball approximation of S . Consider a pencil segment $[b_-b_+]$, and a medial ball b with $b_- \leq_T b <_T b_+$. Then b is critical for $A(S, \varepsilon, \mathcal{B})$ if and only if:*

- $A(S, \varepsilon, \mathcal{B}, b_+) \neq \emptyset$, and
- b is critical for $A(S, \varepsilon, \mathcal{B}, b_+)$.

Remark 58. If the union happens to be empty, then the subset $A(S, \varepsilon, \mathcal{B})$ does not admit any critical ball in the range $[b_-b_+]$.

Proof. We argue that $A(S, \varepsilon, \mathcal{B}, b_+)$ is a superset of $A(S, \varepsilon, \mathcal{B}) \cap C(S, b_+, -)$, which concludes the proof. Indeed, recall that $A(S, \varepsilon, \mathcal{B})$ is composed of some vertices of $\bigcup \mathcal{B}$, and $(\partial S^{\ominus \varepsilon}) \setminus \bigcup \mathcal{B}$. By definition, $\mathcal{V}(S, \varepsilon, \mathcal{B}, b_+)$ contains every vertices of $\bigcup \mathcal{B}$ in $A(S, \varepsilon, \mathcal{B}) \cap C(S, b_+, -)$. Remains to prove that $\mathcal{E}(S, \varepsilon, \mathcal{B}, b_+)$ covers $(\partial S^{\ominus \varepsilon}) \cap (\bigcup \mathcal{B})^c \cap C(S, b_+, -)$.

Consider a circular arc e of $\partial S^{\ominus \varepsilon}$. We show that if $e \notin \mathcal{E}(S, \varepsilon, \mathcal{B}, b_+)$, then $e \subseteq b_+ \cup C(S, b_+, +)$. Thus we have $(\partial S^{\ominus \varepsilon}) \setminus (\bigcup \mathcal{E}(S, \varepsilon, \mathcal{B}, b_+)) \subseteq b_+ \cup C(S, b_+, +)$, which implies by Proposition 4.10 that $(\partial S^{\ominus \varepsilon}) \cap C(S, b_+, -) \subseteq \bigcup \mathcal{E}(S, \varepsilon, \mathcal{B}, b_+)$. Let $x \in e$ and its projection $\pi(x) \in \text{MA}(S)$. Denote by b_x the medial ball of S centered at $\pi(x)$. Since $e \notin \mathcal{E}(S, \varepsilon, \mathcal{B}, b_+)$, we have either $b_+ <_T b_x$, or b_+ and b_x are T -incomparable. Therefore, $x \in b_+ \cup C(S, b_+, +)$. This applies for every $x \in e$, thus $e \subseteq b_+ \cup C(S, b_+, +)$. \square

The lemma above suggests that when computing a critical ball, we can maintain a local set of constraints that is specific to the branch of the medial axis currently being processed. This

local constraint set is, potentially, much smaller than the complete set $A(S, \varepsilon, \mathcal{B})$. Because we can decompose $\partial S^{\ominus \varepsilon}$ into circular arcs, $A(S, \varepsilon, \mathcal{B})$ is a union of arcs and singletons. As we shall see in the following section, it is convenient to process these arcs and singletons separately. Having fewer arcs to process per pencil will thus be beneficial.

4.6.3 Reducing to primitive constraints

As per the previous sections, the constraint sets whose critical balls we have to compute can always be described as a finite union of circular arcs (with or without their endpoints) and singletons. We now argue that it is possible to process each of these primitive arcs and singletons separately. The balls critical for the overall union can indeed all be deduced from the balls critical for the individual primitive constraints. Formally, we rely on the following.

Lemma 4.26. *Let A_1, \dots, A_k be k subsets of \mathring{S} and consider $A = \bigcup_{i=1}^k A_i$. Then a medial ball b is critical for A if and only if it is T -minimal among the collection of balls critical for any A_i , $1 \leq i \leq k$.*

Proof. Recall that for any subset $X \subseteq \mathring{S}$, we have

$$F_{\text{pcrit}}(S, X) = \{b \text{ medial in } S \mid \exists x \in X, b \text{ critical for } x\}$$

By definition, a ball b is critical for X if and only if it is T -minimal in $\overline{F_{\text{pcrit}}(S, X)}$. Notice that $\overline{F_{\text{pcrit}}(S, A)} = \overline{\bigcup_{i=1}^k F_{\text{pcrit}}(S, A_i)} = \bigcup_{i=1}^k \overline{F_{\text{pcrit}}(S, A_i)}$. The property immediately follows. \square

When looking for a ball critical for $A = \bigcup_{i=1}^k A_i$ within a pencil segment $[b_- b_+]$, we can thus look for the balls critical for the subsets A_i in that range. There are at most k such critical balls, and their T -minimum will be critical for A . If none of the k subsets yield a critical ball, then A does not admit a critical ball within the pencil segment. Because the constraints subsets we have to handle are finite unions of circular arcs and singletons, we will see in Chapter 5 how to compute critical balls for these two types of constraints.

4.7 Algorithm

We are now ready to specify our greedy algorithm. Recall that we want to build a sequence of T -small $(0, \varepsilon)$ -ball approximations by computing T -maximal candidates. From the previous sections, this can be done by computing critical balls for the appropriate local constraints, and these local constraints can always be decomposed into a finite union of singletons and circular arcs.

4.7.1 Specification

Our proposed algorithm is based on a loop over the collection of all pencil segments of $MA(S)$ in a topological order (as specified in Section 4.6.1). As we loop over the pencil segments, we maintain a collection of medial balls \mathcal{B} which is a T -small $(0, \varepsilon)$ -ball approximation, while looking for T -maximal candidates to add to said collection.

When we process a pencil segment, we compute a collection of primitive constraints to pass on to the next incident pencil segment. These primitives are as specified by Lemma 4.25. A ball in the pencil will be a T -maximal candidate to \mathcal{B} if and only if it is critical for (the union of) these primitive constraints (Lemma 4.26). If no critical ball is found within the currently processed pencil segment, the set of constraints it will pass on to its incident T -larger pencil segment, is the collection of all constraints it itself inherited from incident T -smaller pencil segments, plus new constraints specific to the current pencil segment. When, a critical ball b is found, we require an additional updating step. We clip constraint primitives against the complement of b . We also investigate the new vertices of $(\bigcup \mathcal{B}) \cup b$, and test if $S^{\ominus \varepsilon}$ is properly covered by the collection of balls in their neighbourhoods. If not, the infringing vertices are added as primitive constraints. The overall approach is summarized in Algorithm 1. See Figure 4.21.

Algorithm 1 Greedy $(0, \varepsilon)$ -ball approximation

Input: A finite union of balls S

Output: A $(0, \varepsilon)$ -ball approximation \mathcal{B} of S

- 1: Compute a topological ordering of $MA(S)$
 - 2: Compute $\partial S^{\ominus \varepsilon}$ and find the T -order of its circular arcs
 - 3: $\mathcal{B} \leftarrow \emptyset$
 - 4: **Loop** over pencil segments of $MA(S)$ in topological order
 - 5: Retrieve incident constraints
 - 6: Include pencil specific constraints
 - 7: Search for a critical ball in the pencil
 - 8: **If** a critical ball b is found **then**
 - 9: $\mathcal{B} \leftarrow \mathcal{B} \cup \{b\}$
 - 10: Update constraints
 - 11: **end If**
 - 12: **end Loop**
 - 13: **If** the set of constraints is non-empty **then**
 - 14: $\mathcal{B} \leftarrow \mathcal{B} \cup \{b_{\text{root}}\}$
 - 15: **end If**
 - 16: **Return** \mathcal{B}
-

4.7.2 Complexity

As per the previous sections, we know that Algorithm 1 will find T -maximal candidates, and will eventually converge to a $(0, \varepsilon)$ -ball approximation of S . If the algorithm terminates, we




Figure 4.21 – Illustration of the different steps of the algorithm.

have the guarantee that it has found a $(0, \varepsilon)$ -ball approximation of S , but we do not yet know whether it will terminate (in finite time). We show that termination will occur in polynomial time.

Theorem 9. Under the Real-RAM model, Algorithm 1 terminates in time polynomial with respect to $|\text{MA}(S)|$.

Proof. Let $n = |\text{MA}(S)|$. First, we show that Algorithm 1 outputs a collection $\mathcal{B}_{\text{algo}}$ with size at most $2n$, and then that the complexity is polynomial in n .

For any partial T -small $(0, \varepsilon)$ -ball approximation \mathcal{B} , let b be T -maximal in \mathcal{B} . There is a unique pencil segment incident to b that contains balls T -larger than b . Let $b_+ >_T b$ be the T -large endpoint of that pencil segment. We show that b_+ always preserves the coverage property of $S^{\ominus \varepsilon}$ in the T -small component. Although b_+ may not strictly increase the coverage of $S^{\ominus \varepsilon}$, and thus may not be a proper candidate to \mathcal{B} , this still implies that there can be at most two medial balls from the same pencil segment in $\mathcal{B}_{\text{algo}}$. Thus $|\mathcal{B}_{\text{algo}}| \leq 2n$.

Let $\mathcal{B}_+ = \mathcal{B} \cup \{b_+\}$. Because b and b_+ belong to the same pencil, $C(S, \mathcal{B}_+, -)$ and $C(S, \mathcal{B}, -)$ only differ in the domain of S covered by the pencil segment $[bb_+]$. This implies the equalities $(C(S, \mathcal{B}_+, -) \cup (\bigcup \mathcal{B}_+)) \setminus C(S, \mathcal{B}, -) = (C(S, b_+, -) \cup b_+) \setminus C(S, b, -) = \bigcup_{b \leq_T b' \leq_T b_+} b' = b \cup b_+$.

In particular, $C(S, \mathcal{B}_+, -) \subseteq C(S, \mathcal{B}, -) \cup b \cup b_+$. Hence, we obtain

$$S^{\ominus \varepsilon} \cap C(S, \mathcal{B}_+, -) \subseteq S^{\ominus \varepsilon} \cap (C(S, \mathcal{B}, -) \cup b \cup b_+) \subseteq \left(\bigcup \mathcal{B} \right) \cup b \cup b_+ = \bigcup \mathcal{B}_+$$

thus b_+ preserves the coverage of $S^{\ominus \varepsilon}$ in the T -small component of collection \mathcal{B} . Therefore, we indeed have $|\mathcal{B}_{\text{algo}}| \leq 2n$.

We now analyse the complexity. Computing a topological ordering [Cor+01] is linear in n . The boundaries of S and $S^{\ominus \varepsilon}$ have a (worst case) combinatorial complexity quadratic in n . It follows that there can be at most $|\partial S| \times |\text{MA}(S)| = \mathcal{O}(n^3)$ distinct subdivided face joins. Figuring out the T -order over $\partial S^{\ominus \varepsilon}$ thus takes polynomial time. Enforcing a single constraint takes constant time (see Chapter 5), hence the time spent to search for a critical ball depends on the number of constraints. This number of constraints cannot exceed the combinatorial complexity of the arrangement of the boundaries of $S^{\ominus \varepsilon}$ and $\bigcup \mathcal{B}_{\text{algo}}$. The latter is (worst case) quadratic in the size of $\mathcal{B}_{\text{algo}}$, which is itself linear in n . Hence the number of constraints is at worst quadratic in n . Regarding the update procedure, clipping the constraints is clearly linear in the number of constraints. Recall that in Section 3.3.1 we explained how to capture $S^{\ominus \varepsilon}$ through an arrangement. Using a similar approach with arrangements, we can investigate which vertices of $\partial \bigcup \mathcal{B}_{\text{algo}}$ need to be taken as constraint or not. Note that the number of curves (and circles) to include in that arrangement is again linear with respect to n , hence using that particular structure implies polynomial costs. Thus, Algorithm 1 is overall polynomial in n . \square

4.7.3 Optimality

In order to prove that our algorithm reaches an optimal, we rely on the lemma below.

Lemma 4.27. *Consider a $(0, \varepsilon)$ -ball approximation \mathcal{B} . Let $\mathcal{B}_- \subsetneq \mathcal{B}$ be a partial T -small $(0, \varepsilon)$ -ball approximation, and let b be any T -maximal candidate to \mathcal{B}_- . Then $\mathcal{B} \setminus \mathcal{B}_-$ contains a candidate to \mathcal{B}_- that is T -smaller than or equal to b .*

Proof. Let $\mathcal{B}_+ = \mathcal{B} \setminus \mathcal{B}_-$. First we prove that \mathcal{B}_+ always contains candidates to \mathcal{B}_- , and then that one of these candidates is T -smaller than or equal to b .

By contradiction, assume that \mathcal{B}_+ is void of candidate to \mathcal{B}_- . Consider $b_+ \in \mathcal{B}_+$, T -minimal in \mathcal{B}_+ . Note that because \mathcal{B}_- is only a partial $(0, \varepsilon)$ -ball approximation, $\bigcup \mathcal{B}_-$ alone does not entirely covers $S^{\ominus \varepsilon}$. It follows that \mathcal{B}_+ is necessarily non-empty, and always admits such a T -minimal ball b_+ . By T -minimality of b_+ , for all $b' \in \mathcal{B}_+$, we have $b' \subseteq b_+ \cup C(S, b_+, +)$. Thus $\bigcup \mathcal{B}_+ \subseteq b_+ \cup C(S, b_+, +)$. By Proposition 4.10 we deduce that

$$S^{\ominus \varepsilon} \cap C(S, b_+, -) \subseteq S^{\ominus \varepsilon} \cap \left(\bigcup \mathcal{B}_+ \right)^c = S^{\ominus \varepsilon} \setminus \bigcup \mathcal{B}_+$$

which we can rewrite as

$$S^{\ominus \varepsilon} \cap C(S, b_+, -) = S^{\ominus \varepsilon} \cap C(S, b_+, -) \setminus \bigcup \mathcal{B}_+ \quad (4.5)$$

By assumption, b_+ cannot be a candidate to \mathcal{B}_- , hence by Definitions 15 and 16 $(S^{\ominus\varepsilon} \cap C(S, b_+, -)) \setminus (b_+ \cup (\bigcup \mathcal{B}_-))$ is non empty. With the previous equality (4.5), we get the following development:

$$\begin{aligned} (S^{\ominus\varepsilon} \cap C(S, b_+, -)) \setminus (b_+ \cup (\bigcup \mathcal{B}_-)) &= (S^{\ominus\varepsilon} \cap C(S, b_+, -)) \setminus (b_+ \cup (\bigcup \mathcal{B}_-) \cup (\bigcup \mathcal{B}_+)) \\ &= (S^{\ominus\varepsilon} \cap C(S, b_+, -)) \setminus \bigcup \mathcal{B} \\ &\subseteq S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B} \end{aligned}$$

Hence, $S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B} \neq \emptyset$ which is impossible since \mathcal{B} is a $(0, \varepsilon)$ -ball approximation. Thus, \mathcal{B}_+ must contain a candidate to \mathcal{B}_- .

Because \mathcal{B}_- may have several distinct T -maximal candidates, \mathcal{B}_+ may only contain candidates T -incomparable to b . By contradiction assume \mathcal{B}_+ is void of candidate to \mathcal{B}_- T -smaller than or equal to b . Since \mathcal{B}_+ contains at least one candidate to \mathcal{B}_- , b cannot be centered at the root of T . By Proposition 4.1, any ball T -smaller than or equal to b will preserve the covering property of $S^{\ominus\varepsilon}$ in the T -small component. That is, if $b_- \leq_T b$ then $S^{\ominus\varepsilon} \cap C(S, \mathcal{B}_- \cup \{b_-\}, -) \subseteq (\bigcup \mathcal{B}_-) \cup b_-$. Hence \mathcal{B}_+ only contains balls that are either strictly T -larger than or T -incomparable to b . Therefore, there exists a medial ball $b_+ >_T b$, with b_+ also strictly T -smaller or T -incomparable to balls in \mathcal{B}_+ . By T -maximality of b , b_+ cannot be a candidate to \mathcal{B}_- . Once again, we have a ball b_+ , T -minimal in $\mathcal{B}'_+ = \mathcal{B}_+ \cup \{b_+\}$ which is not a candidate. The same development as above using Proposition 4.10 yields $S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B} \neq \emptyset$, which is still impossible. Therefore, \mathcal{B}_+ must contain a candidate to \mathcal{B}_- that is T -smaller than or equal to b . \square

Theorem 10. Algorithm 1 outputs a cardinal minimum $(0, \varepsilon)$ -ball approximation.

Proof. We denote by $\mathcal{B}_{\text{algo}}$ the $(0, \varepsilon)$ -ball approximation found by Algorithm 1. We number the balls of $\mathcal{B}_{\text{algo}}$ by b_1, \dots, b_k , such that for $i \leq j$, b_i was found before b_j . Consider any cardinal minimum $(0, \varepsilon)$ -ball approximation \mathcal{B}_{opt} . We assume without loss of generality that \mathcal{B}_{opt} only contains medial balls. Indeed, S is a finite union of balls, hence any ball $b \in \mathcal{B}_{\text{opt}}$ is wholly contained in a medial ball. Using consecutive substitutions, we want to build a finite sequence of collections of balls $\mathcal{B}_0, \dots, \mathcal{B}_k$ that satisfies the properties:

- (a) $\mathcal{B}_0 = \mathcal{B}_{\text{opt}}$,
- (b) \mathcal{B}_i is a $(0, \varepsilon)$ -ball approximation $\forall i \in \llbracket 0, k \rrbracket$
- (c) $|\mathcal{B}_{i+1}| = |\mathcal{B}_i|$, $\forall i \in \llbracket 0, k-1 \rrbracket$,
- (d) $\{b_1, \dots, b_i\} \subseteq \mathcal{B}_i$, $\forall i \in \llbracket 1, k \rrbracket$,

The idea is to replace one by one the balls of \mathcal{B}_{opt} by balls of $\mathcal{B}_{\text{algo}}$. If such a sequence exists, we immediately deduce that $|\mathcal{B}_{\text{algo}}| = |\mathcal{B}_{\text{opt}}|$, and $\mathcal{B}_{\text{algo}}$ is also cardinal minimum.

We proceed by induction. Assume that for $0 \leq i < k$, we have built $\mathcal{B}_0, \dots, \mathcal{B}_i$ with the above properties. Consider $b_{i+1} \in \mathcal{B}_{\text{algo}}$. Let $\mathcal{B}_- = \{b_1, \dots, b_i\}$. By construction,

$\mathcal{B}_- \subsetneq \mathcal{B}_i$. Let $\mathcal{B}_+ = \mathcal{B}_i \setminus \mathcal{B}_-$. \mathcal{B}_+ contains the balls that have not been replaced yet. Algorithm 1 guarantees that \mathcal{B}_- is a partial T -small $(0, \varepsilon)$ -ball approximation and that b_{i+1} is a T -maximal candidate for \mathcal{B}_- . Hence we can apply Lemma 4.27, \mathcal{B}_+ contains a ball b that is a candidate to \mathcal{B}_- and such that $b \leq_T b_{i+1}$. Then, let $\mathcal{B}_{i+1} = (\mathcal{B}_i \cup \{b_{i+1}\}) \setminus \{b\}$. \mathcal{B}_{i+1} satisfies both properties (c) and (d), we must prove that it also satisfies (b). To do so, it suffices to prove that both $C(S, b_{i+1}, -) \cap S^{\ominus\varepsilon}$ and $C(S, b_{i+1}, +) \cap S^{\ominus\varepsilon}$ are contained in $\bigcup \mathcal{B}_{i+1}$. Because b_{i+1} is a candidate to \mathcal{B}_- , we have $C(S, b_{i+1}, -) \cap S^{\ominus\varepsilon} \subseteq b_{i+1} \cup (\bigcup \mathcal{B}_-) \subseteq \bigcup \mathcal{B}_{i+1}$. Also, $b \leq_T b_{i+1}$, hence by Property 4.1 we have $C(S, b_{i+1}, +) \subseteq C(S, b, +)$. This implies $C(S, b_{i+1}, +) \cap S^{\ominus\varepsilon} \subseteq C(S, b, +) \cap S^{\ominus\varepsilon} \subseteq (\bigcup \mathcal{B}_i) \setminus b \subseteq \bigcup \mathcal{B}_{i+1}$. Thus, \mathcal{B}_{i+1} is a $(0, \varepsilon)$ -ball approximation, and $\mathcal{B}_{\text{algo}}$ is a cardinal minimum $(0, \varepsilon)$ -ball approximation. \square

Theorem 7 immediately follows from Theorems 9 and 10.

4.8 Conclusion

In order to reach Algorithm 1, we started from the intuition of candidate balls and visited the more general concept of critical balls. From the several properties of these balls, we were able to show that our greedy algorithm achieves a cardinal-minimum solution for a certain class of shapes. We now discuss some of the assumptions in relation to these shapes.

First we address the topic of the partial order T over the medial axis. From the description we provided back in Section 4.1, any point of the medial axis can be chosen as the root of T , and thus define an orientation and partial order. This has little bearing on our method, and the choice of this root is completely free. The resulting orientation and partial order will indeed have all the desired properties for our algorithm to run properly. In practice, changing the root may change the position of the chosen balls, but not the cardinal of the collection.

If we relax the assumption that $\text{MA}(S)$ is cycle-free, we can easily extend Algorithm 1 so as to handle cycles. Note that this extension does not preserve the guarantee of optimality. To see this, recall that $\text{MA}(S)$ can be seen as the embedding of a graph. In graph theory, a *feedback vertex set* of a graph $G = (V, E)$ is a subset F of the vertices that contains at least one vertex of any cycle in G . Computing a feedback vertex set of minimum cardinality is known to be NP-complete [Kar72]. Given such a feedback vertex set F , we can then start our algorithm from the collection \mathcal{B}_F of medial balls whose center is a vertex in F . This effectively punctures every cycle of the medial axis, and we can then define an orientation on $\text{MA}(S)$ by choosing a root from $\text{MA}(S) \setminus F$. From there, we can proceed with our method as usual after computing the initial constraint $A(S, \varepsilon, \mathcal{B}_F)$. Denote by $\mathcal{B}_{\text{algo}}$ the resulting collection of balls. The previous sections ensure that $\mathcal{B}_{\text{algo}} \setminus \mathcal{B}_F$ has minimum cardinal among the collections of balls in S whose union contains $S^{\ominus\varepsilon} \setminus \bigcup \mathcal{B}_F$. It follows that $\mathcal{B}_{\text{algo}}$ overshoots optimal $(0, \varepsilon)$ -ball approximations of S by at most $|\mathcal{B}_F| = |F|$ balls.

Although we made the simplifying assumption $\delta = 0$, note that it has close to no consequences for our method beside simplifying the notations. Indeed, so long as the medial

axis of $S^{\oplus\delta}$ is cycle-free, we can then substitute S by $S^{\oplus\delta}$ in the chapter above. Because the cycle-free assumption must be satisfied by $\text{MA}(S^{\oplus\delta})$ rather than $\text{MA}(S)$, working with $\delta > 0$ only restricts the input shapes S for which we reach an optimal.

In light of the previous observations, Algorithm 1 really only requires to loop over two lists: a list of pencils of a containing shape with a topological order, and a list of boundary elements of a shape to be covered. As long as the predicates are available, our method can thus be adapted to a wider class of shapes. We address this topic in the following Chapter 5.

Generalization to other shapes

Contents

5.1	Generalized problem and algorithm	150
5.2	Generalities on computing critical balls	153
5.3	Classic point-point pencil	155
5.3.1	Parameterization	155
5.3.2	Point constraint	156
5.3.3	Ball constraint	156
5.3.4	Arc constraint	162
5.4	Line-line pencil	164
5.4.1	Parameterization	164
5.4.2	Point constraint	165
5.4.3	Ball constraint	167
5.4.4	Arc constraint	169
5.5	Point-line pencil	170
5.5.1	Parameterization	170
5.5.2	Point constraint	171
5.5.3	Ball constraint	174
5.5.4	Arc constraint	177
5.6	Conclusion	180

In this chapter, we finally show how to implement predicates for computing critical balls. In fact, we shall see that these predicates can be implemented for a wider class of family of balls than just the elliptic 1-pencils that are needed for shapes defined as finite unions of disks. Therefore, we in fact extend the algorithm presented previously to a wider class of shapes, that will in particular cover polygonal shapes.

We start with some definitions in Section 5.1 in order to extend our approximation scheme and the notions related to critical balls to these new shapes. In Section 5.2 we explain the general strategy to establish the predicates needed for the algorithm, that we then apply to various families of medial balls in Sections 5.3 through 5.5.

5.1 Generalized problem and algorithm

§ Generalized approximation problem

Definition 20. Consider two subsets $I \subseteq O \subseteq \mathbb{R}^d$. A collection of balls \mathcal{B} is an (O, I) -ball approximation of O and I if $I \subseteq \bigcup \mathcal{B} \subseteq O$. We say that I is the *inner* subset, and O the *outer* subset of the (O, I) -ball approximation.

Note that I and O can be anything so long as the condition $I \subseteq O$ is satisfied. In particular, the subsets I, O can be a deformation of some original shape S , in which case $\bigcup \mathcal{B}$ is an approximation of S by union of balls. The specific properties of $\bigcup \mathcal{B}$ with respect to S will of course depend on the transformation used to derive I and O from S . The (δ, ε) -ball approximations we presented in the previous chapters are obviously a special case of (O, I) -ball approximation, where $I = S^{\ominus \varepsilon}$ and $O = S^{\oplus \delta}$. But the framework of (O, I) -ball approximations allows for a much wider range of transformations.

Remark 59. Similarly to Remark 45 for (δ, ε) -ball approximations, O and I admits (O, I) -ball approximations of finite cardinality whenever I is bounded and $d(I, O^c) > 0$.

Suppose I is bounded and $d(I, O^c) > 0$. Let $r = d(I, O^c)/2$ and notice that $\bar{I} \subseteq O$ is compact. To build a finite (O, I) -ball approximation, consider the collection of open balls with radius r centered at points in \bar{I} , these balls cover \bar{I} by construction. Since \bar{I} is compact, the collection of balls has a finite subcover. The closure of balls in that subcover is still contained within the outer subset O , they thus form a finite (O, I) -ball approximation.

§ Generalized algorithm Thereafter, we will show how our previous algorithm can be extended to the computation of (O, I) -ball approximations, given some assumptions on I and O . To properly state these assumptions, we require some additional notions. The medial axis of finite unions of balls in \mathbb{R}^2 can be fully described with a finite number of line segments, that corresponds to path-connected subsets of elliptic 1-pencils. In this context, because the centers' space of these elliptic 1-pencils is the bisector of the radical 0-sphere of the pencil, we refer to them as *point-point pencils*, or pp-pencils. We claim now, and justify later in Sections 5.4 and 5.5, that we can also define *line-line pencils* (ll-pencils) and *point-line pencils* (pl-pencils).

Remark 60. Outside of this chapter, a “pencil” will always refer to a pencil of balls, as defined in Chapter 2.

These two new pencils are families of inclusion-wise maximal balls, and their path-connected subsets can be used to describe portions of a medial axis. Intuitively, the centers' space of a ll-pencil is the bisector of two lines, whereas the centers' space of a pl-pencil is the bisector of a point and a line.

Conjecture 2 (Termination and optimality of the generalized algorithm). Let $I \subseteq O \subseteq \mathbb{R}^2$ such that $\partial I \cap \partial O = \emptyset$. Suppose that ∂I can be described as a finite collection

of points, segments and circular arcs. Suppose also that $\text{MA}(O)$ is cycle-free, and can be split into finitely many closed path-connected subsets of point-point, line-line, and point-line pencils. Given finite descriptions of ∂I and $\text{MA}(O)$, there is an algorithm that computes cardinal minimum (O, I) -ball approximations. Its time and space complexity is polynomial with respect to $|\partial I|$, $|\text{MA}(O)|$, and opt , where opt is the minimum cardinal of a (O, I) -ball approximation.

Remark 61. The algorithm is output-size sensitive.

Remark 62. In fact Algorithm 1 is also output-size sensitive, but we have shown in the proof of Theorem 9 that when O is a finite union of balls, we have $\text{opt} = \mathcal{O}(n)$.

Remark 63. With the assumptions of Conjecture 2, O and I always admit a finite (O, I) -ball approximation. Indeed, $\text{MA}(O)$ is closed, the two notions of medial and inclusion-wise maximal balls coincide for O . It follows that O is compact, and I is bounded. Finally the assumptions $\partial I \cap \partial O = \emptyset$ and $I \subseteq O$ imply that $d(I, O^c) > 0$, therefore Remark 59 applies.

Note that the class of shapes we allow for the outer subset O covers both finite unions of balls and polygons. In fact, it also covers some shapes whose boundaries are composed of both line segments and circular arcs (see Figure 5.1), but not all such shapes are admissible. Indeed, in order to apply our method to all these shapes, the algorithm would require to process a wider variety of pencils, such as “arc-point”, “arc-line”, and “arc-arc” pencils. On the contrary, the class of eligible inner shapes I is precisely those shapes whose boundary is composed of finitely many line segments and circular arcs.

Remark 64. Although the boundary of finite unions of balls is a collection of open circular arcs and vertices, describing their medial axis does not require “arc-arc” or “arc-point” pencils, as per Theorem 3.




Figure 5.1 – The medial axis of this shape can be split into three pencils: one pp-pencil (red), one ll-pencil (green) and one pl-pencil (blue). Its boundary is composed of both line segments and circular arcs.

§ **Mapping balls in O to medial balls of O** Recall from Section 4.2 that we have defined a projection function π that maps balls of a finite union of balls, and whose center lies in the interior, onto (inclusion-wise) maximal balls of the union. In general, this mapping is well-defined for any compact subset. Indeed, interior points that do not belong to the medial axis of the shape admit one unique closest point on the boundary, that we can use to define the projection. Consider an outer subset O that satisfies the conditions of Conjecture 2, and π_O its associated mapping to medial balls. The definition of π_O ensures that for any $b \subseteq O$, $c(b) \in \mathring{O}$, we have $b \subseteq \pi_O(b)$. We claim the following.

Conjecture 3. If O is an outer subset that satisfies the conditions of Conjecture 2, then the mapping π_O is continuous.

Recall that for finite union of balls, we have shown in Chapter 2 that they can be split into face joins, whose properties allowed us to conclude on the continuity of mapping π . For the current, broader class of shapes, we can still define the join of a boundary point. Indeed, let $u \in \partial O$ and $J_u \subseteq \text{MA}(O)$ the points of the medial axis that admit u as a closest point in ∂O . The join $\{u\} \star J_u$ is still well-defined, but the notion of join of a face is more ambiguous:

- two distinct points $u \neq v$ of what we would intuitively call a “face” of the boundary may have distinct collections $J_u \neq J_v$, and
- for a boundary face $f \subseteq \partial O$, if we let $J_f \subseteq \text{MA}(O)$ be the points of the medial axis that admit a point of f as a closest point in ∂O , then $f \star J_f$ is in general not a proper join.

Nevertheless, we claim that, up to a proper definition of a face of ∂O , the identity $\{u\} \star J_u = \text{vor}_O(u, \partial O)$ still holds, and that a decomposition of O similar to the face joins decomposition can be achieved through the Voronoi diagram of the faces of ∂O . Conjecture 3 would then follow from the properties of that decomposition.

§ **Specification and complexity** The cycle-free assumption in Conjecture 2 guarantees that we can define a partial T -order on the medial axis of the eligible outer subsets O . It follows that the notions of candidate and critical balls can be adapted. Assuming that Conjecture 3 is true, most of the results obtained in Chapter 4 would then transfer to (O, I) -ball approximations. In its current formulation, Property 4.13 would be incorrect when applied to critical balls of O , but we claim that an adapted reformulation is possible.

If there exist predicates for computing critical balls in every type of pencil of $\text{MA}(O)$, for every type of boundary element of ∂I , then Conjecture 2 could be proved from the various properties of Chapter 4. See Algorithm 2 for a high-level description of the resulting algorithm.

Recall that in the complexity proof of Theorem 9, $S = \bigcup \mathcal{S}$ was a finite union of balls, and we bounded the values of $|\text{MA}(S)|$, $|\partial S^{\ominus \varepsilon}|$, and opt with polynomial expressions of $n = |\mathcal{S}|$. In the context of (O, I) -ball approximations, O plays the role of S , I plays the role of $S^{\ominus \varepsilon}$. It follows that an adaptation of the previous complexity proof yields the polynomial complexity claim of Conjecture 2.

Algorithm 2 (O, I) -ball approximation

Input: Eligible outer shape O and inner shape I
Output: A cardinal minimum (O, I) -ball approximation \mathcal{B}

```

1:  $\mathcal{B} \leftarrow \emptyset$ 
2: constraints  $\leftarrow \partial I$ 
3: while  $I \neq \emptyset$  do
4:   Compute a critical ball  $b$  to  $I$ 
5:    $\mathcal{B} \leftarrow \mathcal{B} \cup \{b\}$ 
6:    $I \leftarrow I \setminus b$ 
7:   Update the set of constraints
8: end while
9: Return  $\mathcal{B}$ 

```

5.2 Generalities on computing critical balls

Given an inner shape I and an outer shape O , the core step of the algorithm is to compute the ball b critical for a set p , with $c(b) \in q$, where $p \subseteq \partial I$ and $q \subseteq \text{MA}(O)$. As a subset of $\text{MA}(O)$, q can be viewed as a family of medial balls. In the following we show how to perform this computation for a variety of simple geometric constraints p and families of medial balls q . More precisely, given a family of medial balls $\{b_\lambda\}$ and a set of points p , we show how to compute the values of λ for which $p \subseteq b_\lambda$. Knowing these values, the partial order on $\text{MA}(O)$ enables to conclude. Note that the parameterization of each family can always be chosen so that for any λ, μ with $b_\lambda, b_\mu \in q$, we have the equivalence $b_\lambda \leq_T b_\mu \iff \lambda \leq \mu$.

We consider the cases where the set of points $p \subseteq \partial I$ is either a single point, a line segment, a ball or a circular arc. Hence, any inner shape I whose boundary can be decomposed into line segments and circular arcs is handled.

Point For any point $x \in \mathbb{R}^2$, we use the equivalence $x \in b_\lambda \iff \text{pow}(x, b_\lambda) \leq 0$ to derive a first or second order polynomial in the variable λ , and then study its roots, bringing to light different configurations.

Line segment By convexity of balls and line segments, the balls b_λ containing a line segment $[xy]$ are the balls containing the points x and y .

Ball For any ball $b = b(c, r)$, we use the equivalence $b \subseteq b_\lambda \iff r_\lambda \geq \|c_\lambda - c\| + r$ to derive a second order polynomial in λ . The study of its roots enables to conclude for all the possible configurations.

Circular arc Let e be a circular arc supported by a ball b , $e \subseteq \partial b$. Let Λ_b be the ranges of values λ such that $\lambda \in \Lambda_b \iff b \subseteq b_\lambda$. There may be values $\lambda \notin \Lambda_b$ such that $e \subseteq b_\lambda$. Thus, we consider the values in $\mathbb{R} \setminus \Lambda_b$, and more precisely each connected component of $\mathbb{R} \setminus \Lambda_b$, that we call a sub-family of balls. The idea is to find a point $x \in \partial b$ such that no ball of the sub-family covers x . We will show in the following sections how to find such a

point. Given this point, any ball b_μ of the sub-family splits ∂b into two path-connected components, one contained in b_μ and the other outside of b_μ . Two cases arise, either $x \in e$ or $x \notin e$. If $x \in e$, then b_μ cannot contain e . Otherwise consider the endpoints of e . Any point of e that is not in b_μ has to be path-connected to x . In particular, that path necessarily goes through one endpoint of e . It follows that if both endpoints of e are actually covered by b_μ , then b_μ contains e entirely. See for instance Figure 5.2.

In certain configurations, there may not exist such a point $x \in \partial b$ that is never covered by the sub-family. In those cases, we establish that the subsets $b_\mu \cap \partial b$ are nested, and then inspect the T -extremal ball b_{μ^*} of the sub-family that contains both endpoints of e . We distinguish two cases. If that T -extremal ball b_{μ^*} does not contain e , then no ball of the sub-family does. Otherwise, b_{μ^*} contains e , and balls of the sub-family contain e if and only if they contain $b_{\mu^*} \cap \partial b$.




Figure 5.2 – Arc constraint in a line-line pencil : the red segment is the collection of centers corresponding to balls of the family that contain the red ball. Point x_1 is not covered by any ball of the pencil whose center is to the left of c_1 : since $x_1 \notin e$, any ball that includes the two endpoints of e includes e itself. Similarly, point x_2 is not covered by any ball of the pencil whose center is to the right of c_2 : since $x_2 \in e$, no ball of this sub-family contains e .

In the next sections, we consider each type of pencil one by one. First, we precise the parameterization used, and second we detail the computation for the point and ball constraints. Finally, for the circular arc constraints, we constructively prove the existence of the specific points mentioned previously, and deal with the exceptional cases.




Figure 5.3 – Point-point pencil.

5.3 Classic point-point pencil

5.3.1 Parameterization

Consider two distinct points u and v . We are interested in the family of balls whose bounding circle goes through both u and v . The collection of their centers is exactly the bisector L of segment $[uv]$. Let $b_0 = b(c_0, r_0)$ be the ball of the pencil with minimal radius. We have $c_0 = \frac{u+v}{2}$. Let \mathbf{e} be a unit direction vector for line L . For $\lambda \in \mathbb{R}$, we denote by c_λ the point of L such that $c_\lambda = c_0 + \lambda\mathbf{e}$. See Figure 5.3.

Notice by Property 2.6 that $r_\lambda^2 = r_0^2 + \|c_\lambda - c_0\|^2 = r_0^2 + \lambda^2$.

The remainder of this section is dedicated to the proof of the below result.

Property 5.1. *Under the Real-RAM model, there are predicates to compute critical balls in pp-pencils for point, segment, ball, and arc constraints in constant time.*

5.3.2 Point constraint

Recall by Lemma 2.8 that we have $\text{pow}(x, b_\lambda) = \text{pow}(x, b_0) - 2\langle x - c_0, c_\lambda - c_0 \rangle$, which yields

$$\text{pow}(x, b_\lambda) = \text{pow}(x, b_0) - 2\lambda\langle x - c_0, \mathbf{e} \rangle \quad (5.1)$$

For any point $x \in \mathbb{R}^2$, we thus have

$$x \in b_\lambda \iff \text{pow}(x, b_\lambda) \leq 0 \iff \text{pow}(x, b_0) - 2\lambda\langle x - c_0, \mathbf{e} \rangle \leq 0$$

We obtain a linear equation in λ . We distinguish three cases.

(a) $\langle x - c_0, \mathbf{e} \rangle = 0$ and $\text{pow}(x, b_0) > 0$. In other words, x belongs to the radical space of the pencil, and does not belong to the radical ball b_0 . Then $\text{pow}(x, b_\lambda) = \text{pow}(x, b_0) > 0$ and no ball of the pencil cover x . See point x in Figure 5.4.

(b) $\langle x - c_0, \mathbf{e} \rangle = 0$ and $\text{pow}(x, b_0) \leq 0$. Again $\text{pow}(x, b_\lambda) = \text{pow}(x, b_0)$. However this time every ball of the pencil covers x . See point y in Figure 5.4.

(c) $\langle x - c_0, \mathbf{e} \rangle \neq 0$. Point x does not lie on the radical space, and there is one unique value λ_* for which $\text{pow}(x, b_{\lambda_*})$ is null. Depending on the sign of $\langle x - c_0, \mathbf{e} \rangle$, the λ values such that b_λ contains x are either the half-line $\lambda \leq \lambda_*$ or the half-line $\lambda \geq \lambda_*$. See point z and the blue balls in Figure 5.4.

5.3.3 Ball constraint

Let $b = b(c, r)$ be a ball.

$$\begin{aligned} b \subseteq b_\lambda &\iff r_\lambda \geq \|c - c_\lambda\| + r \\ &\iff r_\lambda^2 \geq \|c - c_\lambda\|^2 + r^2 + 2r\|c - c_\lambda\| \\ &\iff -\text{pow}(c, b_\lambda) - r^2 \geq 2r\|c - c_\lambda\| \\ &\iff (-\text{pow}(c, b_\lambda) - r^2)^2 \geq 4r^2\|c - c_\lambda\|^2 \quad \& \quad -\text{pow}(c, b_\lambda) - r^2 \geq 0 \\ &\iff (2\lambda\langle c - c_0, \mathbf{e} \rangle - \text{pow}(c, b_0) - r^2)^2 \geq 4r^2\|c - c_0 - \lambda\mathbf{e}\|^2 \\ &\quad \& \quad \text{pow}(c, b_\lambda) + r^2 \leq 0 \\ &\iff 0 \geq A\lambda^2 + B\lambda + C \quad \& \quad \text{pow}(c, b_\lambda) - \text{pow}(c, b) \leq 0 \end{aligned}$$

where

$$\begin{aligned} A &= 4\left(r^2 - \langle c - c_0, \mathbf{e} \rangle^2\right) \\ B &= 4\langle c - c_0, \mathbf{e} \rangle(\text{pow}(c, b_0) - r^2) \\ C &= 4r^2\|c - c_0\|^2 - (\text{pow}(c, b_0) + r^2)^2 \\ &= (r_0 - \|c - c_0\| + r)(r_0 + \|c - c_0\| - r)(\|c - c_0\| + r - r_0)(\|c - c_0\| + r + r_0) \end{aligned}$$




Figure 5.4 – Point constraint in pp-pencils.

It follows that we can obtain the subset $\Lambda \subseteq \mathbb{R}$ of λ values for which the inclusion $b \subseteq b_\lambda$ is verified by comparing the roots of the second order polynomial $A\lambda^2 + B\lambda + C$ with the root of the first order polynomial $\text{pow}(c, b_\lambda) - \text{pow}(c, b)$.

The above two inequalities are sufficient to design a predicate and to find Λ . However, designing an appropriate predicate for circular arc constraints is dependent on the geometric configurations of the support ball of the arc, with respect to the pencil. Thus we now endeavor to provide a geometric intuition of the possible configurations. Ultimately, our analysis will yield the following (see Figure 5.5):

Property 5.2. *Let b a ball and consider an elliptic 1-pencil parameterized by λ , $\{b_\lambda\}_\lambda$. Let L^\perp the line through the radical sphere of this 1-pencil. The collection $\Lambda \subseteq \mathbb{R}$ such that $b \subseteq b_\lambda$ is:*

- *the empty set when $b \cap L^\perp \not\subseteq [uv]$ or b is tangent to L^\perp at u or v ,*
- *a closed half-line when $b \cap L^\perp = \emptyset$ or b is tangent to L^\perp at some point in the open segment $]uv[$,*
- *an interval with non-zero length when $b \cap L^\perp \subseteq]uv[$ and $\overset{\circ}{b} \cap L^\perp \neq \emptyset$,*
- *a singleton otherwise. Note that in this case, either u or v is on the boundary of b , and $b \cap L^\perp \subseteq [uv]$, with $\overset{\circ}{b} \cap L^\perp \neq \emptyset$.*

When both u and v belong to ∂b , the ball b belong to the pencil. In every other case, when $\Lambda \neq \emptyset$, the inclusion is strict, with tangency at the endpoint(s) of Λ .

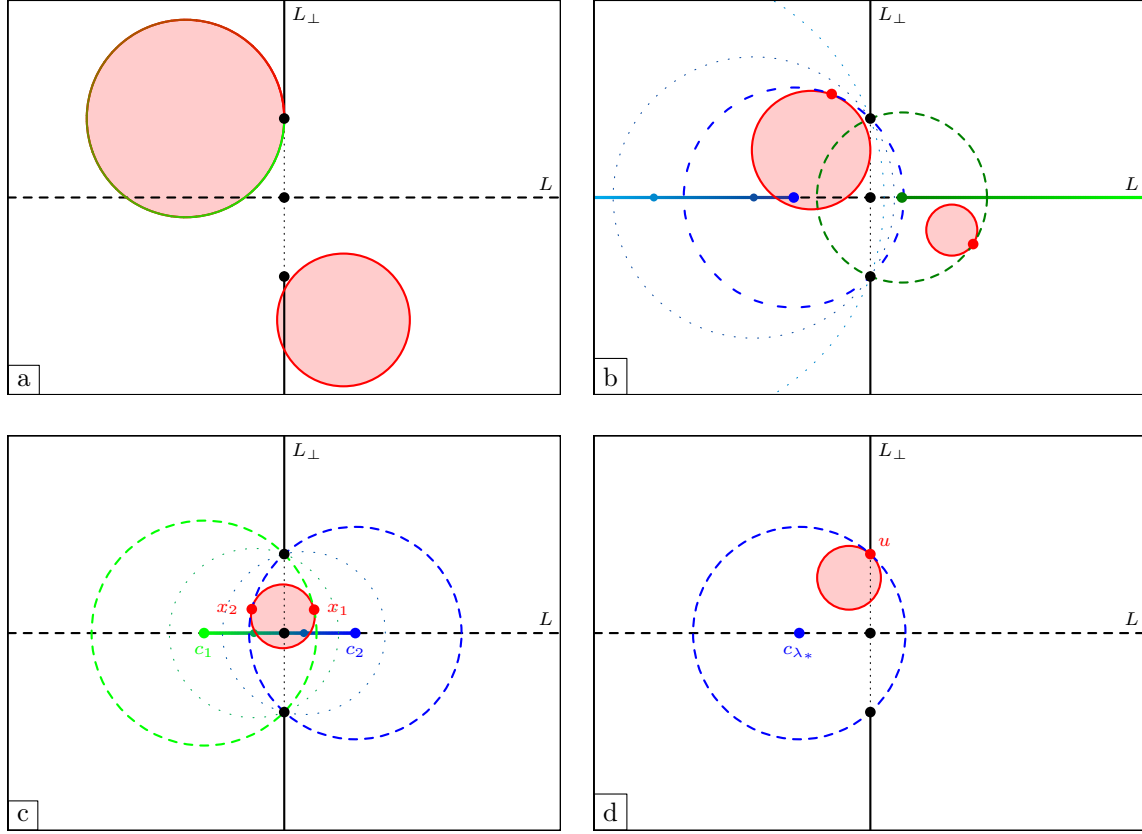


Figure 5.5 – Ball constraint in pp-pencils: (a) Λ_b is empty for the red balls; (b) Λ_b is a half-line; (c) Λ_b is a line segment; (d) Λ_b is a singleton.

To see this, we need an understanding of the geometric configurations that satisfy the inequality $A\lambda^2 + B\lambda + C \leq 0$ alone.

Lemma 5.3. *When λ satisfies the inequality $A\lambda^2 + b\lambda + C \leq 0$, we have one of three cases:*

- $r_\lambda \geq \|c - c_\lambda\| + r \iff b \subseteq b_\lambda$,
- $r \geq \|c - c_\lambda\| + r_\lambda \iff b_\lambda \subseteq b$,
- $\|c - c_\lambda\| \geq r_\lambda + r \iff \overset{\circ}{b} \cap \overset{\circ}{b}_\lambda = \emptyset$,

with equality if and only if b and b_λ are tangent.

Proof.

$$\begin{aligned} A\lambda^2 + B\lambda + C \leq 0 &\iff (-\text{pow}(c, b_\lambda) - r^2)^2 \geq 4r^2\|c - c_\lambda\|^2 \\ &\iff |\text{pow}(c, b_\lambda) + r^2| \geq 2r\|c - c_\lambda\| \end{aligned}$$

If $\text{pow}(c, b_\lambda) + r^2 \leq 0$ then we obtain

$$A\lambda^2 + B\lambda + C \leq 0 \implies r_\lambda \geq \|c - c_\lambda\| + r$$

as above. Otherwise we have $\text{pow}(c, b_\lambda) + r^2 \geq 0$ and

$$\begin{aligned} A\lambda^2 + B\lambda + C \leq 0 &\implies \text{pow}(c, b_\lambda) + r^2 \geq 2r\|c - c_\lambda\| \\ &\iff (\|c - c_\lambda\| - r)^2 \geq r_\lambda^2 \\ &\iff \|c - c_\lambda\| \geq r_\lambda + r \quad \text{or} \quad r \geq \|c - c_\lambda\| + r_\lambda \quad \square \end{aligned}$$

When we additionally have the inequality $\text{pow}(c, b_\lambda) + r^2 \leq 0$, we are in the first case $b \subseteq b_\lambda$ as shown previously. If however $\text{pow}(c, b_\lambda) + r^2 > 0$, we are in one of the two latter cases, which can further be decided depending on the sign of $\|c - c_\lambda\| - r$.

We now examine the geometric configurations in relation to the roots of the polynomial $A\lambda^2 + B\lambda + C$, and start by analyzing the degenerate cases. As a general remark, note that the geometric analysis below can also be done entirely algebraically by relying on a systematic sign analysis of $\text{pow}(c, b_\lambda) - \text{pow}(c, b)$.

§ $A = B = 0$ Recall that L is the mediator of segment $[uv]$. We denote by L^\perp the line (uv) . L and L^\perp are orthogonal and intersect at c_0 . In particular we have $L^\perp = \{x \in \mathbb{R}^2 \mid \langle x - c_0, \mathbf{e} \rangle = 0\}$. Since $\|\mathbf{e}\| = 1$, we obtain $|\langle c - c_0, \mathbf{e} \rangle| = d(c, L^\perp)$. Hence A is null if and only if $d(c, L^\perp) = r$, that is if and only if b is tangent to L^\perp . For B to be simultaneously zero, we must then have $\text{pow}(c, b_0) - r^2 = 0$. This simplifies to $d(c, L) = r_0$, that is the tangency point between b and L^\perp is either u or v . Intuitively, there are then no balls in the pencil that may contain b . Indeed, assume that $u \in \partial b$. Every ball b_λ goes through u , thus to contain b , the two balls b and b_λ must be tangent at u . Therefore b_λ is also tangent to L^\perp at u , but this is impossible because $[uv] \subseteq b_\lambda$, hence b_λ can never be tangent to L^\perp . Algebraically, $A = 0 = B$ implies

$$C = 4r^2(r^2 + r_0^2) - (r^2 + r_0^2)^2 = 4r^2r_0^2 > 0$$

which indicates that there is indeed no value of λ for which we have inclusion.

§ $A = 0, B \neq 0$ If $A = 0$ but $B \neq 0$, we obtain a linear equation that always has one root. Hence there are always some λ 's such that $B\lambda + C \leq 0$. We must determine whether the inclusion $b \subseteq b_\lambda$ is true for those values or not. We denote by y the tangency point between b and L^\perp . We distinguish two cases depending on the position of y on L^\perp . If $y \notin [uv]$, then the constraint ball b contains a point that is never covered by any b_λ , and we never have inclusion. If $y \in]uv[$, we argue that we always have the inclusion of interest. Indeed, for every λ we have $y \in b_\lambda$. Hence we cannot have $b_\lambda \subseteq b$ or $\dot{b} \cap \dot{b}_\lambda = \emptyset$. We deduce by Lemma 5.3 that we necessarily have $b \subseteq b_\lambda$ in every case. In practice, to distinguish between the two cases, it suffices to examine the second condition, that is the sign of $\text{pow}(c, b_\lambda) - \text{pow}(c, b)$, which is linear in λ as per Equation (5.1).

§ $A \neq 0$ We now address the case when $A \neq 0$. We start by computing the discriminant then inspect the geometric configurations of interest.

Lemma 5.4. *The discriminant is equal to $\Delta = 16r^2 \text{pow}(u, b) \text{pow}(v, b)$.*

Proof.

$$\begin{aligned}
\frac{\Delta}{4} &= (2\langle c - c_0, \mathbf{e} \rangle (\text{pow}(c, b_0) - r^2))^2 - 4(r^2 - \langle c - c_0, \mathbf{e} \rangle^2)C \\
\frac{\Delta}{16} &= \langle c - c_0, \mathbf{e} \rangle^2 (C + (\text{pow}(c, b_0) - r^2)^2) - r^2C \\
&= \langle c - c_0, \mathbf{e} \rangle^2 (4r^2\|c - c_0\|^2 - (\text{pow}(c, b_0) + r^2)^2 + (\text{pow}(c, b_0) - r^2)^2) - r^2C \\
&= \langle c - c_0, \mathbf{e} \rangle^2 (4r^2\|c - c_0\|^2 - 4r^2 \text{pow}(c, b_0)) - r^2C \\
&= \langle c - c_0, \mathbf{e} \rangle^2 (4r^2r_0^2) - r^2C \\
\frac{\Delta}{16r^2} &= 4r_0^2\langle c - c_0, \mathbf{e} \rangle^2 - 4r^2\|c - c_0\|^2 + (\text{pow}(c, b_0) + r^2)^2
\end{aligned}$$

We let \mathbf{u} be the unit vector from c_0 to u , it is a direction vector of line $L^\perp = (uv)$. Thus we have $\|c - c_0\|^2 = \langle c - c_0, \mathbf{e} \rangle^2 + \langle c - c_0, \mathbf{u} \rangle^2$. In addition, recall that $\|u - c_0\| = r_0 = \|v - c_0\|$, therefore we have $r_0\mathbf{u} = u - c_0$.

$$\begin{aligned}
\frac{\Delta}{16r^2} &= 4r_0^2 (\|c - c_0\|^2 - \langle c - c_0, \mathbf{u} \rangle^2) - 4r^2\|c - c_0\|^2 + (\|c - c_0\|^2 - r_0^2 + r^2)^2 \\
&= -4\langle c - c_0, r_0\mathbf{u} \rangle^2 + 4r_0^2\|c - c_0\|^2 - 4r^2\|c - c_0\|^2 + \|c - c_0\|^4 + r_0^4 + r^4 \\
&\quad - 2r_0^2\|c - c_0\|^2 + 2r^2\|c - c_0\|^2 - 2r^2r_0^2 \\
&= \|c - c_0\|^4 + 2r_0^2\|c - c_0\|^2 + r_0^4 - 4\langle c - c_0, u - c_0 \rangle^2 \\
&\quad + r^4 - 2r^2\|c - c_0\|^2 - 2r^2r_0^2 \\
&= (\|c - c_0\|^2 + r_0^2)^2 - 4\langle c - c_0, u - c_0 \rangle^2 + r^4 - r^2(2\|c - c_0\|^2 + 2r_0^2) \\
&= (\|c - c_0\|^2 + r_0^2 - 2\langle c - c_0, u - c_0 \rangle) (\|c - c_0\|^2 + r_0^2 + 2\langle c - c_0, u - c_0 \rangle) \\
&\quad + r^4 - r^2(2\|c - c_0\|^2 + 2r_0^2)
\end{aligned}$$

Note that $\|c - c_0\|^2 + r_0^2 - 2\langle c - c_0, u - c_0 \rangle = \|c - u\|^2$. Also, because $c_0 = \frac{u+v}{2}$, we have $u - c_0 = -(v - c_0)$. Hence we likewise have $\|c - c_0\|^2 + r_0^2 + 2\langle c - c_0, u - c_0 \rangle = \|c - v\|^2$. Finally notice that $\|c - u\|^2 + \|c - v\|^2 = 2\|c - c_0\|^2 + 2r_0^2 - \langle c - c_0, u - c_0 \rangle - \langle c - c_0, v - c_0 \rangle = 2\|c - c_0\|^2 + 2r_0^2$. We have

$$\begin{aligned}
\frac{\Delta}{16r^2} &= \|c - u\|^2\|c - v\|^2 + r^4 - r^2(\|c - u\|^2 + \|c - v\|^2) \\
&= (\|c - u\|^2 - r^2)(\|c - v\|^2 - r^2) \\
&= \text{pow}(u, b) \text{pow}(v, b) \quad \square
\end{aligned}$$

The sign of Δ can thus be deduced from the sign of $\text{pow}(u, b)$ and $\text{pow}(v, b)$. We distinguish four main cases.

(a) When $\text{pow}(u, b) < 0$ or $\text{pow}(v, b) < 0$, that is when u or v lies in the interior of b , we cannot have the inclusion $b \subseteq b_\lambda$. Indeed, for all λ , $b_\lambda \cap L^\perp =]uv]$, hence b necessarily meets the complement of b_λ . For the remaining cases, we can assume that both $\text{pow}(u, v)$ and $\text{pow}(v, b)$ are non-negative.

(b) If both $\text{pow}(u, b)$ and $\text{pow}(v, b)$ are null, then ∂b goes through both u and v and b actually belongs to the family. The only ball for which we have inclusion is b itself.

(c) If $\text{pow}(u, b) = 0$ but $\text{pow}(v, b) > 0$, then $\Delta = 0$ and the polynomial $A\lambda^2 + B\lambda + C$ has exactly one root. Notice that $|\langle c - c_0, \mathbf{e} \rangle| = d(c, L^\perp)$, as such because $b \cap L^\perp \neq \emptyset$ we have $r \geq |\langle c - c_0, \mathbf{e} \rangle|$ and $A \geq 0$. It follows that this unique root is the only λ value for which the polynomial $A\lambda^2 + B\lambda + C$ is non-positive. Also, because $v \in b_\lambda \setminus b$, we cannot have $b_\lambda \subseteq b$. Following Lemma 5.3, we must distinguish between the two cases $b \subseteq b_\lambda$ and $\overset{\circ}{b} \cap \overset{\circ}{b}_\lambda = \emptyset$. We claim that geometrically, they corresponds to whether b intersects the open segment $]uv[$ or not. In practice, this can be determined based on the sign of $\text{pow}(c, b_\lambda) - \text{pow}(c, b)$.

We know that b and b_λ must be tangent at u . Since b_λ cannot be tangent to L^\perp (the line through u and v), hence L^\perp cannot be the common tangent line to b and b_λ . Thus in the neighbourhood of u , the half-lines originating from u and supported by L^\perp are either inside or outside of b and b_λ . By definition of b_λ , the half-line from u through v is locally inside b_λ . It follows that if b does not meet $]uv[$, it must locally contain the opposite half-line and thus cannot be contained in b_λ . If however it intersects $]uv[$, then the interior of b and b_λ meet and they cannot be interior disjoint.

(d) If both $\text{pow}(u, b)$ and $\text{pow}(v, b)$ are positive, then $\Delta > 0$ and the polynomial $A\lambda^2 + B\lambda + C$ has two distinct roots $\lambda_1 < \lambda_2$. We have three more sub-cases.

- $A < 0$ and b does not meet the line L^\perp . The λ values for which the polynomial $A\lambda^2 + B\lambda + C$ is non-positive belong to $(-\infty, \lambda_1] \cup [\lambda_2, +\infty)$. Note that b does not contain u or v , therefore we cannot have $b_\lambda \subseteq b$. By Lemma 5.3 we must still distinguish between the two cases $b \subseteq b_\lambda$ and $\overset{\circ}{b} \cap \overset{\circ}{b}_\lambda = \emptyset$. Because we can assume $r > 0$ (otherwise we reduce to the point constraint), these two cases respectively correspond to having the inequality $r_\lambda - \|c - c_\lambda\| > 0$ or the inequality $r_\lambda - \|c - c_\lambda\| < 0$. Since the function $\lambda \mapsto r_\lambda - \|c - c_\lambda\|$ is continuous with respect to λ , its zero(s) must belong to the open interval (λ_1, λ_2) . It follows that this function has constant sign over $(-\infty, \lambda_1]$ and $[\lambda_2, +\infty)$.

As b is entirely contained in one of the two open half-spaces delimited by L^\perp , we know that exactly one of $b_{\lambda=-\infty}$ or $b_{\lambda=+\infty}$ contains b , while the other ball is disjoint from b . Thus, the collection Λ of values such that $b \subseteq b_\lambda$ is precisely one of the two intervals $(-\infty, \lambda_1]$ or $[\lambda_2, +\infty)$. The correct interval can be determined with the use of the second inequality $\text{pow}(c, b_\lambda) - \text{pow}(c, b)$.

For the remaining sub-cases we can assume $A > 0$, which implies that $b \cap L^\perp \neq \emptyset$. Because neither u nor v belong to b , and b is convex, either $b \cap L^\perp \subseteq]uv[$ or $b \cap L^\perp \cap]uv[= b \cap]uv[= \emptyset$. Note also that the only λ values that may satisfy $b \subseteq b_\lambda$ lie in the interval $[\lambda_1, \lambda_2]$.

- $b \cap L^\perp \subseteq]uv[$. Because $b \cap L^\perp \neq \emptyset$ and $]uv[\subseteq \overset{\circ}{b}_\lambda$ for all λ , we know that b and b_λ cannot be interior disjoint. In addition, we cannot have $b_\lambda \subseteq b$ because b does not contain u or v . By Lemma 5.3 we have $b \subseteq b_\lambda$ for all $\lambda \in [\lambda_1, \lambda_2]$.
- $b \cap [uv] = \emptyset$. Because $b \cap L^\perp$ is non-empty, b contains point(s) of L^\perp that no b_λ covers, and there are no value λ such that $b \subseteq b_\lambda$.

5.3.4 Arc constraint

We now address arc constraints. Consider a circular arc e and let b be its supporting ball, $e \subseteq \partial b$. Because balls are convex, the case $e = \partial b$ reduces to a ball constraint. We thus assume without loss of generality that $e \subsetneq \partial b$ has two distinct endpoints. Let $\Lambda_b \subseteq \mathbb{R}$ the collection of λ values such that $b \subseteq b_\lambda$. We distinguish two main cases depending on whether Λ_b is empty.

1. Suppose that $\Lambda_b \neq \emptyset$, and consider $\lambda_* \in \partial\Lambda_b$ an endpoint of Λ_b . Such an endpoint always exist by Property 5.2 since Λ cannot be the entire real line. This property also implies that Λ_b is convex. By symmetry, suppose without loss of generality that for all $\mu < \lambda_*$, $b \not\subseteq b_\mu$. We investigate the values of μ such that $e \subseteq b_\mu$. By Property 5.2, we know that b and b_{λ_*} are tangent with $b \subseteq b_{\lambda_*}$. From the previous section, we distinguish three sub-cases.

- (a) $b = b_{\lambda_*}$, and b belongs to the pencil. Using properties of elliptic pencils, it is easy to find a point of ∂b that does not belong to any b_μ , $\mu < \lambda_*$. For instance we can take the point $x = c + re$, it is one of the point of intersection between ∂b and line L . We can then apply the method outlined in Section 5.2 using this point x , see Figure 5.6.a.

For the remaining cases, we can assume that for all $\lambda \in \Lambda_b$, we have $b \subsetneq b_\lambda$. In particular, b and b_{λ_*} have one unique tangency point that we call x . The two remaining sub-cases depend on the nature of x .

- (b) $x \notin \{u, v\}$, then for $\mu < \lambda_*$ we necessarily have $x \notin b_\mu$. We can thus apply the method outlined in Section 5.2 using this tangency point x . See Figures 5.5.b and 5.5.c.
- (c) $x \in \{u, v\}$, and λ_* is in fact the unique value in Λ_b . Without loss of generality, suppose $x = u$. Because the power of any point y to a ball b_λ is linear with respect to λ , we can show that for any $\mu < \lambda < \lambda_*$, we have $b \cap b_\mu \subseteq b \cap b_\lambda$. Indeed, if we inspect the power diagram of two balls b_μ and b_λ , $\mu < \lambda$, then the power cell of b_μ is the half-space $H_- = \{y \in \mathbb{R}^2 \mid \langle x - c_0, \mathbf{e} \rangle \leq 0\}$, whereas the power cell of b_λ is the half-space $H_+ = \{y \in \mathbb{R}^2 \mid \langle x - c_0, \mathbf{e} \rangle \geq 0\}$. In particular, the portion of e covered by b_μ increases as μ goes to λ_* . We have three more cases depending on the position of u with respect to the endpoints of e , see Figure 5.6.b.
- $u \in \overset{\circ}{e}$, then no ball b_μ , $\mu < \lambda_*$, covers e .
 - $u \notin \bar{e}$, then b_μ contains e if and only if b_μ contains both endpoint of e .
 - $u \in \partial e$ is an endpoint of e . Let μ_* be the smallest parameter value such that b_{μ_*} contains both endpoints of arc e . For this particular μ_* value, ∂b_{μ_*} splits ∂b into

two circular arcs, and one of these two arcs is e itself. If e is the arc inside b_{μ_*} , then $e \subseteq b_\mu$ for all $\mu_* \leq \mu < \lambda_*$. Otherwise if e is the arc outside b_{μ_*} , then no ball b_μ , $\mu < \lambda_*$, contains e .

2. Suppose now that $\Lambda_b = \emptyset$. By Property 5.2, we have two sub-cases.

- (a) b intersects $L^\perp \setminus [uv]$. Note that necessarily, $(\partial b \cap L^\perp) \setminus [uv]$ is non-empty. We can take any point x in that intersection.
- (b) b is tangent to L^\perp at either u or v . Without loss of generality, assume the point of tangency is u . Because every ball b_λ contains the segment $[uv]$, b and b_λ are never tangent. It follows that $b_\lambda \cap \partial b$ is an arc of non-zero length that admits u as an endpoint. We can conclude as in the previous case 1.(c). Properties of the family guarantee that we still have $b \cap b_\mu \subseteq b \cap b_\lambda$ for $\mu \leq \lambda < \lambda_*$. Let μ_* the smallest parameter value such that b_{μ_*} contains both endpoint of e . Either b_{μ_*} contains e , or it doesn't. In the former case, every ball b_μ with $\mu_* \leq \mu < \lambda_*$ will contain e . In the latter, no ball b_μ with $\mu < \lambda_*$ contains e .

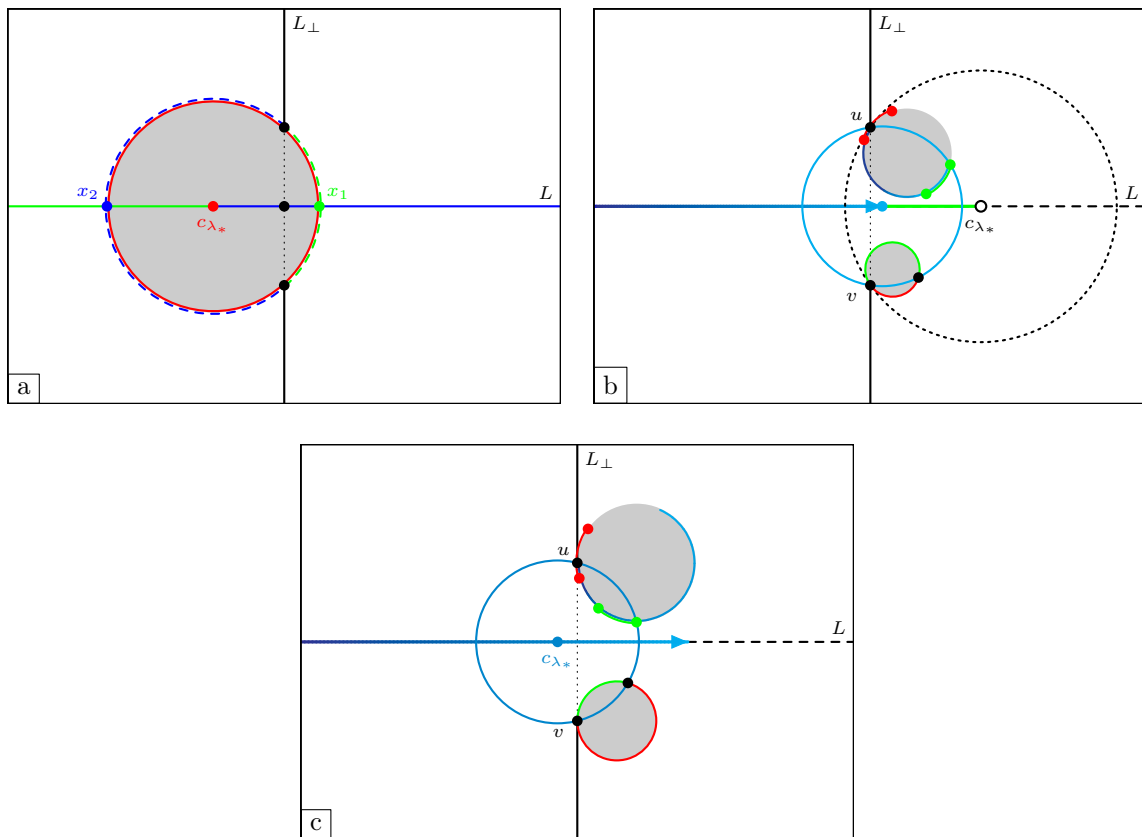


Figure 5.6 – Arc constraint in pp-pencils: (a) the blue dashed arc is never covered by balls centered on the blue half-line; (b) and (c) the red arcs cannot be covered by any ball b_μ , $\mu < \lambda_*$, but the green arcs are covered by any ball of the family centered on the green segment.

In every possible case, we can find every λ value such that $e \subseteq b_\lambda$ by using previous predicates a constant number of time.

5.4 Line-line pencil




Figure 5.7 – Line-line pencil.

5.4.1 Parameterization

The locus of points equidistant from two segments contains a line segment portion that coincides with the angle bisector of the lines supporting the segments. As such, it is of interest to study the family of balls centered on such a bisector.

Consider two non-parallel lines L_+ and L_- that intersect at point c_0 , and L one of their bisector. By definition, any point of L is equidistant to both L_+ and L_- . We are interested in the balls centered at points of bisector L , and whose radius is equal to the distance from their center to any of the two lines. Such balls are tangent to both lines, there exists a linear relationship between their center's positions and radii as per Equation (5.2). The union over all balls of this family yields the double cone delimited by L_+ and L_- that contains the bisector L . In fact, this is the family of maximal balls for this cone, and L is its skeleton.

In practice we will always deal with segment subsets of the bisector whose interior does not contain c_0 , we can thus reduce from the complete bisector L to a half-line originating from c_0 . Let \mathbf{e} be a unit direction vector of L . For $\lambda \in \mathbb{R}$, we denote by $c_\lambda = c_0 + \lambda \mathbf{e}$ the point of L at distance λ from c_0 , b_λ the ball of the family centered at c_λ and r_λ its radius. Let α be

the measure of the angle between L_+ and L_- bisected by L . We have

$$r_\lambda = |\lambda| \sin \frac{\alpha}{2} \quad (5.2)$$

Refer to Figure 5.7 for an illustration of various features of ll-pencils. The remainder of this section is dedicated to the proof of the below result.

Property 5.5. *Under the Real-RAM model, there are predicates to compute critical balls in ll-pencils for point, segment, ball, and arc constraints in constant time.*

5.4.2 Point constraint

Let $x \in \mathbb{R}^2$. Then we have

$$\begin{aligned} x \in b_\lambda &\iff \text{pow}(x, b_\lambda) \leq 0 \\ &\iff \|x - c_\lambda\|^2 - r_\lambda^2 \leq 0 \\ &\iff \|x - c_0\|^2 - 2\langle x - c_0, c_\lambda - c_0 \rangle + \|c_\lambda - c_0\|^2 - r_\lambda^2 \leq 0 \\ &\iff A\lambda^2 + B\lambda + C \leq 0 \end{aligned} \quad (5.3)$$

where

$$\begin{aligned} A &= \cos^2 \frac{\alpha}{2} \\ B &= -2\langle x - c_0, \mathbf{e} \rangle \\ C &= \|x - c_0\|^2 \end{aligned}$$

Computing the λ values such that $x \in b$ can thus be reduced to computing the roots of a quadratic function.

Notice that $A > 0$, hence the polynomial never degenerates to a lower degree, and the only λ values that satisfy the inequality lie between the roots of the polynomial. Recall that we can always reduce to the half-line of non-negative values of λ , thus we should discard the negative solutions to the inequality.

Remark 65. Because $A \times C \geq 0$, note that when the polynomial has two distinct roots λ_1 and λ_2 , they have the same sign. Restricting to non-negative values thus imply discarding every solutions, or keeping every solutions.

We claim the following, see Figure 5.8.

Property 5.6. *Consider the polynomial $A\lambda^2 + B\lambda + C$ for a point constraint x in a ll-pencil. Its number of roots is tied to the position of x with respect to the double cone as follows:*

- *0 root: x lies outside of the double cone and no ball of the family contains x .*
- *1 root: x lies on the boundary of the double cone, exactly one ball of the family contains x .*

- 2 roots: x lies in the interior of the double cone, every ball with λ parameter in between these roots contain x .




Figure 5.8 – Point constraint in ll-pencils.

To see this, it suffices to compute the discriminant. Let \mathbf{e}_\perp be a unit vector orthogonal to \mathbf{e} . In particular, we have $\|x - c_0\|^2 = \langle x - c_0, \mathbf{e} \rangle^2 + \langle x - c_0, \mathbf{e}_\perp \rangle^2$.

$$\begin{aligned} \frac{\Delta}{4} &= \langle x - c_0, \mathbf{e} \rangle^2 - \cos^2 \frac{\alpha}{2} \|x - c_0\|^2 \\ &= \langle x - c_0, \mathbf{e} \rangle^2 - \cos^2 \frac{\alpha}{2} \left(\langle x - c_0, \mathbf{e} \rangle^2 + \langle x - c_0, \mathbf{e}_\perp \rangle^2 \right) \\ &= \sin^2 \frac{\alpha}{2} \langle x - c_0, \mathbf{e} \rangle^2 - \cos^2 \frac{\alpha}{2} \langle x - c_0, \mathbf{e}_\perp \rangle^2 \\ &= \left(\sin \frac{\alpha}{2} \langle x - c_0, \mathbf{e} \rangle - \cos \frac{\alpha}{2} \langle x - c_0, \mathbf{e}_\perp \rangle \right) \left(\sin \frac{\alpha}{2} \langle x - c_0, \mathbf{e} \rangle + \cos \frac{\alpha}{2} \langle x - c_0, \mathbf{e}_\perp \rangle \right) \end{aligned}$$

Let

$$\mathbf{n}_+ = \cos \left(\frac{\alpha - \pi}{2} \right) \mathbf{e} + \sin \left(\frac{\alpha - \pi}{2} \right) \mathbf{e}_\perp = \sin \frac{\alpha}{2} \mathbf{e} - \cos \frac{\alpha}{2} \mathbf{e}_\perp$$

and likewise

$$\mathbf{n}_- = \cos \left(\frac{\pi - \alpha}{2} \right) \mathbf{e} + \sin \left(\frac{\pi - \alpha}{2} \right) \mathbf{e}_\perp = \sin \frac{\alpha}{2} \mathbf{e} + \cos \frac{\alpha}{2} \mathbf{e}_\perp$$

Note that \mathbf{n}_+ and \mathbf{n}_- are unit normal vectors to lines L_+ and L_- , and we have

$$\Delta = 4 \langle x - c_0, \mathbf{n}_+ \rangle \langle x - c_0, \mathbf{n}_- \rangle$$

In addition, let H_+ be the closed half-space $\{x \in \mathbb{R}^2 \mid \langle x - c_0, \mathbf{n}_+ \rangle \geq 0\}$ and likewise $H_- = \{x \in \mathbb{R}^2 \mid \langle x - c_0, \mathbf{n}_- \rangle \geq 0\}$. Then the cone containing the half-line $c_0 + \mathbb{R}^+ \mathbf{e}$ is precisely the intersection between H_+ and H_- , whereas the cone containing the symmetric half-line $c_0 + \mathbb{R}^- \mathbf{e}$ is $\overline{H_+^c} \cap \overline{H_-^c}$. The double cone delimited by L_+ and L_- that contains the line L is the union of $H_+ \cap H_-$ and $\overline{H_+^c} \cap \overline{H_-^c}$. We can thus discuss the sign of Δ depending on the position of x with respect to these two cones:

- Δ is null when x lies on either L_+ or L_- ,
- Δ is positive when x lies in the interior of one of these two cones,
- Δ is negative when x is outside both cones.

5.4.3 Ball constraint

Let $b = b(c, r)$ be a ball. Then we have

$$\begin{aligned}
 b \subseteq b_\lambda &\iff r_\lambda \geq \|c_\lambda - c\| + r \\
 &\iff r_\lambda - r \geq \|c_\lambda - c\| \\
 &\iff (r_\lambda - r)^2 \geq \|c_\lambda - c\|^2 \quad \& \quad r_\lambda - r \geq 0 \\
 &\iff 0 \geq A\lambda^2 + B\lambda + C \quad \& \quad r_\lambda - r \geq 0
 \end{aligned}$$

where

$$\begin{aligned}
 A &= \cos^2 \frac{\alpha}{2} \\
 B &= -2\langle c - c_0, \mathbf{e} \rangle + 2r \sin \frac{\alpha}{2} \\
 C &= \|c_0 - c\|^2 - r^2 = \text{pow}(c_0, b)
 \end{aligned}$$

The two inequalities above suffice to design a predicate. Note that if we consider singletons to be balls of radius zero, the above coefficients coincide with the previous section, as opposed to the classic pencil of balls. Once again $A > 0$, thus we only have to investigate the λ values in between the roots to determine the balls b_λ that contain b .

We will again need some understanding of the possible geometric configurations for a ball constraint in order to derive the predicate for arc constraints. We thus study these geometric configurations. We eventually obtain the below property, see Figure 5.9.

Property 5.7. *Let b a ball and consider a half ll-pencil parameterized by $\lambda \in \mathbb{R}^+$, $\{b_\lambda\}_\lambda$. Let H_+ and H_- the two half-spaces such that $H_+ \cap H_- = \bigcup_{\lambda \geq 0} b_\lambda$, and L_+ and L_- their bounding lines. The collection $\Lambda \subseteq \mathbb{R}^+$ such that $b \subseteq b_\lambda$ is:*

- the empty set when b is not entirely in the cone $H_+ \cap H_-$,
- a singleton when $b \subseteq H_+ \cap H_-$ and b is tangent to L_+ and/or L_- ,
- a closed interval when b is in the interior of $H_+ \cap H_-$.

When the constraint ball b is tangent to both L_+ and L_- , it belongs to the ll-pencil. In every other case, when $\Lambda \neq \emptyset$, the inclusion is strict, with tangency at the endpoint(s) of Λ .

To prove this property, we will rely on this next lemma.




Figure 5.9 – Ball constraint in l -pencils: (a) Λ_b is empty for the red ball; (b) Λ_b is a singleton; (c) Λ_b is a line segment.

Lemma 5.8. *Let $b_\lambda = b(c_\lambda, r_\lambda)$ and $b = b(c, r)$ be two balls. The inequality $(r_\lambda - r)^2 \geq \|c_\lambda - c\|^2$ holds if and only if we have either $b_\lambda \subseteq b$ or $b \subseteq b_\lambda$, with equality when the balls are tangent.*

The additional inequality $r_\lambda - r \geq 0$ enables us to distinguish between the two possibilities

of the lemma above. We inspect now the discriminant:

$$\begin{aligned}
\frac{\Delta}{4} &= \left(-\langle c - c_0, \mathbf{e} \rangle + r \sin \frac{\alpha}{2} \right)^2 - \cos^2 \frac{\alpha}{2} \left(\|c - c_0\|^2 - r^2 \right) \\
&= \langle c - c_0, \mathbf{e} \rangle^2 - 2r \sin \frac{\alpha}{2} \langle c - c_0, \mathbf{e} \rangle + r^2 - \cos^2 \frac{\alpha}{2} \left(\langle c - c_0, \mathbf{e} \rangle^2 + \langle c - c_0, \mathbf{e}_\perp \rangle^2 \right) \\
&= \sin^2 \frac{\alpha}{2} \langle c - c_0, \mathbf{e} \rangle^2 - \cos^2 \frac{\alpha}{2} \langle c - c_0, \mathbf{e}_\perp \rangle^2 - 2r \sin \frac{\alpha}{2} \langle c - c_0, \mathbf{e} \rangle + r^2 \\
&= \left(r - \sin \frac{\alpha}{2} \langle c - c_0, \mathbf{e} \rangle \right)^2 - \cos^2 \frac{\alpha}{2} \langle c - c_0, \mathbf{e}_\perp \rangle^2 \\
&= \left(r - \sin \frac{\alpha}{2} \langle c - c_0, \mathbf{e} \rangle - \cos \frac{\alpha}{2} \langle c - c_0, \mathbf{e}_\perp \rangle \right) \times \left(r - \sin \frac{\alpha}{2} \langle c - c_0, \mathbf{e} \rangle + \cos \frac{\alpha}{2} \langle c - c_0, \mathbf{e}_\perp \rangle \right) \\
&= (r - \langle c - c_0, \mathbf{n}_- \rangle) \times (r - \langle c - c_0, \mathbf{n}_+ \rangle)
\end{aligned}$$

Note that $|\langle c - c_0, \mathbf{n}_+ \rangle| = d(c, L_+)$, therefore the factor $r - \langle c - c_0, \mathbf{n}_+ \rangle$ is negative when b is in the interior of H_+ , zero when b is in H_+ and tangent to L_+ , positive otherwise. The sign of the other factor has a similar interpretation with respect to L_- and H_- . We have four cases:

(a) Whenever one factor is positive, b intersects the complement of the family's domain and we never have the desired inclusion.

For the remaining cases, we can assume that both factors are non-positive, thus b is entirely within the cone covered by the half-family. From Lemma 5.8 we deduce that whenever λ satisfies the inequality, we necessarily have $b \subseteq b_\lambda$. Indeed, the b_λ 's are medial for the cone, thus they are inclusion-wise maximal.

(b) If both factor are null, $\Delta = 0$, and b is tangent to both L_+ and L_- . Thus c belongs to the bisector of these two lines and b is actually a ball of the family. The only solution to the quadratic inequality therefore corresponds to b itself.

(c) If one factor is null and the other negative, we have $\Delta = 0$. The unique λ_* solution corresponds to the unique b_{λ_*} that covers the tangency point between b and one of L_+ or L_- .

(d) If both factors are negative, then for every λ value between the roots we have $b \subseteq b_\lambda$.

5.4.4 Arc constraint

We now address arc constraints. Let e be a circular arc supported by ball b , and Λ_b the collection of λ values for which we have the inclusion $b \subseteq b_\lambda$. Recall that we can in general derive the inclusion λ values for e from those of the endpoints of e , and Λ_b . To do so it suffices to find a point in ∂b that is never covered by any ball of the sub-families that do not contain b .

By relying on Property 5.7, we distinguish three cases:

(a) $\Lambda_b = \emptyset$ and b is not contained in the cone. Thus there exists $x \in b$ not covered by any b_λ . In particular, either $x_+ = c - r\mathbf{n}_+$ or $x_- = c - r\mathbf{u}_-$ is never covered by any b_λ but is still in ∂b . See Figure 5.9.a.

(b) $\Lambda_b = \{\lambda_*\}$ is a singleton and b is tangent to either L_+ or L_- . Suppose b is tangent to L_+ . The tangency point between b and L_+ is uniquely covered by b in the family. It can thus be used to derive the inclusion λ values for arc e . See Figure 5.9.b.

(c) $\Lambda_b = [\lambda_1, \lambda_2]$ with $\lambda_1 < \lambda_2$. We know by Property 5.7 that b and b_{λ_2} are tangent with $b \subsetneq b_{\lambda_2}$. Let x_2 be the unique point of tangency between b and b_{λ_2} . We investigate the sub-family of balls $\{b_\mu \mid \mu > \lambda_2\}$. We argue that for $\mu > \lambda_2$, we have $x_2 \notin b_\mu$. Indeed by Equation (5.3) we know that $\text{pow}(x_2, b_\lambda)$ is quadratic with respect to λ . Because $x_2 \in \partial b_{\lambda_2}$, and $x_2 \in b_\lambda$ for $\lambda \in [\lambda_1, \lambda_2]$, it follows that $\text{pow}(x_2, b_\mu) > 0$ for all $\mu > \lambda_2$. We obtain a similar result for the other endpoint λ_1 and the tangency point between b and b_{λ_1} . See Figure 5.9.c.

Again, the predicate for arc constraints can be derived from the predicates for point and ball constraints.

5.5 Point-line pencil




Figure 5.10 – Point-line pencil.

5.5.1 Parameterization

The locus of points equidistant from a point and a line segment contains a portion that coincides with the bisector of that point and the support line of the segment. Because this locus occurs in the medial axis of polygons, we study the corresponding family of medial balls.

Let \mathcal{D} be a line and F a point not on \mathcal{D} . The bisector of F and \mathcal{D} is the parabola with focus F and directrix \mathcal{D} . We study the family of balls whose center lies on that parabola and whose bounding circle goes through F . Because the center of these balls are equidistant from F and \mathcal{D} , they are tangent to \mathcal{D} . The union over all balls of that family exactly covers the closed half-space delimited by \mathcal{D} and that contains F .

Remark 66. In order for a path-connected subset of the family to form a valid collection of medial balls, and thus to be used to describe the medial axis of a shape, the focus F must remain in the boundary of the union.

To simplify the discussion, we choose to work in the Cartesian coordinate system such that the directrix coincides with the horizontal axis, and the focus lies above it on the vertical axis. The focus then has coordinates $F(0, 2f)$ and the parabola is the collection of points

$$\mathcal{D} = \left\{ (x, y) \mid y = \left(\frac{x}{2\sqrt{f}} \right)^2 + f \right\}$$

where f denotes the focal length of the parabola and is a strictly positive parameter.

We parameterize the centers by their abscissa in the above coordinate system. For $\lambda \in \mathbb{R}$, we denote by c_λ the point of the parabola with abscissa λ . As usual, we then denote the ball of the family centered at c_λ by b_λ , and its radius by r_λ . A basic property of the family is that $r_\lambda = \|c_\lambda - F\|$ is equal to the ordinate of c_λ . Hence it has coordinates (λ, r_λ) with

$$r_\lambda = \frac{\lambda^2}{4f} + f \quad (5.4)$$

Refer to Figure 5.10 for an illustration of various features of pl-pencils. The remainder of this section is dedicated to the proof of the below result.

Property 5.9. *Under the Real-RAM model, there are predicates to compute critical balls in pl-pencils for point, segment, ball, and arc constraints in constant time.*

5.5.2 Point constraint

We can express the power of a point $M(x_M, y_M) \in \mathbb{R}^2$ to b_λ with the following:

$$\text{pow}(M, b_\lambda) = \|M - c_\lambda\|^2 - r_\lambda^2 = \|M - c_\lambda\|^2 - \|F - c_\lambda\|^2 = 2\langle c_\lambda, F - M \rangle + \|M\|^2 - \|F\|^2$$

Since the ordinate of c_λ is quadratic in λ as per Equation (5.4), the power of a point to a ball of the family is also quadratic in λ . Hence

$$\begin{aligned} M \in b_\lambda &\iff \text{pow}(M, b_\lambda) \leq 0 \\ &\iff 2(-\lambda x_M + r_\lambda(2f - y_M)) + \|M\|^2 - 4f^2 \leq 0 \\ &\iff A\lambda^2 + B\lambda + C \leq 0 \end{aligned} \quad (5.5)$$

where

$$\begin{aligned} A &= 1 - \frac{y_M}{2f} \\ B &= -2x_M \\ C &= \|M\|^2 - 2fy_M \end{aligned}$$

Computing the λ values for which $x \in b_\lambda$ thus reduces to finding the sign of a quadratic function. We have the following, see Figure 5.11.

Property 5.10. *Consider the inequality $A\lambda^2 + B\lambda + C \leq 0$ for a point constraint M in a pl-pencil. Let L be the horizontal line through the focus F . The collection $\Lambda \subseteq \mathbb{R}$ such that $M \in b_\lambda$ is:*

- the entire real line \mathbb{R} when $M = F$,
- the union of two disjoint, closed half-lines when M is (strictly) above L ,
- a closed half-line when M is in $L \setminus \{F\}$,
- a closed, bounded interval when M is (strictly) below L but (strictly) above the horizontal axis,
- a singleton when M is on the horizontal axis,
- the empty set when M is (strictly) below the horizontal axis.

To see this, we examine the geometric configurations in relation with the roots of the polynomial, starting with the degenerate cases. Coefficient A is null if and only if M lies on the horizontal line through F , while coefficient B is null if and only if M lies on the ordinate axis.

§ $A = B = 0$ We have $M = F$. It follows that every ball of the family actually contains M .

§ $A = 0, B \neq 0$ M lies on the horizontal line through F , but is distinct from F . We have $y_M = 2f$ and $x_M \neq 0$, thus $C = x_M^2 > 0$. When $x_M > 0$, we then have the inequality

$$-2x_M\lambda + x_M^2 \leq 0 \iff \lambda \geq \frac{x_M^2}{2x_M} = \frac{x_M}{2}$$

In the symmetric case $x_M < 0$, we obtain

$$-2x_M\lambda + x_M^2 \leq 0 \iff \lambda \leq \frac{x_M^2}{2x_M} = \frac{x_M}{2}$$

We can group both case in the following way

$$2\lambda \operatorname{sign}(x_M) \geq |x_M|$$




Figure 5.11 – Point constraint in pl-pencils (a) Λ is the union of two half-lines; (b) Λ is a half-line; (c) Λ is a singleton for points on \mathcal{D} , a segment for points between \mathcal{D} and L .

§ $A \neq 0$ M is not on the horizontal line through F and $y_M \neq 2f$. The polynomial may have 0 to 2 roots. We inspect its discriminant:

$$\begin{aligned}
 \frac{\Delta}{4} &= x_M^2 - \left(1 - \frac{y_M}{2f}\right) (\|M\|^2 - 2fy_M) \\
 &= -2y_M^2 + 2fy_M + \frac{y_M}{2f} \|M\|^2 \\
 &= \frac{y_M}{2f} (-4fy_M + 4f^2 + x_M^2 + y_M^2) \\
 &= \frac{y_M}{2f} (x_M^2 + (y_M - 2f)^2) \\
 &= \frac{y_M}{2f} \|M - F\|^2
 \end{aligned}$$

By the assumption $A \neq 0$, we cannot have $M = F$, thus $\text{sign}(\Delta) = \text{sign}(y_M)$. We distinguish four cases.

(a) $y_M < 0$, which implies $\Delta < 0$ and $A > 0$. M lies outside the half-space spanned by the pencil family and there are no ball that covers the point.

(b) $y_M = 0$, which implies $\Delta = 0$ and $A > 0$. M lies on the directrix and there is exactly one ball of the family that covers it.

(c) $0 < y_M < 2f$, which implies $\Delta > 0$ and $A > 0$. The quadratic function has two roots, and the λ values between those roots correspond to the balls of the family that cover M .

(d) $y_M > 2f$, which implies $\Delta > 0$ and $A < 0$. The quadratic function has two roots, and the λ values between those roots correspond to the balls of the family that do not cover M .

5.5.3 Ball constraint

Let $b = b(c, r)$ be a ball where c has coordinates $c(x_c, y_c)$. Using the expression of the power of a point to b_λ in Equation (5.5) we obtain the below development:

$$\begin{aligned}
 b \subseteq b_\lambda &\iff r_\lambda \geq \|c - c_\lambda\| + r \\
 &\iff (r_\lambda - r)^2 \geq \|c - c_\lambda\|^2 \quad \& \quad r_\lambda - r \geq 0 \\
 &\iff r_\lambda^2 - 2rr_\lambda + r^2 \geq \|c - c_\lambda\|^2 \quad \& \quad r_\lambda - r \geq 0 \\
 &\iff 0 \geq \text{pow}(c, b_\lambda) + 2rr_\lambda - r^2 \quad \& \quad r_\lambda - r \geq 0 \\
 &\iff 0 \geq \lambda^2 \left(1 - \frac{y_c}{2f}\right) - 2\lambda x_c + \|c\|^2 - 2fy_c + 2rr_\lambda - r^2 \quad \& \quad r_\lambda - r \geq 0 \\
 &\iff 0 \geq A\lambda^2 + B\lambda + C \quad \& \quad r_\lambda - r \geq 0
 \end{aligned}$$

where

$$\begin{aligned}
 A &= 1 - \frac{y_c - r}{2f} \\
 B &= -2x_c \\
 C &= \|c\|^2 - 2f(y_c - r) - r^2
 \end{aligned}$$

The two inequalities above are sufficient to design an appropriate predicate. Note that when $r = 0$, the coefficients of the quadratic function coincide with the point constraint case. We again examine the possible geometric configurations and roots.

Property 5.11. *Let b be a ball and consider a pl-pencil parameterized by $\lambda \in \mathbb{R}$, $\{b_\lambda\}_\lambda$. Let F be the focus of the pl-pencil, L the horizontal line through F , and H_L^+ the closed half-space above L . Denote by Λ the collection of λ values such that $b \subseteq b_\lambda$. We distinguish eight distinct geometric configurations.*

(a) $b \subseteq H_L^+$, and we have three sub-cases, see Figure 5.12.

- b is strictly above L , then Λ is the union of two disjoint, closed half-lines.
- b is tangent to L at F (and lies above L), then Λ is empty.

- b is tangent to L at a point distinct from F (and lies above L), then Λ is a closed half-line.

(b) $\overset{\circ}{b} \not\subseteq \bigcup_{\mathbb{R}} \overset{\circ}{b}_{\lambda}$, then Λ is the empty set. Geometrically we have two (non-exclusive) sub-cases (see Figure 5.13):

- b intersects the open half-space below the horizontal axis, or
- b contains F in its interior.

Cases (a) and (b) are mutually exclusive. For the remaining cases, suppose that neither (a) nor (b) apply, thus we have $b \not\subseteq H_L^+$ and $\overset{\circ}{b} \subseteq \bigcup_{\mathbb{R}} \overset{\circ}{b}_{\lambda}$.

(c) ∂b intersects the boundary of $\bigcup_{\mathbb{R}} \overset{\circ}{b}_{\lambda}$, then Λ is a singleton. Geometrically we have two (non-exclusive) sub-cases (see Figure 5.14):

- b is tangent to the horizontal axis, or
- $F \in \partial b$.

Note that when both sub-cases above apply, then b belongs to the pl-pencil.

(d) When none of the above applies, then Λ is a (non-singleton) closed interval.

When b is tangent to the horizontal axis with $F \in \partial b$, it belongs to the pl-pencil. In every other case, when $\Lambda \neq \emptyset$, the inclusion is strict, with tangency at the endpoint(s) of Λ .

To prove this, we rely on the below lemma.

Lemma 5.12. *Let b_{λ} be a ball of a pl-pencil and b a ball of constraint. The ball constraint inequality for pl-pencils $A\lambda^2 + B\lambda + C \leq 0$ holds if and only if we have either $b \subseteq b_{\lambda}$ or $b_{\lambda} \subseteq b$, with equality when the balls are tangent.*

Proof. See the derivation of the inequality. □

The second inequality $r_{\lambda} - r \geq 0$ allows to differentiate between the two possibilities of this lemma.

We now investigate the roots of the polynomial and the corresponding geometric configurations, starting from the degenerate cases. Let L be the horizontal line through F . Observe that $A = 0$ if and only if the center c is above L and b is tangent to L . Indeed that is the only configuration where we have $y_c - r = 2f$. As previously, $B = 0$ if and only if c lies on the ordinate axis.

§ $A = B = 0$ The ball b is tangent to L at point F , and is above L . Because $y_c - r = 2f$ and $x_c = 0$ we get $C = 4fr > 0$. It follows that no ball of the family can cover b . Geometrically, recall that every ball b_{λ} goes through F , hence if we had $b \subseteq b_{\lambda}$, then necessarily b and b_{λ} would be tangent at F . That would imply that b_{λ} admits L as a tangent. However, the only ball that admits L as a tangent is $b_{\lambda=0}$, which lies underneath L . Because b is above L , there are no ball b_{λ} that satisfy the desired inclusion.

§ $A = 0, B \neq 0$ Ball b is tangent to L at some point distinct from F , and is above L . We have $C = x_c^2 + 4fr > 0$ with $x_c \neq 0$, and the inequality becomes:

$$\begin{aligned} -2x_c\lambda + C \leq 0 &\iff -2x_c\lambda \leq -C \\ &\iff -2\operatorname{sign}(x_c)|x_c|\lambda \leq -C \\ &\iff 2\operatorname{sign}(x_c)\lambda \geq \frac{C}{|x_c|} = |x_c| + \frac{4fr}{|x_c|} \end{aligned}$$

We show that the half-line of λ values that satisfy the inequality corresponds to balls b_λ that satisfy the desired inclusion. To do so, we show that the continuous function $\lambda \mapsto r_\lambda - r$ must have constant sign over that half-line, and that some λ values of that half-line satisfy $r_\lambda - r \geq 0$.

By continuity, the sign of $r_\lambda - r$ may only change when r_λ reaches the value r . Suppose for a contradiction that there is some value λ_* such that $r_{\lambda_*} = r$ and $A\lambda_*^2 + B\lambda_* + C = -2x_c\lambda + C \leq 0$. By Lemma 5.12 the inequality $A\lambda_*^2 + B\lambda_* + C \leq 0$ implies that either $b \subseteq b_{\lambda_*}$ or $b_{\lambda_*} \subseteq b$. In both cases this implies that $b = b_{\lambda_*}$. However since b does not belong to the family, this λ_* value cannot exist. It follows that $r_\lambda - r$ must have a constant sign over the half-line of solutions $-2x_c\lambda + C \leq 0$. By Equation (5.4) we have $r_\lambda = \frac{\lambda^2}{4f} + f$, thus r_λ is unbounded over the half-line. It follows that for every λ of the half-line, $r_\lambda - r \geq 0$, and $b \subseteq b_\lambda$.

§ $A \neq 0$ Ball b is not tangent to L , above L . Note that it may still be tangent to L , below L . We inspect the discriminant:

$$\begin{aligned} \frac{\Delta}{4} &= x_c^2 - \left(1 - \frac{y_c - r}{2f}\right) \left(\|c\|^2 - 2f(y_c - r) - r^2\right) \\ &= -y_c^2 + 2f(y_c - r) + r^2 + \frac{y_c - r}{2f} \left(\|c\|^2 - 2f(y_c - r) - r^2\right) \\ &= \frac{y_c - r}{2f} \left(-2f(y_c + r) + 4f^2 + \|c\|^2 - 2f(y_c - r) - r^2\right) \\ &= \frac{y_c - r}{2f} (-4fy_c + 4f^2 + x_c^2 + y_c^2 - r^2) \\ &= \frac{y_c - r}{2f} \left(x_c^2 + (y_c - 2f)^2 - r^2\right) \\ &= \frac{y_c - r}{2f} \left(\|c - F\|^2 - r^2\right) \\ &= \frac{y_c - r}{2f} \operatorname{pow}(F, b) \end{aligned}$$

The sign of the discriminant can thus be deduced from the sign of the two factors $y_c - r$ and $\operatorname{pow}(F, b)$. The former is positive when b is contained in the half-space of positive ordinate points; zero when b is tangent to the abscissae axis and contained in the closed half-space of non-negative ordinate points; negative when b intersects the half-space of negative ordinate points. The latter is positive when b does not contain F ; zero when b goes through F ; negative when b contains F in its interior.

From the previous section, if $y_c - r < 0$, then no ball in the family may cover b , because the lowest point of ball b then lies outside the half-space covered by the family. Likewise, if $\text{pow}(F, b) < 0$, no ball b_λ may contain b because for any λ value, we have $F \in \partial b_\lambda$. In light of these remarks, we can restrict our analysis to $y_c - r \geq 0$ and $\text{pow}(F, b) \geq 0$, which implies that $\Delta \geq 0$.

(a) $y_c - r = 0$ and $\text{pow}(F, b) = 0$, which implies $\Delta = 0$, and the polynomial admits one unique root λ_* . Since $y_c - r = 0$ we have $A = 1 > 0$, thus λ_* is the unique value at which the polynomial is non-positive. In fact, the ball b is tangent to the directrix \mathcal{D} and goes through F , it thus belong to the pl-pencil, it is the unique ball of the family that covers itself.

(b) $y_c - r = 0$ and $\text{pow}(F, b) > 0$, which implies $\Delta = 0$ and $A > 0$. The polynomial again admits one unique root λ_* , which corresponds to its unique non-positive value. Because $\text{pow}(F, b) > 0$, $F \notin b$ and we are guaranteed by Lemma 5.12 that $b \subseteq b_{\lambda_*}$.

(c) $\text{pow}(F, b) = 0$ and $y_c - r > 0$, which implies $\Delta = 0$. The polynomial has one unique root λ_* . Because the distance between c and F is at least the difference in their ordinates, we have $r = \|c - F\| \geq |y_c - 2f|$. This yields $2f \geq y_c - r$ and thus $A \geq 0$. Therefore λ_* is the unique value at which the polynomial is non-positive. Because $y_c - r > 0$, b does not contain any point of the directrix \mathcal{D} , and we obtain $b \subseteq b_{\lambda_*}$ by Lemma 5.12.

(d) $y_c - r > 0$ and $\text{pow}(F, b) > 0$, which implies $\Delta > 0$. The polynomial admits two distinct roots $\lambda_1 < \lambda_2$. Because b does not contain F or any point of the directrix \mathcal{D} , Lemma 5.12 implies that $b \subseteq b_\lambda$ for every λ that satisfy the inequality. Note however that A may be positive or negative, contrary to the previous cases. Observe that if b is contained in the open half-space above L , the horizontal line through point F , then $A < 0$, otherwise $A > 0$. Thus if b intersects the half-space below L , we have inclusion for every $\lambda \in [\lambda_1, \lambda_2]$. If b is above L , we have inclusion for every $\lambda \in (-\infty, \lambda_1] \cup [\lambda_2, +\infty)$.

5.5.4 Arc constraint

We address in this section the problem of arc constraints. We need some additional definitions. We let $b_0 = b_{\lambda=0}$, by Equation (5.4) it is the ball with smallest radius in the family. It is tangent to the horizontal line L at F , and lies beneath L . The ball with parameter value $\lambda = 0$ has several useful properties.

Lemma 5.13. *Consider $\lambda \in \mathbb{R}$ and let $F'_\lambda(2\lambda, 2f)$. We denote by L the horizontal line through the focus. Then we have $b_\lambda \cap L = [FF'_\lambda]$.*

Proof. Recall that $F \in \partial b_\lambda$, hence $\text{pow}(F, b_\lambda) = \lambda^2 + (2f - r_\lambda)^2 - r_\lambda^2 = 0$. Consider $M(x_M, 2f) \in L$.

$$\begin{aligned} \text{pow}(M, b_\lambda) \leq 0 &\iff (x_M - \lambda)^2 + (2f - r_\lambda)^2 - r_\lambda^2 \leq 0 \\ &\iff x_M^2 - 2\lambda x_M + \lambda^2 - \lambda^2 \leq 0 \\ &\iff |x_M| \leq 2\lambda \text{sign}(x_M) \end{aligned}$$

Thus depending on the sign of λ , we deduce that M belongs to b_λ if and only if $0 \leq x_M \leq 2\lambda$ or $2\lambda \leq x_M \leq 0$. Therefore we indeed have $b_\lambda \cap L = [FF'_\lambda]$. \square

Lemma 5.14. *Let H_L^+ be the closed half-space above L . Then for all $0 \leq \lambda \leq \mu$ we have $b_\lambda \cap H_L^+ \subseteq b_\mu \cap H_L^+$. Likewise for $\mu \leq \lambda \leq 0$ we have $b_\lambda \cap H_L^+ \subseteq b_\mu \cap H_L^+$.*

Proof. We only address the case $0 \leq \lambda \leq \mu$. Consider a point $M(x_M, y_M)$ that lies strictly above L . Notice that $M \notin b_0$. We know from Property 5.10 that the collection Λ_M of λ values such that $M \in b_\lambda$ is the union of two disjoint, closed half-lines. In particular, $\mathbb{R} \setminus \Lambda_M$ is an interval, thus it is convex. Consider now $\lambda \geq 0$ such that $M \in b_\lambda$. Then for all $\mu \geq \lambda$, we necessarily have $M \in b_\mu$. Indeed, we know that $0 \notin \Lambda_M$ whereas $\lambda \in \Lambda_M$. By convexity, $\mu \notin \Lambda_M$. It follows from the remark above and Lemma 5.13 that for any $0 \leq \lambda \leq \mu$, we have $b_\lambda \cap H_L^+ \subseteq b_\mu \cap H_L^+$. \square

In the following discussion, it will be useful to consider the two sub-families delimited by b_0 . We let \mathcal{P}^+ the collection of balls with non-negative λ parameter, $\mathcal{P}^+ = \{b_\lambda \mid \lambda \geq 0\}$, and \mathcal{P}^- those with non-positive λ parameter, $\mathcal{P}^- = \{b_\lambda \mid \lambda \leq 0\}$.

The next lemma can be seen as the equivalent of Lemma 5.14 for points below L , but above the horizontal axis.

Lemma 5.15. *Let D be the stripe of points strictly below L and strictly above the horizontal axis. Let $\lambda_* \in \mathbb{R}$. Then for all $\lambda_* \leq \lambda \leq \mu$ we have $b_\lambda \cap (D \cap b_{\lambda_*}) \supseteq b_\mu \cap (D \cap b_{\lambda_*})$. Likewise for $\mu \leq \lambda \leq \lambda_*$ we have $b_\lambda \cap (D \cap b_{\lambda_*}) \supseteq b_\mu \cap (D \cap b_{\lambda_*})$.*

Proof. We only address the case $\lambda_* \leq \lambda \leq \mu$. Consider a point $M(x_M, y_M) \in D \cap b_{\lambda_*}$. We know from Property 5.10 that the collection Λ_M of λ values such that $M \in b_\lambda$ is a closed interval, it is thus convex. Consider now $\lambda \geq \lambda_*$ such that $M \notin b_\lambda$. Then for all $\mu \geq \lambda$, we deduce by convexity of Λ_M that $M \notin b_\mu$. It follows that for any $\lambda_* \leq \lambda \leq \mu$, we have $b_\lambda \cap (D \cap b_{\lambda_*}) \supseteq b_\mu \cap (D \cap b_{\lambda_*})$. \square

We are now ready to look for the inclusion λ values. Consider a circular arc e supported by a ball b , and Λ_b the collection of λ values for which $b \subseteq b_\lambda$. Based on the previous section, there are eight distinct cases to consider: (a) b lies strictly above line L ; (b) b lies above L , and is tangent to L at F ; (c) b lies above L , and is tangent to L at points distinct from F ; (d) b intersects the open half-space of negative ordinates; (e) F lies in the interior of b ; (f) b is above the horizontal axis, tangent to it; (g) b contains F in its boundary; and (h) the remaining case.

(a) When b is strictly above line L , we know that there are $\lambda_1 < \lambda_2$ such that $\Lambda_b = \mathbb{R} \setminus (\lambda_1, \lambda_2)$. We look at the sub-family of balls $\{b_\mu \mid \lambda_1 < \mu < \lambda_2\}$. Let x_i be the (unique) tangent point between b and b_{λ_i} . Because x_i necessarily lies above L , we know by Property 5.10 that the λ values such that $x_i \in b_\lambda$ is the union of two closed half-lines. In particular the collection of λ values for which $x_i \notin b_\lambda$ is convex. A similar reasoning as in Section 5.4.4, case (c), shows

that for $\mu < \lambda_2$ close enough to λ_2 , $x_2 \notin b_\mu$. Because b_0 lies entirely below L , we also have $x_2 \notin b_0$. It follows that for $0 \leq \mu < \lambda_2$, we can use x_2 to apply our usual method. Likewise, for $\lambda_1 < \mu \leq 0$ we can use x_1 instead. See Figure 5.12.a.

(b) When b lies above L , and is tangent to L at F , we know that $\Lambda_b = \emptyset$. Ball b and every ball b_λ contain F in their boundary. Consider b_0 , it is tangent to L and below L . Hence $b \cap b_0 = \{F\}$, and b_0 cannot cover $e \subseteq \partial b$. We can thus exclude b_0 from our considerations. For the remaining λ values, b and b_λ are not tangent, and it follows that $\partial b \cap \partial b_\lambda$ is a circular arc of non-zero length that admits F as an endpoint. We first look at the sub-family \mathcal{P}^+ . Let H_L^+ be the closed half-space above L , we have $e \subseteq \partial b \subseteq H_L^+$. By Lemma 5.14 we deduce $b_\lambda \cap \partial b \subseteq b_\mu \cap \partial b$ for all $0 \leq \lambda \leq \mu$.

We can conclude as in case 2(b) of Section 5.3.4. Let $\lambda_* > 0$ the smallest positive parameter value such that b_{λ_*} contains both endpoints of e . If b_{λ_*} does not contain e , then no ball b_μ with $\mu > 0$ does. Otherwise b_μ , $\mu > 0$, contains e if and only if $\mu \geq \lambda_*$. See Figure 5.16. By symmetry this reasoning can also be applied to negative λ values and we can also find the balls of \mathcal{P}^- that contain e .

(c) When b lies above L , and is tangent to L at a point distinct from F , we know that Λ_b is a closed half-line. By symmetry, suppose $\Lambda_b = (-\infty, \lambda_*]$. We have to look at the sub-family of balls $\mathcal{B} = \{b_\mu \mid \mu > \lambda_*\}$. We further split this family into $\mathcal{B}^- = \mathcal{B} \cap \mathcal{P}^-$ and $\mathcal{B}^+ = \mathcal{B} \cap \mathcal{P}^+$. By Property 5.11, let x be the unique tangency point between b and b_{λ_*} . A reasoning similar as in case (a) (of this section) shows that no ball of \mathcal{B}^- contains x . For the family \mathcal{B}^+ , consider instead the tangency point between b and L . Thus in both case, we can reduce to our usual method. See points x and y in Figure 5.12.c.

For the remaining cases, we can assume that b is not above L .

(d) When b meets the open half-space of points with negative ordinate, then the lowest point of b is never covered by any ball of the pl-pencil. This lowest point can be used to derive the λ inclusion values for e . See point x in Figure 5.13.a.

For the remaining cases, we can assume that b is above the horizontal axis.

(e) When b contains the focus F in its interior, then ∂b intersects line L in two distinct points F_- and F_+ such that $F \in]F_-F_+[$. Suppose that F_- is on the left side of F_+ . See Figure 5.13.b. We inspect the two sub-families \mathcal{P}^- and \mathcal{P}^+ . By Lemma 5.13, balls in \mathcal{P}^+ never cover any point of L that are to the left of F , thus it never covers F_- . Likewise, \mathcal{P}^- never covers F_+ . We can thus use our usual method on the two sub-families and retrieve every inclusion values for e .

For the remaining cases, we can assume that F is not in the interior of b .

(f) If b is tangent to the horizontal axis, consider the point x of tangency. See Figure 5.14.a. Indeed, for any point of the horizontal axis, there is one unique ball of the family that contains it. We can thus use this tangency point x to perform our usual method.

For the remaining cases, we can assume that b is strictly above the horizontal axis.

(g) If b contains F in its boundary, and excluding all the previous cases, we know from Property 5.11 that $\Lambda_b = \{\lambda_*\}$ is a singleton. We have the strict inclusion $b \subsetneq b_{\lambda_*}$, and b and b_{λ_*} are tangent at F . We distinguish two more subcases.

- b lies underneath L , then we necessarily have $\lambda_* = 0$. The whole arc e thus lies below L . By Lemma 5.15 we deduce that for $0 \leq \lambda \leq \mu$ we have $e \cap b_\lambda \supseteq e \cap b_\mu$. Let μ_* the maximum value such that b_{λ_*} still contains both endpoints of e . If b_{μ_*} contains e , then the balls $b_\mu \in \mathcal{P}^+$ that contain e are precisely the balls with $0 \leq \mu \leq \mu_*$. If b_{μ_*} does not contain e , then no ball of \mathcal{P}^+ contain e . By symmetry, this also covers the family \mathcal{P}^- .
- b does not lie below L , and ∂b meets L at two distinct points. We split the family into three sub-families.

Let F' be the second intersection point and suppose that F' lies on the right side of F . Then for the collection of balls \mathcal{P}^- we can use point F' to derive the inclusion values for e , as per the standard method. See Figure 5.18.

For the balls b_μ , $\mu > \lambda_*$, we know by Lemma 5.14 that the portion of b above L will be covered by every b_μ . Lemma 5.15 also indicates that as μ decreases, the portion of b between L and the horizontal axis that is covered by b_μ will increase. It follows that for $\mu \geq \lambda \geq \lambda_*$, we have $e \cap b_\mu \subseteq e \cap b_\lambda$. Let $\mu_* > \lambda_*$ the largest parameter value such that b_{μ_*} contains both endpoints of e . If b_{μ_*} does not contain e , then no ball b_μ with $\mu > \lambda_*$ does. Otherwise b_μ , $\mu > \lambda_*$, covers e if and only if $\mu \leq \mu_*$. See Figure 5.19.a.

For the remaining balls b_μ , we have $0 \leq \mu < \lambda_*$. Through Lemmas 5.14 and 5.15 we deduce that $e \cap b_\lambda \subseteq e \cap b_\mu$ for all $0 \leq \lambda \leq \mu < \lambda_*$. Let μ_* be the minimum parameter value such that b_{μ_*} contains both endpoints of e . Then b_μ contains e if and only if $e \subseteq b_{\mu_*}$ and $\mu \geq \mu_*$. See Figure 5.19.b.

For the last case, we can assume that F does not belong to b .

(h) When excluding all of the previous cases, we know from Property 5.11 that there are $\lambda_1 < \lambda_2$ such that $\Lambda_b = [\lambda_1, \lambda_2]$. Consider x_i the unique tangent point between b and b_{λ_i} . If x_2 lies below L , then no ball b_μ with $\mu > \lambda_2$ contains x_2 . Likewise if x_1 lies below L , no ball b_μ with $\mu < \lambda_1$ contains x_1 . We can then apply our usual method with x_1 and x_2 .

Otherwise, either x_1 or x_2 lies strictly above L , but not both. Without loss of generality, suppose x_1 is above L . Then for all $0 \leq \mu < \lambda_1$ we have $x_1 \notin b_\mu$ and we can apply our usual method on that range. For $\mu \leq 0$, consider instead any intersection point between ∂b and L . See point y in Figure 5.15.

Once again, the arc constraint predicate can be deduced from the point and ball predicates.

5.6 Conclusion

Because the computation of critical balls only relies on a few predicates, it is possible to generalize Algorithm 1, that we used to compute (δ, ε) -ball approximations for finite unions of balls,

to other classes of shapes. Indeed it suffices to show that predicates can be implemented for those new classes of shapes. In order to accommodate for this generalization, it is interesting to decouple the inner constraint set from the outer one as much as possible, thus leading to (O, I) -ball approximations.

It can be shown that the medial axis of a polygon can be decomposed into portions of point-point, line-line, and point-line pencils. Our study of predicates for these families of balls thus extends our approach to polygons. It should be noted that because we have yet to perform an in-depth study of the new pencils, there still remains some uncertainty on the optimality and termination claim, which were both left as conjectures.

A natural extension of our algorithm in \mathbb{R}^2 would be to investigate the existence of predicates for “arc-point”, “arc-line”, and “arc-arc” pencils, thus permitting both the inner and outer subset to span the same class of shapes.

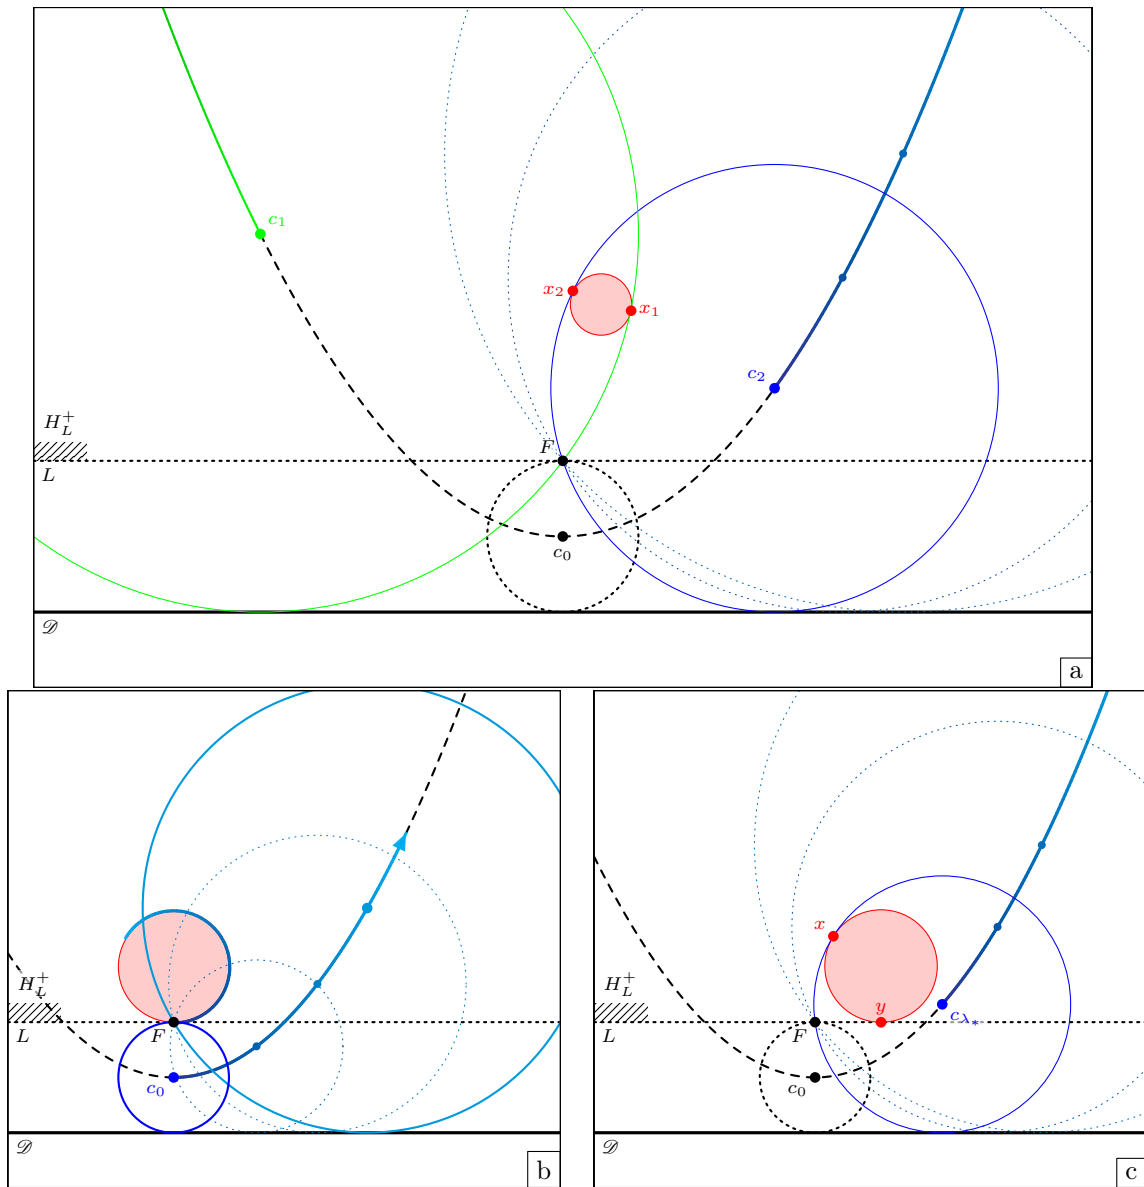


Figure 5.12 – Ball constraint in pl-pencils lying above L : (a) Λ is the union of two half-lines; (b) Λ is empty; (c) Λ is a half-line.




Figure 5.13 – Ball constraint in pl-pencils lying out of bounds: (a) the ball intersects the half-space of negative ordinates; (b) the ball contains the focus in its interior.




Figure 5.14 – Ball constraint in pl-pencils that are tangent to the bounds: (a) the ball is tangent to the directrix; (b) the boundary of the ball goes through F .




Figure 5.15 – Ball constraint in pl-pencils: last case.




Figure 5.16 – Arc constraint in pl-pencils: support ball above and tangent to L at F .




Figure 5.17 – Arc constraint in pl-pencils: support ball below and tangent to L at F .




Figure 5.18 – Arc constraint in pl-pencil whose support ball goes through F : sub-family with $\mu \leq 0$.

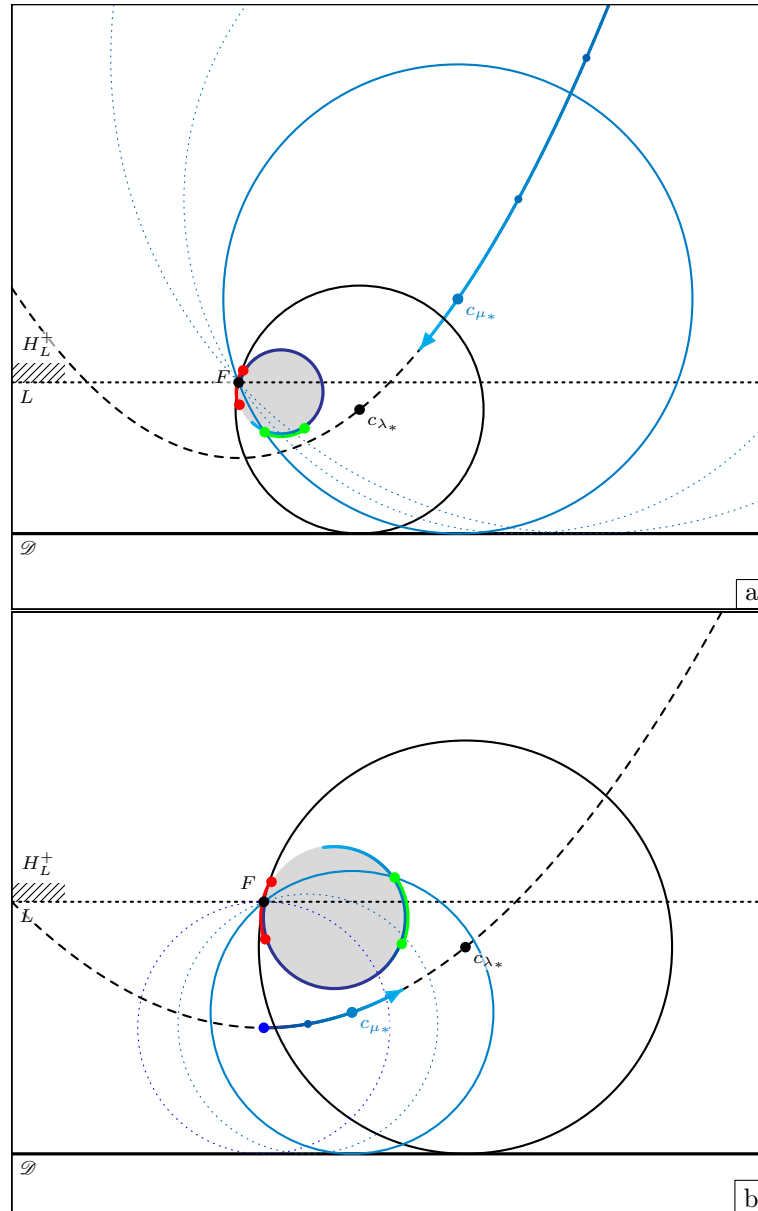


Figure 5.19 – Arc constraint in pl-pencil whose support ball goes through F : (a) sub-family with $\mu > \lambda_*$; (b) sub-family with $0 \leq \mu < \lambda_*$. Only balls of the sub-families centered between c_{λ_*} and c_{μ_*} contain the green arcs. No ball may contain the red arcs.

Conclusion and outlooks

Conclusion

In this thesis, we investigated the problem of shape approximation by finite unions of balls. We based our approach on the use of pencils of balls and their many properties. We first used pencils to robustly describe finite unions of balls, their boundary, and their medial axis. Indeed, the framework of pencils can formally describe and handle every degenerate configuration of balls. We then built on these results to prove a characterization of set inclusion within finite unions of balls. The main appeal of this result is that it provides more easily tractable conditions to test for set inclusion. This characterization was used to prove several properties related to our new approximation scheme, the (δ, ε) -ball approximations problem. Specifically, it showed that it is possible to compute critical balls, which provide an algorithmic solution to the candidate ball heuristic.

Although we proved that the (δ, ε) -ball approximation problem is NP-complete, we have presented a greedy optimal algorithm for shapes of \mathbb{R}^2 whose medial axis is cycle-free. This algorithm is based on the intuitive idea of candidate balls to a T -small approximation. To derive a practical implementation from this heuristic, we studied the properties of critical balls, and proved that they can be computed, given simple predicates. We then extended the scope of this algorithm by showing that predicates can be implemented for other classes of shapes.

There are many questions related to this work that we couldn't address. Regarding the structure of finite unions of balls, the most obvious topic is probably the implementation of an algorithm to compute the medial axis. Recall that every face f of a finite union of balls has a well-defined join, and that it is possible to obtain the contribution J_f of that join to the medial axis. Naively, it is possible to loop over every faces of the union and compute every contribution J_f . However those contributions can overlap significantly. We ask whether it is possible to minimize the number of faces to process to retrieve the complete medial axis.

In our definition of (δ, ε) -ball approximation, we can control the geometric error of the approximation through the two parameters δ and ε . This paradigm however does not provide much guarantee in terms of topology of the resulting approximation. Indeed, if the collection \mathcal{B} is a (δ, ε) -ball approximation of S with $\bigcup \mathcal{B} \subsetneq S^{\oplus \delta}$, then we can then easily derive other approximations of S with arbitrarily many connected components. It indeed suffices to find disjoint balls in $S^{\oplus \delta} \setminus \bigcup \mathcal{B}$ and include them in the approximation. The resulting collection will still be a valid (δ, ε) -ball approximation of S . If we however constrain \mathcal{B} to be a cardinal minimum approximation, then we cannot have these superfluous connected components. We wonder if it is possible to derive any topological guarantees for cardinal minimum approximations.

Although we spent great efforts to describe Algorithm 1, and show that it reaches an

optimal solution, in finite time, for a certain class of shapes, it has not been implemented yet. It would thus be interesting to implement it and compare its output to other approximation methods. Regarding the generalized algorithm, we are also close to proving that it also reaches an optimal solution in finite time.

Finally, the issue of generalizing our method to higher dimensions remains an open problem. We share some of our investigation in the next section.

(O, I) -ball approximations in higher dimensions

In the last few chapters, many of the results we presented were only true for \mathbb{R}^2 . We now provide some insights on the \mathbb{R}^d case when the outer subset O is a finite union of balls. In particular, we investigate how our approach with critical and candidate balls can be adapted. We outline some of the issues with candidate balls and present an alternative heuristic that relies on tools similar to critical balls.

§ Partial order and plane sweep The first difficulty resides in the choice of a partial T -order over the medial axis. We suggest using a plane sweep. Given a sweep direction \mathbf{e} , we let H_t be the hyperplanes with normal \mathbf{e} that intersects the line $\mathbb{R}\mathbf{e}$ at elevation t . If we inspect the intersection between $\text{MA}(O)$ and H_t for various values of $t \in \mathbb{R}$, then there are finitely many singular values $t_1 < t_2 < \dots < t_k$ at which the topology of the intersection will change. If we slice the medial axis $\text{MA}(O)$ at these singular values, we can derive a subdivision of $\text{MA}(O)$ into path-connected components. Assuming that $\text{MA}(O)$ does not contain any cycle (of any dimension), these subdivided components then admit a topological ordering. We can thus obtain a consistent orientation of $\text{MA}(O)$ by choosing between the two sweep directions $+\mathbf{e}$ or $-\mathbf{e}$ for each components, in accordance with their topological order. This yields a partial T -order on $\text{MA}(O)$.

§ Critical ball predicates and convex optimization Given any partial T -order on $\text{MA}(O)$, the notion of critical ball for a point or a set directly carries over from the 2D case. When using the partial T -order that we introduced above, we argue that computing a critical ball for any type of constraint can be reduced to a convex optimization problem. First, observe that we can reduce to finding critical balls for a constraint within finitely many convex subsets of elliptic pencils. Combining these “locally-critical” balls to find a “globally-critical” one reduces to comparing the elevation t of the hyperplanes H_t that contain the locally-critical balls. Within a single (subdivided) face f of $\text{MA}(O)$, the sweep direction is constant. Consider the elliptic pencil \mathcal{P} such that $f \subseteq c(\mathcal{P})$, and let $\mathcal{B}_f = \{b \in \mathcal{P} \mid c(b) \in f\}$ the corresponding balls centered at points of f . Recall that we say a ball b_* is critical for point x if it is T -maximal for the collection $F_x = \{b \text{ medial in } O \mid x \in b\}$. We instead look at $F_x^{\mathcal{P}} = \{b \in \mathcal{P} \mid x \in b\}$. Notice that $F_x \cap \mathcal{P} = F_x^{\mathcal{P}} \cap \mathcal{B}_f$. With the plane sweep T -order, being critical implies that $c(F_x^{\mathcal{P}}) \cap f$ must be beneath the sweeping plane H_{t_*} that contains $c(b_*)$.

Because $F_x^{\mathcal{P}} \subseteq \mathcal{P}$, we can inspect the convex hull $\text{CH}(F_x^{\mathcal{P}}) \subseteq \mathcal{P}$. By convexity of the various subsets involved, we deduce that b_* is also T -maximal for $\text{CH}(F_x^{\mathcal{P}}) \cap \mathcal{B}_f$. We obtain a similar result for a set of constraint.

From the above, computing a critical ball for any constraint is equivalent to finding a tangent hyperplane to a convex subset. Whether that convex subset can easily be computed or not is of course dependant on the specific subset. We claim that the predicates for point constraints can always be implemented in constant time for pencils of balls in the sweep line partial order. Additionally, for an elliptic k -pencil and a point of constraint $x \notin V(\mathcal{P})$, it can be shown that F_x (and $c(F_x)$) is a half k -flat, and ∂F_x is an elliptic $(k - 1)$ -subpencil of \mathcal{P} .

§ Termination of the algorithm ? From the above, it is possible to apply the heuristic of computing critical balls to any dimension, however it is yet unclear whether such an algorithm will always terminate. Indeed in 2D, we relied on the fact that for any given pencil segment $[b_-b_+]$ we have $\bigcup[b_-b_+] = b_- \cup b_+$. This prevented the greedy algorithm from ever choosing more than two balls within the same pencil segment. For higher dimensions, we have convex pencil polytopes rather than just 1-dimensional segments. Although we know that these pencil polytopes can be expressed as the convex hull of finitely many balls, the algorithm computes each critical ball sequentially, and does not make use of that structural property. When running the algorithm, we may thus end up with more balls in a pencil polytope than the number of vertices of that polytope.

§ On the choice of a partial order Even though predicates to compute critical balls can be made available, we cannot guarantee a finite-time algorithm. This approach is thus unlikely to reach an optimal solution. The problem of using this heuristic in higher dimensions is tied to the choice of the partial T -order, but also to the notion of candidate ball. Indeed, the notion of T -criticality will be well-defined whenever the partial T -order is. However, any arbitrary choice of T -order will generally not yield any meaningful property for the corresponding critical balls. In 2D, one key property we relied on to prove the many properties of critical balls, was that the removal of any point of $\text{MA}(O)$ would split $\text{MA}(O)$ in a specific number of (path-)connected components. In general, the medial axis of a shape in \mathbb{R}^d is a $(d - 1)$ -dimensional surface, and that splitting property will rarely hold. Splitting the medial axis indeed implied splitting the interior of the outer shape $\overset{\circ}{O}$, and we could then exploit local properties to derive global properties.

Remark 67. When the medial axis of O is 1-dimensional, this “splitting” property holds and our approach still works.

The \mathbb{R}^2 assumption that $\text{MA}(O)$ should be cycle-free was to guarantee this splitting property. Although the condition that $\text{MA}(O)$ should be free of cycle of any dimension seems attractive, we are unaware of any set of properties that $\text{MA}(O)$ should satisfy in the \mathbb{R}^d case in order to define a good partial T -order, with good properties on the resulting critical balls, and that would then allow us to apply this heuristic properly. In particular, our current pro-

positional is unable to obtain the “substitution” property that we used in the proof of Theorem 10. This property enabled us to transform any (O, I) -ball approximation into an approximation that only contained candidates to its T -small subsets.

§ T -tight collections and substitution heuristic We suggest a change in heuristic. Recall that we introduced critical balls as a computationally-friendly (and equivalent) alternative to candidate balls. We thus tailored the set of constraints for our critical balls so that they would coincide with candidate balls. If we alter the way we define the constraints and compute critical balls, we will then obtain a new heuristic.

Consider a finite collection of balls \mathcal{B} with $I \subseteq \bigcup \mathcal{B} \subseteq O$. Let $b \in \mathcal{B}$ and $\mathcal{B}' = \mathcal{B} \setminus \{b\}$ the balls in \mathcal{B} minus b . We denote the points of I that are exclusively covered by b in $\bigcup \mathcal{B}$ as $A(I, b, \mathcal{B}) = I \setminus \bigcup(\mathcal{B}')$. We then inspect the collection of medial balls that contain $A(I, b, \mathcal{B})$, $F(I, b, \mathcal{B}) = \{b \text{ medial in } O \mid A(I, b, \mathcal{B}) \subseteq b\}$, and look for its T -maximal element b_* . We call b_* a (T) -tight ball for the set $A(I, b, \mathcal{B})$.

Remark 68. Let $F_x = \{b \text{ medial in } O \mid x \in b\}$, and notice that

$$F(I, b, \mathcal{B}) = \bigcap_{x \in A(I, b, \mathcal{B})} F_x$$

Recall that F_x is convex, thus $F(I, b, \mathcal{B})$ is convex. In a plane sweep T -order, computing a tight ball reduces to a convex optimization problem.

Remark 69. Because $b \in F(I, b, \mathcal{B})$, this collection of balls is non-empty and always admits a T -maximal element.

A tight ball b_* will satisfy the inclusion $A(I, b, \mathcal{B}) \subseteq b_*$ and it follows that we can replace b by b_* in \mathcal{B} . The resulting collection \mathcal{B}_* still satisfy the two inclusions $I \subseteq \bigcup \mathcal{B}_* \subseteq O$. We call a collection of ball \mathcal{B} (T) -tight for I if each of its ball $b \in \mathcal{B}$ is tight for $A(I, b, \mathcal{B})$, with $A(I, b, \mathcal{B}) \neq \emptyset$.

This suggests the following substitution heuristic. Starting from an initial (O, I) -ball approximation \mathcal{B}_0 , inspect the set of points exclusively covered by a ball b . If it is empty, remove the corresponding ball. If it is non-tight for $A(I, b, \mathcal{B})$, then replace it by the tight ball for $A(I, b, \mathcal{B})$. See Algorithm 3.

Remark 70. In \mathbb{R}^2 , the substitution heuristic and the candidate ball heuristic coincide.

It should be noted that this substitution heuristic may not terminate, and Algorithm 3 may loop indefinitely. Since the collection of medial balls is compact, we can however show that the procedure admits a limit collection of balls \mathcal{B}_{lim} . We conjecture the following:

Conjecture 4. Algorithm 3 outputs (potentially at infinity) a tight collection for I .

Throughout this thesis, we often had the problem to look for a collection \mathcal{B} with property (i) that is difficult to compute directly, but knowing about another property (ii) with some

Algorithm 3 Substitution heuristic**Input:** \mathcal{B}_0 a finite collection of balls, $I \subseteq \bigcup \mathcal{B}_0 \subseteq O$

```

1:  $\mathcal{B} \leftarrow \mathcal{B}_0$ 
2: While  $\exists b \in \mathcal{B}$  with  $A(I, b, \mathcal{B}) = \emptyset$ , or  $\exists b \in \mathcal{B}$  non-tight for  $A(I, b, \mathcal{B})$  do
3:   If  $A(I, b, \mathcal{B}) = \emptyset$  then
4:     Remove  $b$  from  $\mathcal{B}$ 
5:   else
6:     Compute  $b_*$  tight for  $A(I, b, \mathcal{B})$ 
7:     Replace  $b$  by  $b_*$ 
8:   end If
9: end While
10: Return  $\mathcal{B}$ 

```

representatives that are easy to compute. The recurring strategy we used to deal with this problem was to show that every collection \mathcal{B} with property (i) could be transformed into another collection \mathcal{B}' with property (ii) while preserving property (i). From there, we showed that there was one unique equivalence class of collections that can satisfy property (ii), and that any representative of that equivalence class would also fulfill property (i). Computing a representative for property (i) could thus be accomplished by computing a representative for property (ii).

Assuming that Conjecture 4 holds, we know that there exists cardinal-minimum (O, I) -ball approximations that are tight for I . Even if we disregard the issue of termination, it seems unlikely that any initial collection \mathcal{B}_0 in Algorithm 3 will always yield a cardinal minimum collection. The influence of the choice of \mathcal{B}_0 on the output of the procedure is indeed an open question.

Remark 71. When O is a finite union of balls, the collection of balls that defines O can always be used as the initial collection \mathcal{B}_0 . Algorithm 3 is in fact a method that tries to simplify (O, I) -ball approximations.

Although Algorithm 3 also has no guarantee on termination or optimality, it still presents an interesting property, compared to the candidate ball heuristic. Indeed our initial approach by candidate balls terminates if and only if it reaches a (O, I) -ball approximation. If we interrupt the procedure before termination, we do not have an approximation available. On the other hand, Algorithm 3 terminates if and only if it managed to remove every non-local optimality from its stored collection \mathcal{B} , and \mathcal{B} is always has a valid (O, I) -ball approximation. It is thus always possible to interrupt the procedure early.

Remark 72. After interrupting Algorithm 3, it is also easy to resume the procedure, since the state of the algorithm is entirely encoded within the collection \mathcal{B} itself.

Additional technical results

A.1 Boundary of union and intersection

A.1.1 Intersection

Lemma A.1. *Let $A, B \subseteq \mathbb{R}^d$. Then $\partial(A \cap B) \subseteq (\partial A \cap \bar{B}) \cup (\partial B \cap \bar{A})$. Additionally, if A and B are closed, then we have equality.*

Proof.

$$\begin{aligned}
 \partial(A \cap B) &= \overline{A \cap B} \cap \overline{(A \cap B)^c} \\
 &\subseteq \bar{A} \cap \bar{B} \cap \overline{A^c \cup B^c} \\
 &= \bar{A} \cap \bar{B} \cap (\bar{A}^c \cup \bar{B}^c) \\
 &= (\bar{A} \cap \bar{B} \cap \bar{A}^c) \cup (\bar{A} \cap \bar{B} \cap \bar{B}^c) \\
 &= (\partial A \cap \bar{B}) \cup (\partial B \cap \bar{A})
 \end{aligned}$$

If A and B are closed, then $\bar{A} \cap \bar{B} = A \cap B = \overline{A \cap B}$ and we thus have equality. \square

A.1.2 Union

Lemma A.2. *Let $A, B \subseteq \mathbb{R}^d$. Then $\partial(A \cup B) \subseteq (\partial A \setminus \overset{\circ}{B}) \cup (\partial B \setminus \overset{\circ}{A})$. Additionally, if A and B are open, then we have equality.*

Proof.

$$\begin{aligned}
 \partial(A \cup B) &= \overline{A \cup B} \cap \overline{(A \cup B)^c} \\
 &= \overline{A \cup B} \cap \overline{A^c \cap B^c} \\
 &\subseteq (\bar{A} \cup \bar{B}) \cap (\bar{A}^c \cap \bar{B}^c) \\
 &= (\bar{A} \cap \bar{A}^c \cap \bar{B}^c) \cup (\bar{B} \cap \bar{B}^c \cap \bar{A}^c) \\
 &= (\partial A \setminus \overset{\circ}{B}) \cup (\partial B \setminus \overset{\circ}{A})
 \end{aligned}$$

If A and B are open, then A^c and B^c are closed. Thus $\bar{A}^c \cap \bar{B}^c = A^c \cap B^c = \overline{A^c \cap B^c}$, which makes the above development an equivalence. \square

Lemma A.3. *Let $A, B \subseteq \mathbb{R}^d$, such that A and B are disjoint. Then $\partial A \setminus \overset{\circ}{B} = \partial A$.*

Proof. A and B are disjoint, thus $A \cap B = \emptyset$ and in particular, $A \subseteq B^c$. This implies that $\overline{A \cap B^c} = \overline{A}$. We have

$$\begin{aligned} \partial A \setminus \overset{\circ}{B} &= \overline{A} \cap \overline{A^c} \cap \overline{B^c} \\ &= \overline{A} \cap \overline{A^c} \\ &= \partial A. \end{aligned}$$

□

Corollary A.4. *Let $A, B \subseteq \mathbb{R}^d$ be open and disjoint. Then $\partial(A \cup B) = \partial A \cup \partial B$.*

Proof. Stems from lemmas A.2 and A.3. □

In general, the inclusion $\overline{A \cap B} \subseteq \overline{A} \cap \overline{B}$, used in the proofs of lemmas A.1 and A.2 is strict. Indeed, consider for example $X \subseteq \mathbb{R}^d$ with $X \neq \emptyset \neq X^c$. Then $X \cap X^c = \emptyset$ hence $\overline{X \cap X^c} = \emptyset$. However $\emptyset \neq \partial X = \partial(X^c) = \overline{X} \cap \overline{X^c}$.

A.2 Interior of a cone

Lemma A.5. *Let $u, c_1, \dots, c_k \in \mathbb{R}^d$ be $k + 1$ points. Let $C = \{c_1, \dots, c_k\}$. We have the inclusion*

$$\overset{\circ}{\text{coni}}(u, C) \subseteq \left\{ u + \sum_{i=1}^k \lambda_i (c_i - u) \mid \lambda_i \in \mathbb{R}, \lambda_i > 0 \right\}$$

In order to prove Lemma A.5 we rely on the following technical result:

Lemma A.6. *Let $c \in \overset{\circ}{\text{coni}}(u, C)$ and $l \in \{1, \dots, k\}$. There exist $\lambda_1, \dots, \lambda_k \geq 0$, $\lambda_l > 0$, such that $c - u = \sum_{i=1}^k \lambda_i (c_i - u)$.*

Proof. Let $\mathbf{e}_l = (c_l - u) / \|c_l - u\|$. Let $r > 0$ such that $b(c, r) \subseteq \overset{\circ}{\text{coni}}(u, C)$. Then $c - r\mathbf{e}_l \in \overset{\circ}{\text{coni}}(u, C)$ and there are $\mu_1, \dots, \mu_k \geq 0$ such that

$$c - r\mathbf{e}_l - u = \sum_{i=1}^k \mu_i (c_i - u) \iff c - u = \frac{r}{\|c_l - u\|} (c_l - u) + \sum_{i=1}^k \mu_i (c_i - u)$$

We let $\lambda_i = \mu_i$ for $i \neq l$, and $\lambda_l = \mu_l + r / \|c_l - u\|$. It follows that $c - u = \sum_{i=1}^k \lambda_i (c_i - u)$ with $\lambda_i \geq 0$, and $\lambda_l > 0$. □

We can now prove Lemma A.5.

Proof of Lemma A.5. By Lemma A.6, for each $l \in \{1, \dots, k\}$ there are $\lambda_{l1}, \dots, \lambda_{lk} \geq 0$, $\lambda_{ll} > 0$, such that $c - u = \sum_{i=1}^k \lambda_{li}(c_i - u)$. By summing over $1 \leq l \leq k$ we have

$$k(c - u) = \sum_{i=1}^k \left(\sum_{l=1}^k \lambda_{li} \right) (c_i - u)$$

with $\sum_{l=1}^k \lambda_{li}/k > 0$ for all i . □

Remark 73. One can prove the stronger statement

$$\text{relint}(\text{coni}(u, C)) = \left\{ u + \sum_{i=1}^k \lambda_i (c_i - u) \mid \lambda_i \in \mathbb{R}, \lambda_i > 0 \right\}$$

A.3 An alternative way to test the inclusion $(\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B}$

In this appendix, we provide an alternative polynomial-time algorithm to check whether the erosion of a finite collection of disks is contained into the union of another finite collection of disks. The algorithm relies on a characterization of the inclusion of one set into a finite union of balls given by Lemma 4.3. Recall that this lemma provides conditions that guarantee the inclusion of a bounded set A into a finite union of balls X . Specifically, let \mathcal{V} be the vertex set of ∂X . Then we have inclusion $A \subseteq X$ if:

- (i) $\partial A \subseteq X$, and
- (ii) $\forall v \in \mathcal{V}, \exists N_v$ an open neighbourhood of v such that $N_v \cap A \subseteq X$.

We are now ready to describe an algorithm that tests in polynomial time whether $A \subseteq X$, where A is the erosion of a finite collection of disks and X is the union of a finite collection of disks. The pseudocode of the algorithm is given in Algorithm 4. To describe it, let \mathcal{S} and \mathcal{B} be two finite collections of disks, \mathcal{C} a finite collection of circular arcs, x a point, $\varepsilon \geq 0$ a non-negative real number. We introduce four functions:

BOUNDARY takes as input the pair $(\mathcal{S}, \varepsilon)$ and returns the set of circular arcs that compose the boundary of the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$.

POINTLOCATION takes as input the triplet $(x, \mathcal{S}, \varepsilon)$ and returns **inner** if $x \in \overset{\circ}{A}$, **boundary** if $x \in \partial A$ and **outer** if $x \notin A$, where $A = (\bigcup \mathcal{S})^{\ominus \varepsilon}$.

INCLUSION takes as input $(\mathcal{C}, \mathcal{B})$ and returns **true** iff $(\bigcup \mathcal{C}) \subseteq (\bigcup \mathcal{B})$.

LOCALINCLUSION takes as input $(v, \mathcal{S}, \varepsilon, \mathcal{B})$ where v lies on the boundary of both the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$ and the union $\bigcup \mathcal{B}$. It returns **true** iff there exists an open neighborhood N_v of v such that $N_v \cap (\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq (\bigcup \mathcal{B})$.

Each of those four functions can be implemented in polynomial time and we leave to the reader the choice of an implementation.

Algorithm 4 An alternative way to test the inclusion $(\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B}$.

Input: $(\mathcal{S}, \varepsilon, \mathcal{B})$

Output: true if $(\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B}$, false otherwise

```

1: if not INCLUSION(BOUNDARY( $\mathcal{S}, \varepsilon$ ),  $\mathcal{B}$ ) then return false ;
2: for all vertices  $v$  of the boundary of  $\bigcup \mathcal{B}$  do
3:   switch POINTLOCATION( $v, \mathcal{S}, \varepsilon$ ) do
4:     case outer do return false ;
5:     case boundary do
6:       if not LOCALINCLUSION( $v, \mathcal{S}, \varepsilon, \mathcal{B}$ ) then return false ;
7:     end case
8:   end switch
9: end For
10: return true

```

Bibliography

- [ABE09] Dominique Attali, Jean-Daniel Boissonnat, and Herbert Edelsbrunner. “Stability and Computation of Medial Axes - a State-of-the-Art Report”. In: *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*. Ed. by Torsten Möller, Bernd Hamann, and Robert D. Russell. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 109–125 (cit. on pp. 17, 86, 88).
- [AK00] Nina Amenta and Ravi Krishna Kolluri. “Accurate and Efficient Unions of Balls”. In: *Proceedings of the Sixteenth Annual Symposium on Computational Geometry*. SCG '00. ACM, 2000, pp. 119–128 (cit. on pp. 3, 12, 86).
- [AK01] Nina Amenta and Ravi Krishna Kolluri. “The medial axis of a union of balls”. In: *Computational Geometry* 20.1 (2001), pp. 25–37 (cit. on pp. 64–66).
- [AM97] Dominique Attali and Annick Montanvert. “Computing and simplifying 2D and 3D continuous skeletons”. In: *Computer vision and image understanding* 67.3 (1997), pp. 261–273 (cit. on p. 64).
- [AP14] Pankaj K Agarwal and Jiangwei Pan. “Near-linear algorithms for geometric hitting sets and set covers”. In: *Proceedings of the thirtieth annual symposium on Computational geometry*. ACM. 2014, p. 271 (cit. on p. 88).
- [BCS09] Alain Broutta, David Coeurjolly, and Isabelle Sivignon. “Hierarchical discrete medial axis for sphere-tree construction”. In: *International Workshop on Combinatorial Image Analysis*. Springer. 2009, pp. 56–67 (cit. on p. 88).
- [Ber+02] Eric Berberich et al. “A Computational Basis for Conic Arcs and Boolean Operations on Conic Polygons”. In: *Algorithms — ESA 2002*. Ed. by Rolf Möhring and Rajeev Raman. Springer Berlin Heidelberg, 2002, pp. 174–186 (cit. on p. 95).
- [Ber87] Marcel Berger. *Geometry (vols. 1-2)*. 1987 (cit. on pp. 36, 39).
- [BG95] Hervé Brönnimann and Michael T Goodrich. “Almost optimal set covers in finite VC-dimension”. In: *Discrete & Computational Geometry* 14.4 (1995), pp. 463–479 (cit. on p. 88).
- [BO04] Gareth Bradshaw and Carol O’Sullivan. “Adaptive medial-axis approximation for sphere-tree construction”. In: *ACM Transactions on Graphics (TOG)* 23.1 (2004), pp. 1–26 (cit. on pp. 3, 12, 85, 88).
- [BY95] Jean-Daniel Boissonnat and Mariette Yvinec. “Géométrie algorithmique”. In: (1995) (cit. on pp. 8, 18, 22, 104).
- [Cab+09] Sergio Cabello et al. “Maximizing the Area of Overlap of two Unions of Disks under Rigid Motion”. In: *International Journal of Computational Geometry and Applications* 19.06 (2009), pp. 533–556 (cit. on pp. 3, 12, 86).

- [Caz+14] Frédéric Cazals et al. “Greedy Geometric Algorithms for Collection of Balls, with Applications to Geometric Approximation and Molecular Coarse-Graining”. In: *Computer Graphics Forum*. Vol. 33. 6. Wiley Online Lib. 2014, pp. 1–17 (cit. on pp. 3, 12, 85, 88).
- [CCT11] John Chaussard, Michel Couprie, and Hugues Talbot. “Robust skeletonization using the discrete λ -medial axis”. In: *Pattern Recognition Letters* 32.9 (2011), pp. 1384–1394 (cit. on pp. 17, 88).
- [Chv79] Vasek Chvatal. “A greedy heuristic for the set-covering problem”. In: *Mathematics of operations research* 4.3 (1979), pp. 233–235 (cit. on p. 88).
- [CKM99] Tim Culver, John Keyser, and Dinesh Manocha. “Accurate Computation of the Medial Axis of a Polyhedron”. In: *Proceedings of the Fifth ACM Symposium on Solid Modeling and Applications*. SMA '99. ACM, 1999, pp. 179–190 (cit. on pp. 3, 12, 86).
- [CL05] Frédéric Chazal and André Lieutier. “The λ -medial axis”. In: *Graphical Models* 67.4 (2005), pp. 304–331 (cit. on pp. 17, 88).
- [CM07] D. Coeurjolly and A. Montanvert. “Optimal Separable Algorithms to Compute the Reverse Euclidean Distance Transformation and Discrete Medial Axis in Arbitrary Dimension”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29.3 (2007), pp. 437–448 (cit. on pp. 3, 12, 86).
- [Cor+01] Thomas H. Cormen et al. *Introduction to Algorithms*. 2nd. McGraw-Hill Higher Education, 2001 (cit. on p. 145).
- [CP95] Tiziana Calamoneri and Rossella Petreschi. “An efficient orthogonal grid drawing algorithm for cubic graphs”. In: *Computing and Combinatorics*. Springer, 1995, pp. 31–40 (cit. on p. 99).
- [DMT92] Olivier Devillers, Stefan Meiser, and Monique Teillaud. “The space of spheres, a geometric tool to unify duality results on Voronoi diagrams”. PhD thesis. INRIA, 1992 (cit. on p. 22).
- [DZ04] Tamal K. Dey and Wulue Zhao. “Approximate medial axis as a Voronoi subcomplex”. In: *Computer-Aided Design* 36.2 (2004), pp. 195–202 (cit. on pp. 17, 88).
- [Ede92] Herbert Edelsbrunner. *Weighted alpha shapes*. University of Illinois at Urbana-Champaign, Department of Computer Science, 1992 (cit. on p. 64).
- [Ede93] Herbert Edelsbrunner. “The union of balls and its dual shape”. In: *Proceedings of the ninth annual symposium on Computational geometry*. ACM. 1993, pp. 218–231 (cit. on pp. 65, 107).
- [EK05] Herbert Edelsbrunner and Patrice Koehl. “The geometry of biomolecular solvation”. In: *Combinatorial and computational geometry* 52 (2005), pp. 243–275 (cit. on pp. 3, 12, 85).
- [Fed59] H. Federer. “Curvature measures”. In: *Trans. Amer. Math. Soc.* 93 (1959), pp. 418–491 (cit. on pp. 7, 86).

- [Fei+15] Julian Feinauer et al. “Structural characterization of particle systems using spherical harmonics”. In: *Materials Characterization* 106 (2015), pp. 123–133 (cit. on pp. 3, 12, 86).
- [FHK02] Gerald E Farin, Josef Hoschek, and Myung-Soo Kim. *Handbook of computer aided geometric design*. Elsevier, 2002 (cit. on pp. 1, 11).
- [GJ77] Michael R Garey and David S. Johnson. “The rectilinear Steiner tree problem is NP-complete”. In: *SIAM Journal on Applied Mathematics* 32.4 (1977), pp. 826–834 (cit. on p. 95).
- [Hah01] Stefanie Hahmann. “Geometric modeling of smooth surfaces: Design and Fairing”. date de rédaction: Juin 2001. Habilitation à diriger des recherches. Institut National Polytechnique de Grenoble - INPG, Dec. 2001 (cit. on pp. 1, 11).
- [HP11] Sariel Har-Peled. *Geometric approximation algorithms*. Vol. 173. American mathematical society Boston, 2011 (cit. on p. 88).
- [Hub95] P. M. Hubbard. “Collision detection for interactive graphics applications”. In: *IEEE Transactions on Visualization and Computer Graphics* 1.3 (1995), pp. 218–230 (cit. on p. 88).
- [Hub96] Philip M Hubbard. “Approximating polyhedra with spheres for time-critical collision detection”. In: *ACM Transactions on Graphics (TOG)* 15.3 (1996), pp. 179–210 (cit. on pp. 3, 12, 86).
- [Joh74] David S Johnson. “Approximation algorithms for combinatorial problems”. In: *Journal of computer and system sciences* 9.3 (1974), pp. 256–278 (cit. on p. 88).
- [Kar72] Richard M Karp. “Reducibility among combinatorial problems”. In: *Complexity of computer computations*. Springer, 1972, pp. 85–103 (cit. on pp. 9, 147).
- [Ker39] Richard Kershner. “The Number of Circles Covering a Set”. In: *American Journal of Mathematics* 61.3 (1939), pp. 665–671 (cit. on p. 88).
- [Lie04] André Lieutier. “Any open bounded subset of \mathbb{R}^n has the same homotopy type as its medial axis”. In: *Computer-Aided Design* 36.11 (2004), pp. 1029–1046 (cit. on pp. 8, 103).
- [Mat88] Georges Matheron. “Examples of topological properties of skeletons”. In: *Image analysis and mathematical morphology* 2 (1988), pp. 217–238 (cit. on pp. 18, 65).
- [MGP10] Balint Miklos, Joachim Giesen, and Mark Pauly. “Discrete scale axis representations for 3D geometry”. In: *ACM Transactions on Graphics (TOG)*. Vol. 29. 4. ACM. 2010, p. 101 (cit. on pp. 3, 12, 88).
- [Pac12] J. Pach. *New Trends in Discrete and Computational Geometry*. Algorithms and Combinatorics. Springer Berlin Heidelberg, 2012 (cit. on p. 93).
- [Req77] Aristides Requicha. “Mathematical models of rigid solid objects”. In: (1977) (cit. on pp. 2, 11).
- [RF96] V. Ranjan and A. Fournier. “Matching and Interpolation of Shapes using Unions of Circles”. In: *Computer Graphics Forum* 15.3 (1996), pp. 129–142 (cit. on pp. 3, 12, 86).

- [RMS13] Pedro J. de Rezende, Flavio K. Miyazawa, and Anderson T. Sasaki. “A {PTAS} for the disk cover problem of geometric objects”. In: *Operations Research Letters* 41.5 (2013), pp. 552–555 (cit. on p. 88).
- [RR92] Aristides AG Requicha and Jarek R Rossignac. “Solid modeling and beyond”. In: *IEEE Computer Graphics and Applications* 12.5 (1992), pp. 31–44 (cit. on pp. 2, 11).
- [Sch79] Hans Schwerdtfeger. *Geometry of complex numbers: circle geometry, Moebius transformation, non-euclidean geometry*. Courier Corporation, 1979 (cit. on p. 22).
- [Ser82] Jean Serra. *Image analysis and mathematical morphology, v. 1*. Academic press, 1982 (cit. on pp. 6, 86).
- [T04] G. Fejes Tóth. “Handbook of Discrete and Computational Geometry”. In: ed. by Joseph O’Rourke Jacob E. Goodman. Chapman & Hall, CRC, 2004. Chap. Packing and covering (cit. on pp. 8, 88, 104).
- [Wei+] Ron Wein et al. *2D Arrangements*. CGAL: Computational Geometry Algorithms Library (cit. on p. 95).
- [WZ06] Ron Wein and Baruch Zukerman. *Exact and efficient construction of planar arrangements of circular arcs and line segments with applications*. Tech. rep. Tech. Rep. ACS-TR-121200-01, Tel Aviv University, Israel, 2006 (cit. on p. 93).
- [WZ09] Rene Weller and Gabriel Zachmann. “Inner sphere trees for proximity and penetration queries.” In: *Robotics: science and systems*. Vol. 2. 2009 (cit. on pp. 3, 12).

Résumé — Représenter un objet géométrique complexe par un ensemble de primitives simples est une tâche souvent fondamentale, que ce soit pour la reconstruction et la réparation de données, ou encore pour faciliter la visualisation ou la manipulation des données. Le choix de la ou les primitives, ainsi que celui de la méthode d'approximation, impactent fortement les propriétés de la représentation de forme qui sera obtenue.

Dans cette thèse, nous utilisons les boules comme seule primitive. Nous prenons ainsi un grand soin à décrire les unions finies de boules et leur structure. Pour cela, nous nous reposons sur les faisceaux de boules. En particulier, nous aboutissons à une description valide en toute dimension, sans hypothèse de position générale. En chemin, nous obtenons également plusieurs résultats portant sur les tests d'inclusion locale et globale dans une union de boules.

Nous proposons également une nouvelle méthode d'approximation par union finie de boules, l'approximation par boules à (δ, ϵ) -près. Cette approche contraint l'union de boules à couvrir un sous-ensemble de la forme d'origine (précisément, un ϵ -érosion), tout en étant contenu dans un sur-ensemble de la forme (un δ -dilaté). En nous appuyant sur nos précédents résultats portant sur les unions de boules, nous démontrons plusieurs propriétés de ces approximations. Nous verrons ainsi que calculer une approximation par boules à (δ, ϵ) -près qui soit de cardinal minimum est un problème NP-complet. Pour des formes simples dans le plan, nous présentons un algorithme polynomial en temps et en espace qui permet de calculer ces approximations de cardinal minimum. Nous concluons par une généralisation de notre méthode d'approximation pour une plus large variété de sous-ensembles et sur-ensembles.

Mots clés : géométrie algorithmique, géométrie discrète, approximation de forme, axe médian, union finie de boules, faisceau de sphères/boules

Abstract — Describing a complex geometric shape with a set of simple primitives is often a fundamental task for shape reconstruction, visualization, analysis and manipulation. The type of primitives, as well as the choice of approximation scheme, both greatly impact the properties of the resulting shape representation.

In this PhD, we focus on balls as primitives. Using pencils of balls, we carefully describe finite unions of balls and their structure. In particular, our description holds in all dimension without assuming general position. On our way, we also establish various results and tools to test local and global inclusions within these unions.

We also propose a new approximation scheme by union of balls, the (δ, ϵ) -ball approximation. This scheme constrains the approximation to cover a core subset of the original shape (specifically, an ϵ -erosion), while being contained within a superset of the shape (a δ -dilatation). Using our earlier results regarding finite unions of balls, we prove several properties of these approximations. We show that computing a cardinal minimum (δ, ϵ) -ball approximation is an NP-complete problem. For simple planar shapes however, we present a polynomial time and space algorithm that outputs a cardinal minimum approximation. We then conclude by generalizing the approximation scheme to a wider range of core subsets and bounding supersets.

Keywords: computational geometry, digital geometry, shape approximation, medial axis, finite union of balls, pencil of spheres/balls
