Algebraic cryptanalysis of the shortest vector problem in ideal lattices

Olivier Bernard

DOCTORAT BRETAGNE LOIRE / MATHSTIC

UNIVERSITÉ DE RENNES 1

THÈSE DE DOCTORAT DE

L'UNIVERSITÉ DE RENNES 1

ÉCOLE DOCTORALE N° 601
Mathématiques et Sciences et Technologies de l'Information et de la Communication
Spécialité : *Informatique*

Par
Olivier BERNARD

Algebraic Cryptanalysis of the Shortest Vector Problem in Ideal Lattices

Thèse présentée et soutenue à Rennes, le 13 juin 2022
Unité de recherche : IRISA, UMR 6074

Rapporteurs avant soutenance :

Guillaume HANROT Professeur, ENS Lyon, LIP, France
Emmanuel THOMÉ Directeur de Recherche, INRIA, LORIA, Nancy, France

Composition du Jury :

Examinateurs :	Gildas AVOINE	Professeur, INSA Rennes, IRISA, France
	Léo DUCAS	Tenured Researcher, CWI, Cryptology Group, Amsterdam, Pays-Bas
	Guillaume HANROT	Professeur, ENS Lyon, LIP, France
	Emmanuel THOMÉ	Directeur de Recherche, INRIA, LORIA, Nancy, France
	Brigitte VALLÉE	Directrice de Recherche Émérite, CNRS, GREYC, Caen, France
Dir. de thèse :	Pierre-Alain FOUQUE	Professeur, Université de Rennes 1, IRISA, France
	Adeline ROUX-LANGLOIS	Chargée de Recherche HDR, CNRS, IRISA, Rennes, France

Acknowledgements

Successful works always require the support of many people, and this is especially true for completing a PhD thesis, a particularly heavy life project. My first thanks go to my supervisors Pierre-Alain Fouque and Adeline Roux-Langlois, who welcomed me in Rennes at the IRISA laboratory, and without whom this journey would not have been possible. In particular, I really appreciate all the time Adeline dedicated to me, as well as the freedom she gave me in my research topics and numerous discussions that helped me to obtain a comprehensive view of academic research. This epic submission at Asiacrypt at nearly 6am will certainly go down in history! My gratitude also goes to Pierre-Alain for his unfailing kind help and support. I will remember each of our fruitful discussions as clearly as I remember you trying to convince me to follow this PhD track almost 11 years ago when mentoring me at the Éns Paris. Nothing would also have been possible without the unwavering support of Éric Garrido, head of the *CHiffre Laboratory* at Thales, who sold the financial and work arrangement upstairs.[1]

Secondly, I am very grateful to Guillaume Hanrot and Emmanuel Thomé, who thoroughly reviewed this work and wrote very encouraging and heartwarming reports. Likewise, many thanks go to Gildas Avoine, Léo Ducas and Brigitte Vallée who kindly agreed to participate to my PhD defense jury. Anonymous proofreaders of the name of Alexandre W., Thomas R. and Simon A. also significantly helped "off-the-cuff" to improve the quality of various parts of this manuscript through many insightful comments, and I am very thankful for their responsiveness.

My co-authors moreover played a crucial role. First, I am deeply indebted to Radan Kučera for thorough and invaluable discussions about the Stickelberger ideal, and for his very accurate and patient supervision. I am also thankful to Andrea Lesavourey and Tuong-Huy Nguyen, whose enthusiasm in a very tense and competitive context was considerably helpful. A special thought is addressed to Thomas Ricosset for invaluable discussions about the geometry of lattices — and life, as well as to Alice Pellet-Mary and Damien Stehlé who warmly welcomed me at the LIP, ENS Lyon for six weeks, and later at the LIFANT, Bordeaux.

Furthermore, I had the chance to evolve in a very fertile research environment, both at IRISA at Rennes and at the *CHiffre Laboratory* at Thales, where many remarkable people helped me through this adventure: Katharina for sharing her office and taking care of Bada55,[2] Solène for beautifully playing with me Brahms' first violin sonata at the piano, not to mention all the Corgis lovers of the SPICY/CAPSULE teams, and of course all the Cryptosaurs™ at Thales, whose jolly company make work such a pleasant place. Even though I always was between two suitcases catching a train, I enjoyed every moment with all of you.

My last thoughts, but not least, are leaned towards all my close friends, musicians, bridge players, former colleagues, for their precious support, and especially as the glaring Covid-19

[1] *In fine*, this PhD thesis has been largely financed thanks to the European Union PROMETHEUS project (Horizon 2020 Research and Innovation Program, grant 780701).

[2] My PhD thesis mascot, a now-vegetarian Cryptosaur named after a maliciously-designed elliptic curve family, now trying to survive the quantum apocalypse.

pandemic situation made everything more chaotic than necessary. Thanks to you all for putting up with me during this 3-years fulltime-absorbing task, and my most sincere apologies to the ones that were lost on track. I hope we will catch up again! Needless to say, I am immensely grateful for the exceptional love stars[3] that successively stood by me, pushed me and made me happy through thick and thin.

As a final wink, it is impossible for me not to deliver a tribute to all my plush animals, whose unconditional support even in the harshest moments was very comforting.[4]

Rennes, June 2022

[3] Alas, I was not able to refute LCH's infamous statistics.
[4] Now, there will be plenty of time for you to shower.

Contents

Résumé en français

Poussée par la menace hypothétique de la construction dans les prochaines décennies d'un ordinateur quantique à grande échelle, la communauté cryptographique a été amenée à considérer de nouveaux problèmes mathématiques sur lesquels fonder la sécurité des cryptosystèmes à clé publique dits *post-quantiques*. En 2016, l'agence américaine *National Institute of Standards and Technology* (NIST) a lancé une compétition de standardisation pour la cryptographie post-quantique, afin d'évaluer et standardiser des algorithmes à clé publiques résistants à l'ordinateur quantique. Pas loin de 70 propositions ont été reçues, utilisant plusieurs objets mathématiques comme, pour n'en nommer que quelques-uns, les réseaux euclidiens, les codes correcteurs d'erreurs ou les graphes d'isogénies entre courbes elliptiques supersingulières.

La famille des réseaux euclidiens, qui fait l'objet d'un grand nombre de soumissions, représente l'une des solutions post-quantiques les plus prometteuses. Plusieurs problèmes difficiles sont utilisés afin de prouver la sécurité de ces cryptosystèmes, comme le problème Ntru [HPS98], le problème *Short Integer Solution* (Sis) [Ajt96] ou le problème *Learning With Errors* (Lwe) [Reg05], ainsi que leurs variantes algébriquement structurées *ring* (Ring-Sis [LM06, PR06], Ring-Lwe [SSTX09, LPR10]) ou *Module* (Module-Sis, Module-Lwe [LS15]). Typiquement, les variantes algébriquement structurées ont l'avantage d'offrir de meilleures performances, au prix d'une possible perte de sécurité. En fin de compte, leur sécurité repose sur la difficulté de résoudre le problème du presque plus court vecteur, ou *Approximate Shortest Vector Problem* (Approx-Svp), dans la classe réduite des réseaux euclidiens algébriquement structurés.

Dans le cas de réseaux arbitraires, Svp est un problème NP-difficile [Ajt98] extensivement étudié. Sa version approchée consiste, pour tout réseau de rang n, à trouver un vecteur non nul du réseau dont la norme euclidienne diffère d'un petit facteur multiplicatif de la longueur du plus court vecteur non nul du réseau. Le meilleur compromis dans ce cas est donné par la hiérarchie de Schnorr [Sch87], qui permet d'atteindre un facteur d'approximation $2^{\widetilde{O}(n^\omega)}$ en temps $2^{\widetilde{O}(n^{1-\omega})}$, pour tout $\omega \in (0,1)$. En pratique, le meilleur algorithme connu qui est proche de ce compromis est l'algorithme *Block Korkin-Zolotarev* (BKZ) [SE94], qui peut être vu comme une amélioration de l'algorithme bien connu LLL [LLL82], dû à A. Lenstra, H. Lenstra et L. Lovász.

Cependant, ces hypothèses structurées (p. ex., Ring-Lwe) pourraient se révéler triviales si les variantes sous-jacentes d'Approx-Svp s'avéraient plus faciles dans le cas spécifique des réseaux algébriquement structurés. Ainsi, une cible naturelle pour la cryptanalyse est le problème du plus court vecteur dans les *réseaux idéaux*, ou *Ideal Shortest Vector Problem* (id-Svp), c.-à-d., restreint aux réseaux images par le plongement de Minkowski d'idéaux fractionnaires de l'anneau des entiers $\mathcal{O}_K$ d'un corps de nombres K. Pendant une longue période, le meilleur algorithme connu pour résoudre Approx-Svp dans les réseaux idéaux a été le même que pour les réseaux non structurés. Cependant, une série récente de travaux [CGS14, EHKS14, BS16, CDPR16, CDW17, DPW19, PHS19a] tend à montrer que la résolution de ce problème pourrait se révéler plus facile dans les réseaux idéaux, en particulier dans un monde quantique.

Algorithmes quantiques pour la théorie des nombres

En effet, la découverte de nouveaux algorithmes *quantiques* en temps polynomial pour la théorie des nombres a attiré de plus en plus d'attention sur la manière dont la forte structure algébrique de ces réseaux idéaux pourrait être utilisée pour s'attaquer à id-SVP plus efficacement que par le truchement des algorithmes traditionnels de réduction de réseaux comme LLL ou BKZ.

Tout a commencé avec la note de CAMPBELL, GROVES et SHEPHERD [CGS14], qui a fait grand bruit et a revendiqué, toutefois sans preuves, une attaque en temps polynomial quantique contre un schéma nommé Soliloquy, qui résout des instances très spécifiques d'Approx-SVP sur des réseaux idéaux *principaux*. Leur algorithme comporte deux étapes successives :

- la première résout le problème de l'idéal principal, ou *Principal Ideal Problem* (PIP), qui consiste à trouver n'importe quel générateur d'un idéal principal,
- la deuxième réduit ce générateur autant que possible, grâce aux unités algébriques du corps de nombres, ce qui revient à résoudre un problème du plus proche vecteur, ou *Closest Vector Problem* (CVP) dans le réseau log-unité.

Les auteurs affirment que la première étape peut être effectuée en temps quantique polynomial, et que la deuxième étape est suffisamment facile pour permettre de casser le schéma dans le cas des corps cyclotomiques, grâce aux *unités circulaires*.

La première partie de leurs revendications a été prouvée indépendamment dans [EHKS14], qui décrit une généralisation de l'algorithme de Shor [Sho97], pour calculer le groupe des unités de corps de nombres de degrés arbitraires en temps quantique polynomial. Plus tard, en se basant sur [EHKS14], BIASSE et SONG [BS16] ont étendu ce résultat au calcul du groupe des classes et des $\mathcal{S}$-unités. Plus précisément, ils montrent comment calculer des $\mathcal{S}$-unités, une généralisation des unités algébriques d'un corps de nombres dépendant d'un ensemble $\mathcal{S}$ d'idéaux premiers, en temps quantique polynomial en la taille du discriminant Δ_K du corps de nombres K, et en la taille de la *base de facteurs* $\mathcal{S}$. Ils montrent également comment la résolution du PIP, ainsi que le calcul du groupe des classes ou du groupe des unités, peuvent être réduits à ces calculs de $\mathcal{S}$-unités pour des bases de facteurs $\mathcal{S}$ convenablement choisies.

Cryptanalyses algébriques de id-SVP

En ce qui concerne la seconde revendication de [CGS14], CRAMER, DUCAS, PEIKERT et REGEV [CDPR16] ont prouvé que, dans le cas des corps cyclotomiques de conducteur égal à une puissance d'un nombre premier, les plongements logarithmiques de l'ensemble des unités circulaires [Was97, §8] induisent une base de suffisamment bonne qualité d'un sous-réseau d'indice relativement petit dans le réseau log-unité. Cette propriété fondamentale leur permet de conclure qu'il existe un algorithme quantique polynomial qui, en moyenne, résout Approx-SVP dans des réseaux idéaux principaux pour un facteur d'approximation $2^{\widetilde{O}(\sqrt{n})}$, où n est la dimension de l'idéal.

Cette première cryptanalyse algébrique a ensuite donné lieu à plusieurs généralisations, à toute classe d'idéaux fractionnaires [CDW17], tous corps cyclotomiques [CDW21] et tout corps de nombres [PHS19a]. Pour tout idéal challenge $\mathfrak{b}$ d'un corps de nombres K, toutes ces approches partent d'une solution au problème du logarithme discret dans le groupe des classes, ou *Class Group Discrete Logarithm Problem* (CLDLP). Ce problème de représentation consiste, à partir d'un ensemble fixé de places finies correspondant à des idéaux premiers $\{\mathfrak{p}_1, \ldots, \mathfrak{p}_k\}$ de K, à trouver, s'ils existent, $\alpha \in K$ et $e_1, \ldots, e_k \in \mathbb{Z}$ tels que :

$$\langle \alpha \rangle = \mathfrak{b} \cdot \prod_{1 \leq i \leq k} \mathfrak{p}_i^{e_i}.$$

Ce problème revient à calculer un certain groupe de $\mathcal{S}$-unités, ce qui selon la discussion précédente

peut donc être résolu facilement dans un monde quantique. Ainsi, la partie la plus délicate de ces cryptanalyses réside dans le fait de réduire la norme euclidienne de α, c.-à-d., de trouver une *plus courte* solution au ClDlp ci-dessus, ou de manière équivalente, le plus court représentant de la classe de α modulo le groupe multiplicatif H finiment engendré par les relations de classes entre les $\mathfrak{p}_i$'s. En utilisant un plongement logarithmique adapté, cela revient à résoudre une instance Cvp dans le réseau image de H par ce plongement logarithmique. À la fin, l'espoir est que ce plus court représentant soit un élément suffisamment petit de l'idéal challenge. Par conséquent, il est particulièrement important de choisir soigneusement le plongement logarithmique, de telle sorte qu'il convoie toutes les informations utiles sur la taille de α, et de sorte que la base du réseau obtenu ne soit pas de trop mauvaise qualité pour l'oracle Cvp. Remarquons que l'algorithme de [CGS14, CDPR16] suit exactement cette procédure pour $k = 0$, auquel cas une solution au ClDlp existe si et seulement si $\mathfrak{b}$ est principal.

Ces cryptanalyses algébriques peuvent se regrouper en deux lignes de travaux, qui utilisent des outils différents pour estimer et garantir la taille de leurs sorties, et n'ont pas la même portée :

- L'algorithme CDW, par Cramer, Ducas et Wesolowski [CDW17, CDW21], résout id-Svp pour un facteur d'approximation $\exp \widetilde{O}(\sqrt{n})$ dans les corps cyclotomiques de degrés n, en temps *quantique* polynomial. Ce compromis est prouvé à l'aide d'heuristiques soigneusement justifiées. L'algorithme utilise l'idéal de Stickelberger d'un corps cyclotomique, un idéal spécial qui fournit gratuitement des relations *courtes* dans la partie relative du groupe des classes. Ces relations courtes permettent de trouver un *proche* multiple principal de tout idéal challenge, c.-à-d., un multiple principal dont la norme algébrique, divisée par la norme de l'idéal challenge, est relativement petite. De là, la routine de [CDPR16] peut être appliquée à un générateur de ce multiple principal, dans l'espoir que sa sortie soit suffisamment petite.
- Ces deux étapes peuvent en fait être combinées dans une unique instance Cvp, ce qui a donné naissance à ce qui est maintenant appelé les *attaques par $\mathcal{S}$-unités* : l'idée est de trouver de cette manière un multiple principal qui n'est pas seulement de petite norme algébrique, mais qui est également généré par un *petit* élément. C'est l'idée centrale de l'algorithme PHS par Pellet-Mary, Hanrot and Stehlé [PHS19a], qui s'applique à tout corps de nombres et que nous détaillons dans la section suivante.

En ce qui concerne l'algorithme CDW, son impact en pratique a été évalué dans [DPW19] grâce à de nombreuses simulations pour la résolution du Cvp dans chacun des deux réseaux impliqués. À partir de ces résultats expérimentaux, les auteurs dérivent une *borne inférieure volumétrique* [DPW19, Eq. (5) et Tab. 1] et en concluent que l'algorithme CDW devrait battre BKZ_{300} pour des corps cyclotomiques de degrés plus grands que 7000.[5]

Attaques par $\mathcal{S}$-unités

Nous décrivons plus en détails l'algorithme PHS [PHS19a], d'après Pellet-Mary, Hanrot et Stehlé, qui est à notre connaissance la première attaque par $\mathcal{S}$-unités décrite et prouvée dans la littérature, même si le formalisme des $\mathcal{S}$-unités n'est pas directement utilisé dans [PHS19a].

La principale caractéristique de leur algorithme est de combiner dans une unique instance Cvp les deux étapes principales de l'algorithme CDW [CDW17, CDW21], plus précisément le problème du proche multiple principal ou *Close Principal Multiple Problem* (Cpmp) d'une part, et le problème du plus court générateur, ou *Shortest Generator Problem* (Sgp) d'autre part. Ceci garantit dans une certaine mesure que la sortie de l'algorithme de résolution du Cpmp

[5] La première version publiée de [DPW19] présentait un point de rencontre au degré 12000. Après la correction d'une erreur d'implémentation, découverte par Bernstein et rendue publique le 20 Août 2021 dans une présentation à la conférence SIAM, ce point de rencontre a été réévalué à 7000 [DPW19, Fig. 5].

possède un générateur qui n'est "*pas beaucoup plus grand*" que son plus court élément non nul. Malheureusement, cela n'est rendu possible qu'au prix d'un précalcul exponentiel, dépendant uniquement du corps de nombres K. En effet, afin de garantir la taille de la sortie et la complexité temporelle de l'algorithme, un ingrédient fondamental de la preuve réside dans l'utilisation d'un oracle CVP avec données de précalcul dû à LAARHOVEN [Laa16]. Plus formellement, l'algorithm PHS se divise en deux phases :

1. La phase de précalcul construit un réseau spécifique, ne dépendant que du corps K, qui peut être vu comme un réseau log-$\mathcal{S}$-unité sous un plongement logarithmique particulier, ainsi qu'une donnée permettant de résoudre efficacement Approx-CVP dans ce réseau. En notant Δ_K le discriminant de K, cette phase s'exécute en temps $2^{\widetilde{O}(\log|\Delta_K|)}$ et produit une donnée $\mathcal{V}$ de taille $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$, où $\omega \in \left[0, \frac{1}{2}\right]$ paramétrise le compromis entre le temps d'exécution et le facteur d'approximation obtenu par la phase suivante.
2. La phase de requête réduit chaque challenge pour Approx-id-SVP à la résolution d'une instance Approx-CVP dans ce réseau log-$\mathcal{S}$-unité fixé. Elle prend en entrée n'importe quel idéal de $\mathcal{O}_K$, dont la norme algébrique est de taille bornée par $2^{\mathrm{poly}(\log|\Delta_K|)}$, ainsi que la donnée précalculée $\mathcal{V}$, et s'exécute en temps $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)} + \mathrm{T}_{\mathsf{Su}}(K)$. La sortie est un élément non nul de l'idéal qui est une solution d'Approx-SVP pour un facteur d'approximation $2^{\widetilde{O}(\log^{\omega+1}|\Delta_K|/n)}$, où n est le degré de K.

Ici, $\mathrm{T}_{\mathsf{Su}}(K)$ désigne le temps d'exécution des calculs de groupes de $\mathcal{S}$-unités, c.-à-d., dans un monde quantique, $\mathrm{T}_{\mathsf{Su}}(K) = \widetilde{O}\big(\ln|\Delta_K|\big)$ est polynomial [BS16], tandis que dans un monde classique, $\mathrm{T}_{\mathsf{Su}}(K)$ reste sous-exponentiel en $\ln|\Delta_K|$, soit $\mathrm{T}_{\mathsf{Su}}(K) = \exp\widetilde{O}(\ln^{\alpha}|\Delta_K|)$, où $\alpha = 1/2$ pour les corps cyclotomiques [BEF+17],[6] et $\alpha = 2/3$ dans le cas général [BF14], récemment réduit à $\alpha = 3/5$ par GÉLIN [Gél17].

En omettant le coût exponentiel du précalcul, la phase de requête bat la traditionnelle hiérarchie de Schnorr [Sch87] quand $\log|\Delta_K| \leq \widetilde{O}(n^{1+\varepsilon})$ avec $\varepsilon = 1/3$ dans le cas quantique, et $\varepsilon = 1/11$ dans le cas classique [PHS19a, Fig. 5.3]. Cependant, ces bornes sur le discriminant ne sont pas homogènes quand le facteur d'approximation varie, c.-à-d., pour un facteur d'approximation fixé à $2^{\sqrt{n}}$, la complexité temporelle de l'algorithme PHS bat asymptotiquement la hiérarchie de Schnorr uniquement dans le cas quantique et uniquement pour $\varepsilon \leq 1/6$.

Contributions de cette thèse

Les contributions de cette thèse se placent dans le contexte des attaques par $\mathcal{S}$-unités. Tout d'abord, nous utilisons le formalisme des $\mathcal{S}$-unités pour définir l'algorithme Twisted-PHS, une version pondérée de l'algorithme PHS qui se révèle extrêmement puissante en pratique. Puis, les contributions suivantes font, pour tous les corps cyclotomiques, la jonction entre les deux lignes décrites précédemment de travaux de cryptanalyse : en utilisant des techniques avancées portant sur le réseau de Stickelberger, nous supprimons d'une part des étapes quantiques de l'algorithme CDW, et d'autre part approchons expérimentalement l'algorithme Twisted-PHS en moyenne dimension, où les phénomènes asymptotiques commencent à s'exprimer pleinement.

Twisted-PHS : utilisation de la Formule du Produit

En fait, le réseau particulier utilisé dans l'algorithme PHS correspond à un réseau spécial appelé le réseau log-$\mathcal{S}$-unité, c.-à-d., un réseau obtenu à partir des images de $\mathcal{S}$-unités par un plongement

[6] Pour des raisons historiques, l'article [BEF+17] est écrit spécifiquement pour les corps cyclotomiques de conducteurs une puissance de nombre premier, mais s'adapte directement au cas général pour le calcul des groupes de $\mathcal{S}$-unités.

logarithmique adapté, où $\mathcal{S}$ peut être identifié à une *base de facteurs* FB d'idéaux premiers. Il s'avère que choisir soigneusement le plongement logarithmique utilisé est particulièrement important en pratique.

Ainsi, notre première contribution consiste à proposer une nouvelle version *tordue* de l'algorithme PHS, dénommée *Twisted-PHS*, dont l'idée principale consiste à identifier un plongement logarithmique préservant les propriétés algébriques naturelles des $\mathcal{S}$-unités. Plus précisément, nous incluons les poids standards de la théorie des nombres induits par la *Formule du Produit* aux coordonnées du plongement logarithmique. Ainsi, pour tout α du corps de nombres K, nous partons de la formule suivante :

$$\mathrm{Log}_{\mathcal{S}}\,\alpha = \Big(\big\{[K_\sigma : \mathbb{R}] \cdot \ln|\sigma(\alpha)|\big\}_{\sigma\in\mathcal{S}_\infty}, \big\{-v_{\mathfrak{p}}(\alpha)\cdot \ln \mathcal{N}(\mathfrak{p})\big\}_{\mathfrak{p}\in \mathrm{FB}}\Big),$$

où $K_\sigma = \mathbb{R}$ (resp. $\mathbb{C}$) pour tout plongement $\sigma \in \mathcal{S}_\infty$ réel (resp. complexe) allant de K dans $\mathbb{R}$ (resp. $\mathbb{C}$), et pour tout idéal premier $\mathfrak{p} \in \mathrm{FB}$, $\mathcal{N}(\mathfrak{p})$ désigne sa norme algébrique et $v_{\mathfrak{p}}(\alpha)$ désigne la valuation de α en $\mathfrak{p}$. Par contraste, le plongement logarithmique sur lequel se base [PHS19a] n'inclut pas les poids $\ln \mathcal{N}(\mathfrak{p})$ sur les dernières coordonnées.

En utilisant ce formalisme des $\mathcal{S}$-unités, nous prouvons que notre algorithme Twisted-PHS réalise le même compromis temps d'exécution v.s. facteur d'approximation que l'algorithme PHS, grâce au même oracle Cvp avec précalcul dû à Laarhoven [Laa16] pour résoudre efficacement les instances Approx-Cvp dans le réseau log-$\mathcal{S}$-unité. À titre de contribution secondaire, nous proposons également plusieurs améliorations de l'algorithme PHS.

Intuitivement, le fait d'ajouter des poids $\ln \mathcal{N}(\mathfrak{p})$ aux valuations entières pour chaque idéal premier $\mathfrak{p}$ capture l'idée qu'utiliser une relation augmentant les valuations pour un idéal de grande norme est plus coûteux qu'utiliser une relation impliquant des idéaux de plus petite norme. Ceci encode également dans le réseau log-$\mathcal{S}$-unité l'information sur la longueur et la norme algébrique des $\mathcal{S}$-unités, contrairement au plongement logarithmique utilisé dans [PHS19a] n'impliquant que les valuations entières. *In fine*, ces éléments tendent à indiquer que l'oracle Cvp dans le réseau log-$\mathcal{S}$-unité *tordu* combine plus efficacement l'objectif de chercher un idéal multiple principal de petite norme algébrique tout en minimisant la longueur de son générateur.

Une autre conséquence fondamentale de l'utilisation d'un plongement logarithmique convenablement normalisé tient à ce que nous appelons le *phénomène de base de facteurs optimale*, c.-à-d., nous prouvons qu'il existe une base de facteurs $\mathcal{S}$ pour laquelle la densité du réseau log-$\mathcal{S}$-unité est maximale, et donnons un algorithme pour la calculer.

Sur le plan pratique, nous fournissons une implémentation de bout en bout de l'algorithme Twisted-PHS, où l'oracle Cvp de Laarhoven est remplacé par l'algorithme Nearest Plane de Babai [Bab86]. Cette implémentation est publiquement disponible sur GitHub: ob3rnard/Twisted-PHS[7]. Pour la première fois, ceci a permis d'exécuter complètement des attaques par $\mathcal{S}$-unités sur une palette significative d'exemples concrets. Les résultats de nos expériences suggèrent, pour des corps cyclotomiques de conducteurs premiers et des corps NTRU Prime de petites dimensions, plus précisément jusqu'en dimension 70, que :

- avec la normalisation standard de la théorie de nombres, les réseaux log-$\mathcal{S}$-unités présentent des caractéristiques géométriques très particulières et semblent extrêmement faciles à réduire avec BKZ ;
- les facteurs d'approximation *exacts* obtenus sont particulièrement petits et croissent très lentement avec la dimension, "*de manière potentiellement sous-exponentielle ou même meilleure*".

À notre connaissance, il s'agit des toutes premières preuves expérimentales de la particularité géométrique des réseaux log-$\mathcal{S}$-unités tordus ainsi que du potentiel des attaques par $\mathcal{S}$-unités en

[7] https://github.com/ob3rnard/Twisted-PHS

pratique. Malheureusement, à cause de la complexité du calcul des $\mathcal{S}$-unités dans un monde classique, les dimensions atteintes ne permettent pas de conjecturer concrètement le comportement asymptotique de l'algorithme Twisted-PHS.

Une base courte de l'idéal de Stickelberger

Dans la contribution suivante, nous verrons que calculer explicitement les générateurs de Stickelberger est utile dans au moins deux situations : la première intervient pour supprimer la dernière étape quantique dans l'algorithme CDW [CDW21], la seconde intervient pour approcher le réseau log-$\mathcal{S}$-unité utilisé dans l'algorithme Twisted-PHS. Dans cette dernière situation, certaines étapes de calcul, en particulier la procédure de 2-*saturation* utilisée pour densifier le réseau, deviennent rapidement irréalisables avec la croissance des coefficients des éléments. Ceci incite à contraindre à la fois le nombre et la taille des générateurs de Stickelberger impliqués.

C'est ici que notre seconde contribution s'avère particulièrement utile. Notre résultat principal consiste à décrire pour la première fois une *base courte* explicite de l'idéal de Stickelberger $\mathcal{S}_m$ du m-ième corps cyclotomique pour *tout* conducteur m, c.-à-d., une base qui n'est constituée *que* d'éléments courts. Par définition, un élément de $\mathbb{Z}[G_m]$, où G_m désigne le groupe de Galois du m-ième corps cyclotomique, est dit court s'il s'écrit sous la forme :

$$\sum_{\sigma \in G_m} \varepsilon_\sigma \cdot \sigma \in \mathcal{S}_m \subset \mathbb{Z}\big[G_m\big], \qquad \text{avec } \varepsilon_\sigma \in \{0,1\} \text{ pour tout } \sigma \in G_m.$$

Dans le cas où le conducteur est premier, notre base courte coïncide avec la base donnée dans [Sch08, Th. 9.3(i)]. Un ingrédient d'intérêt indépendant de la preuve consiste à décrire une vaste famille d'éléments courts de $\mathcal{S}_m$, qui contient l'ensemble identifié dans [CDW21, §4.2]. Cette description utilise un critère arithmétique très simple, dans l'esprit de [Was97, Lem. 16.3] quand m est une puissance d'un nombre premier impair. Nous obtenons notre base courte en choisissant astucieusement certains éléments $\alpha_m(b)$ parmi cette grande famille d'éléments courts.

Nous montrons également comment calculer explicitement les entiers algébriques qui génèrent $\mathfrak{L}^{\alpha_m(b)}$, pour tout idéal premier non ramifié $\mathfrak{L}$ et tout élément $\alpha_m(b)$ de notre base courte. Ces générateurs s'expriment comme des sommes de Jacobi qui s'avèrent extrêmement plus efficaces à calculer que les générateurs donnés p. ex., dans [Was97, §6.2].

Pour terminer, une conséquence théorique intéressante de notre résultat consiste à dériver une borne supérieure sur le nombre de classes relatives du m-ième corps cyclotomique. La preuve de notre borne donne également un algorithme pour calculer le nombre de classes relatives grâce au déterminant d'un multiple d'une matrice de Hadamard : incidemment, cette méthode semble significativement plus efficace que d'utiliser la formule analytique traditionnelle quand le nombre de facteurs premiers distincts de m est petit.

Réseaux log-$\mathcal{S}$-unités à partir de générateurs explicites de Stickelberger

Dans notre dernière contribution, nous étendons les expériences de Twisted-PHS à tous les corps cyclotomiques de degrés allant jusqu'à 210. Ceci fait sauter la barrière des petites dimensions et permet d'atteindre des tailles de paramètres où les phénomènes asymptotiques, p. ex., la croissance exponentielle du nombre de classes, commencent à s'exprimer pleinement.

Cette percée est obtenue grâce à des améliorations à la fois théoriques et d'implémentation. Tout d'abord, nous montrons comment obtenir une famille de $\mathcal{S}$-unités indépendantes et de rang plein à partir d'un ensemble de $\mathcal{S}^+$-unités fondamentales du sous-corps réel maximal, par l'adjonction de générateurs explicites correspondant à une base de l'idéal de Stickelberger. Grâce aux techniques avancées développées précédemment sur l'idéal de Stickelberger, il s'avère que

ces générateurs s'expriment toujours comme des sommes de Jacobi, qui sont particulièrement petites et faciles à calculer. Cette famille de rang plein génère un sous-groupe des $\mathcal{S}$-unités d'indice explicitement calculable, dont nous explicitons et prouvons la valeur exacte. Cet indice contient une large puissance de 2 qui peut être retirée grâce à des techniques classiques de 2-*saturation*, pour lesquelles les nouveaux résultats sur l'idéal de Stickelberger sont essentiels. Ainsi, nous obtenons des sous-réseaux du réseau log-$\mathcal{S}$-unité complet sur lesquels tester l'algorithme Twisted-PHS. Nous fournissons ici aussi une implémentation complète, publiquement disponible à l'adresse GitHub: ob3rnard/Tw-Sti[8].

Les facteurs d'approximation obtenus par nos expériences ne montrent ni un impact catastrophique des attaques par $\mathcal{S}$-unités, ni ne permettent d'écarter la menace. En effet, le mode approché utilisé au-delà de la dimension 80 donne seulement une borne supérieure sur les performances de l'algorithme Twisted-PHS. Néanmoins, nous observons une forte corrélation entre la densité du sous-réseau log-$\mathcal{S}$-unité utilisé et le facteur d'approximation obtenu par l'algorithme Twisted-PHS : plus le réseau est dense, meilleures sont les performances. Nous sommes également en mesure de confirmer la nature géométrique très particulière du réseau log-$\mathcal{S}$-unité déjà observée en petite dimension, pour tous les corps cyclotomiques, tous les sous-réseaux log-$\mathcal{S}$-unités et toutes les bases de facteurs considérés. Ces observations récurrentes dans des régimes très différents suggèrent que ceci est possiblement l'émanation de phénomènes algébriques plus profonds, une observation qui a été récemment développée par Bernstein et Lange [BL21].

Quoi qu'il en soit, le fait de rassembler toutes ces données en dimensions suffisamment grandes est d'une importance capitale afin de mieux comprendre les performances des attaques par $\mathcal{S}$-unités, et doit être vu comme une première étape avant d'obtenir une estimation fiable du comportement asymptotique de l'algorithme Twisted-PHS, ou plus généralement des attaques par $\mathcal{S}$-unités dans n'importe quel régime.

Dans un résultat complémentaire, nous utilisons la connaissance explicite de ces générateurs de Stickelberger, ainsi que le réseau de toutes les relations de classes réelles, pour enlever presque toutes les étapes quantiques de l'algorithme CDW tout en prouvant le même facteur d'approximation, sous l'hypothèse relativement inoffensive que la partie réelle du nombre de classes vérifie $h_m^+ \leq O(\sqrt{m})$.

[8] https://github.com/ob3rnard/Tw-Sti

Publications

On this page are summarized the contributions of this thesis, together with their publication status as of June, 2022.

> **[BR20] Twisted-PHS: Using the Product Formula to Solve Approx-SVP in Ideal Lattices.**
> Olivier Bernard and Adeline Roux-Langlois.

Published in the proceedings of Asiacrypt 2020, Part II, vol. 12492 of *Lecture Notes in Computer Science* (LNCS) Series, pp.349–380, Springer.

Keywords: Ideal lattices, Approx-SVP, S-unit attacks, Twisted-PHS algorithm.
Links: [ePrint: 2020/1081[9] | GitHub: ob3rnard/Twisted-PHS[7]]

> **[BK21] A short basis of the Stickelberger ideal of a cyclotomic field.**
> Olivier Bernard and Radan Kučera.

Submitted to AMS :: Mathematics of Computation (*American Mathematical Society*).

2010 MSC classes: 11R18 (Primary), 11R29, 11Y40 (Secondary).
Keywords: Cyclotomic fields, Stickelberger ideal, short basis, relative class number.
Links: [arXiv: 2109.13329 [math.NT][10]]

> **[BLNR21] Log-$\mathcal{S}$-unit lattices using Explicit Stickelberger Generators to solve Approx Ideal-SVP.**
> Olivier Bernard, Andrea Lesavourey, Tuong-Huy Nguyen and Adeline Roux-Langlois.

Accepted for publication in the proceedings of Asiacrypt 2022, in *Lecture Notes in Computer Science* (LNCS) Series, Springer..

Keywords: Ideal lattices, Approx-SVP, Stickelberger, S-unit attacks, Twisted-PHS algorithm.
Links: [ePrint: 2021/1384[11] | GitHub: ob3rnard/Tw-Sti[8] | Blog: H2020 Prometheus[12]]

[9] https://eprint.iacr.org/2020/1081
[7] https://github.com/ob3rnard/Twisted-PHS
[10] https://arxiv.org/abs/2109.13329
[11] https://eprint.iacr.org/2021/1384
[8] https://github.com/ob3rnard/Tw-Sti
[12] https://www.h2020prometheus.eu/dissemination/blog

Chapter 1

Introduction

Obliged by the hypothetic threat of the construction of a large scale quantum computer in the next few decades, the cryptographic community has been driven to consider new mathematical problems to serve as the security foundations for so-called *post-quantum* public-key cryptosystems. In 2016, the U.S. *National Institute of Standards and Technology* (NIST) launched the *Post-Quantum Cryptography Standardization* competition to evaluate and standardize quantum-resistant public-key algorithms. Around 70 proposals were received, involving several mathematical objects such as, to name a few, Euclidean lattices, error-correcting codes or supersingular isogeny graphs.

As shown by the large number of submissions for this family, one of the most promising post-quantum solution is based on Euclidean lattices. Several hard problems are used to prove the security of these cryptosystems, such as the Ntru problem [HPS98], the *Short Integer Solution* (Sis) problem [Ajt96] or the *Learning With Errors* (Lwe) problem [Reg05], and their algebraically structured variants *Ring* (Ring-Sis [LM06, PR06], Ring-Lwe [SSTX09, LPR10]) or *Module* (Module-Sis, Module-Lwe [LS15]). Typically, the structured variants offer the advantage of a better efficiency, at the price of possibly losing some of the security, which is ultimately relying on the hardness of the *Approximate Shortest Vector Problem* (Approx-Svp) in the restricted corresponding class of algebraically structured Euclidean lattices.

In the case of arbitrary lattices, Svp is a well-studied NP-hard problem [Ajt98]. Its Approximate version consists, for any lattice of rank n, in finding a non-zero vector of the lattice, whose Euclidean norm is within a small multiplicative factor from the length of the shortest non-zero vector in the lattice. The best trade-off in this case is given by Schnorr's hierarchy [Sch87], which allows to reach an approximation factor $2^{\widetilde{\mathcal{O}}(n^{\omega})}$ in time $2^{\widetilde{\mathcal{O}}(n^{1-\omega})}$ for any $\omega \in (0,1)$, as represented on Fig. 1.1a. In practice, the best known algorithm that is close to this trade-off is the *Block Korkin-Zolotarev* (BKZ) algorithm [SE94], which can be seen as an improvement of the well-known LLL algorithm [LLL82] due to A. Lenstra, H. Lenstra and L. Lovász.

However, these structured assumptions (e.g., Ring-Lwe) could become vacuous if the underlying variants of Approx-Svp become easier on the specific class of algebraically structured lattices. Hence, a natural target for cryptanalysis is the *Ideal Shortest Vector Problem* (id-Svp) which focuses on *ideal lattices* corresponding, under the Minkowski embedding, to fractional ideals of the ring of integers $\mathcal{O}_K$ of a number field K. For a long time, the best known algorithm to solve Approx-Svp in ideal lattices was the same as for arbitrary lattices, but recently, a series of works [CGS14, EHKS14, BS16, CDPR16, CDW17, DPW19, PHS19a] tends to show that solving this problem could be easier in ideal lattices, in particular in the quantum setting.

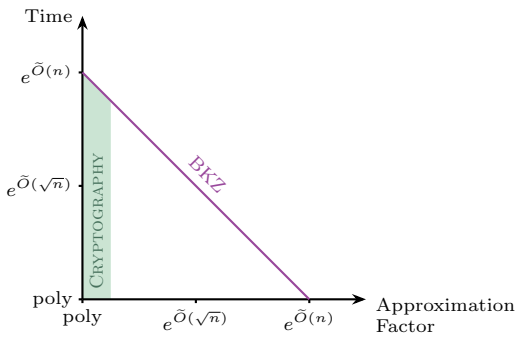

(a) Schnorr's hierarchy [Sch87] (unstructured case)

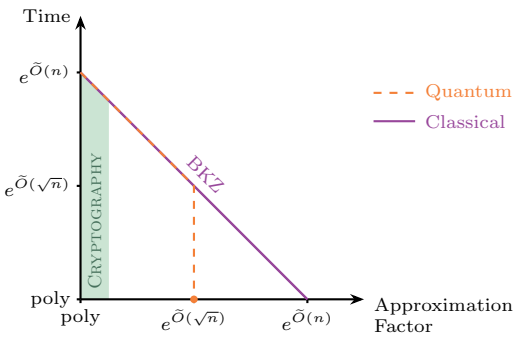

(b) CDW algorithm [CDW21] (cyclotomic fields)

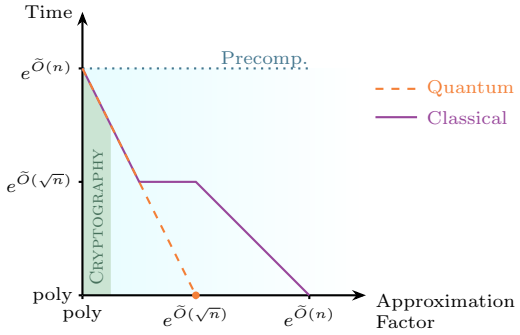

(c) PHS algorithm [PHS19a] (cyclotomic fields)

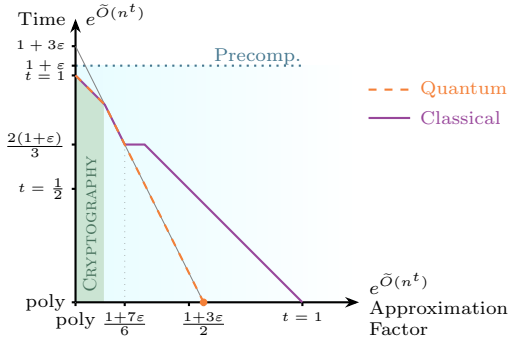

(d) PHS algorithm [PHS19a] for number fields K s.t. $\log|\Delta_K| = \widetilde{O}(n^{1+\varepsilon})$

FIGURE 1.1 – Trade-offs between runtime and approximation factors reached by algebraic cryptanalyses of id-SVP.

Quantum Algorithms for Number Theory

Indeed, the discovery of new number-theoretic polynomial-time *quantum* algorithms showcased how the strong algebraic structure of these ideal lattices could be used to tackle id-SVP more efficiently than by relying on traditional lattice reduction algorithms.

Everything started with the buzzing note of CAMPBELL, GROVES and SHEPHERD [CGS14], that claimed, without proofs, a quantum polynomial-time attack against a scheme named Soliloquy, solving specific instances of the Approx-SVP on *principal* ideal lattices. Their algorithm has two successive steps:

- the first one is solving the *Principal Ideal Problem* (PIP) that asks for any generator of a principal ideal,
- the second one is shortening this generator as much as possible using the algebraic units of the field, which reduces to solving a *Closest Vector Problem* (CVP) in the log-unit lattice.

The former is claimed to run in quantum polynomial-time, and the latter is claimed to be sufficiently easy in the case of cyclotomic fields using *circular units* to practically break the scheme.

The first claim was proven independently in [EHKS14], where the authors described a generalization of Shor's algorithm [Sho97], to compute unit groups of number fields of arbitrary degree in quantum polynomial time. Later on, building upon [EHKS14], BIASSE and SONG [BS16] extended this result to the computation of class groups and $\mathcal{S}$-unit groups of arbitrary degree number fields. More precisely, they showed how to compute $\mathcal{S}$-units, a generalization of

the algebraic units of a number field depending on a set $\mathcal{S}$ of prime ideals, in quantum polynomial time in the size of the discriminant Δ_K of the number field K and in the size of the *factor base* $\mathcal{S}$. They also showed how the Pip resolution, as well as computing class groups or unit groups, can be reduced to these $\mathcal{S}$-unit computations for adequately chosen prime ideals in $\mathcal{S}$.

Algebraic Cryptanalyses of id-SVP

As for the second claim of [CGS14], Cramer, Ducas, Peikert and Regev [CDPR16] proved that, in prime-power cyclotomic fields, logarithmic embeddings of circular units [Was97, §8] yield a sufficiently good basis of a sublattice of relatively small finite index inside the log-unit lattice. This key property allowed them to conclude that there exists a polynomial-time quantum algorithm that, on average, solves Approx-Svp on principal ideal lattices for an approximation factor $2^{\widetilde{O}(\sqrt{n})}$, where n is the dimension of the ideal.

Subsequently, this first algebraic cryptanalysis led to several generalizations, extending to any class of fractional ideal [CDW17], any cyclotomic fields [CDW21] and to any number field [PHS19a]. For any challenge ideal $\mathfrak{b}$ of a number field K, all approaches start from a solution to the *Class Group Discrete Logarithm Problem* (ClDlp). This representation problem asks, given a fixed set of finite places corresponding to prime ideals $\{\mathfrak{p}_1, \ldots, \mathfrak{p}_k\}$ of K, to find, if they exist, $\alpha \in K$ and $e_1, \ldots, e_k \in \mathbb{Z}$ such that:

$$\langle \alpha \rangle = \mathfrak{b} \cdot \prod_{1 \leq i \leq k} \mathfrak{p}_i^{e_i}.$$

This problem reduces to some $\mathcal{S}$-unit group computation, hence is not hard to solve in a quantum world according to the previous discussion. So, the most difficult part of these cryptanalyses resides in reducing the Euclidean norm of α, i.e., to find a *shortest* solution to the ClDlp above, or equivalently, the shortest coset representative modulo the finitely generated multiplicative group H of class group relations between the $\mathfrak{p}_i$'s. Using a suitable logarithmic embedding, this boils down to solve a Cvp instance in the image lattice of H under this logarithmic embedding. At the end, this shortest coset representative is hoped to be a sufficiently small element of the challenge ideal. Therefore, the choice of the logarithmic embedding is particularly important, since it must convey all useful informations on the size of α, and since the obtained lattice basis must not be too bad for the Cvp solver. Note that [CGS14, CDPR16] exactly follow this procedure for $k = 0$, in which case a solution to the ClDlp exist if and only if $\mathfrak{b}$ is principal.

For our purpose, we will separate these algebraic cryptanalyses between two lines of work, that use different tools to guarantee the output size, and have different scopes:

- The CDW algorithm, by Cramer, Ducas and Wesolowski [CDW17, CDW21], solves id-Svp for approximation factors $\exp \widetilde{O}(\sqrt{n})$ in cyclotomic fields of degree n, in *quantum* polynomial time. This trade-off, depicted in Fig. 1.1b, is proven under "*carefully justified heuristics*". The algorithm uses the Stickelberger ideal of a cyclotomic field, a special ideal providing free *short* relations in the relative part of the ideal class group. These short relations allow to find a *close* principal multiple for any challenge ideal, i.e., a principal multiple whose algebraic norm is relatively small when divided by the challenge ideal norm. Then, the [CDPR16] routine is applied to a generator of this multiple, hoping that its output is sufficiently short.
- These two steps can actually be combined in a single Cvp instance, giving rise to what are now called $\mathcal{S}$-*unit attacks*: the idea is to find in this way a principal multiple which is not only of small algebraic norm, but is also generated by a *small* element. This was the core idea of the algorithm of Pellet-Mary, Hanrot and Stehlé (PHS) [PHS19a], which applies to any number field, and which we detail in the next section.

The practical impact of the CDW algorithm was evaluated in [DPW19] by running numerous simulations for the CVP in each of the two lattices involved. From these experimental results, they heuristically derive a *volumetric lower bound* [DPW19, Eq. (5) and Tab. 1] and conclude that the CDW algorithm should beat BKZ_{300} for cyclotomic fields of degree larger than 7000.[13]

$\mathcal{S}$-unit Attacks

We describe in more detail the PHS algorithm [PHS19a], by PELLET-MARY, HANROT and STEHLÉ, which is to our knowledge the first $\mathcal{S}$-unit attack described and proven in the literature, even though this $\mathcal{S}$-unit formalism was not directly used in [PHS19a].

The main feature of their algorithm is to combine in a single CVP instance the two principal resolution steps of the CDW algorithm [CDW17, CDW21], namely the CPMP (*Close Principal Multiple Problem*) and the SGP (*Shortest Generator Problem*). This provides some guarantee that the output of the CPMP solver has a generator which is "*not much larger*" than its shortest non-zero vector. Unfortunately, this comes at the price of an exponential amount of preprocessing, depending only on the number field K. Indeed, in order to guarantee the output size and the running time of the algorithm, a key ingredient is to use a CVP with preprocessing hint algorithm due to LAARHOVEN [Laa16]. More formally, the PHS algorithm is split in two phases:

1. The preprocessing phase builds a specific lattice, depending only on the field K, which can be viewed as a log-$\mathcal{S}$-unit lattice under a particular logarithmic embedding, together with some hint allowing to efficiently solve Approx-CVP instances inside this lattice. Denoting by Δ_K the discriminant of K, this phase runs in time $2^{\widetilde{O}(\log|\Delta_K|)}$ and outputs a hint $\mathcal{V}$ of bit-size $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$, where $\omega \in \left[0, \frac{1}{2}\right]$ is the trade-off parameter.
2. The query phase reduces each Approx-id-SVP challenge to an Approx-CVP instance in this fixed lattice. It takes as inputs any ideal of $\mathcal{O}_K$, whose algebraic norm has bit-size bounded by $2^{\mathrm{poly}(\log|\Delta_K|)}$, the hint $\mathcal{V}$, and runs in time $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)} + \mathrm{T}_{\mathsf{Su}}(K)$. It outputs a non-zero element of the ideal which solves Approx-SVP with an approximation factor $2^{\widetilde{O}(\log^{\omega+1}|\Delta_K|/n)}$, where n is the degree of K.

Here, $\mathrm{T}_{\mathsf{Su}}(K)$ denotes the running time for $\mathcal{S}$-unit groups related computations, i.e., in a quantum world, $\mathrm{T}_{\mathsf{Su}}(K) = \widetilde{O}\big(\ln|\Delta_K|\big)$ is polynomial [BS16], whereas in a classical world, it remains subexponential in $\ln|\Delta_K|$, i.e., $\mathrm{T}_{\mathsf{Su}}(K) = \exp\widetilde{O}(\ln^{\alpha}|\Delta_K|)$, where $\alpha = 1/2$ for cyclotomic fields [BEF+17],[14] and $\alpha = 2/3$ in the general case [BF14], recently lowered to 3/5 by GÉLIN [Gél17].

This trade-off is shown on Fig. 1.1c and 1.1d on resp. cyclotomic fields and number fields K with $\log|\Delta_K| \leq \widetilde{O}(n^{1+\varepsilon})$. Ignoring the preprocessing cost, the query phase beats the traditional Schnorr's hierarchy [Sch87] when $\log|\Delta_K| \leq \widetilde{O}(n^{1+\varepsilon})$ with $\varepsilon = 1/3$ in the quantum case, and $\varepsilon = 1/11$ in the classical case [PHS19a, Fig. 5.3]. It should be noted however that these bounds on the discriminant are not uniform as the approximation factor varies, e.g., for an approximation factor set to $2^{\sqrt{n}}$, the time complexity of the PHS algorithm asymptotically beats Schnorr's hierarchy only in the quantum case and only for $\varepsilon \leq 1/6$.

Contributions of this Thesis

The contributions of this thesis take place in the context of $\mathcal{S}$-unit attacks. First, the $\mathcal{S}$-unit formalism is used to propose a twisted version of the PHS algorithm that reveals extremely

[13] The first published version of [DPW19] reported a crossover point at degree 12 000. After fixing a bug in the implementation, which was pointed out by BERNSTEIN on 20th August 2021 in a talk at SIAM Conference, this crossover point has been reevaluated to 7000 [DPW19, Fig. 5].

[14] The article [BEF+17] is written for prime-power cyclotomic fields for historical reasons, but readily adapts to the general case for class group computations.

powerful in practice. Then, the following contributions join for all cyclotomic fields the two lines of cryptanalyses described above, by using extended techniques related to the Stickelberger lattice to both remove quantum steps from the CDW algorithm and experimentally approximate the Twisted-PHS algorithm in medium dimensions, where asymptotic phenomena start to express.

Twisted-PHS: using the product formula

In fact, the particular lattice used in the PHS algorithm corresponds to a special lattice called the *log-$\mathcal{S}$-unit* lattice, i.e., a lattice obtained by applying some logarithmic embedding on $\mathcal{S}$-units, where $\mathcal{S}$ can be identified to a *factor base* FB of prime ideals. As it turns out, choosing carefully the used logarithmic embedding is particularly important in practice.

Hence, our first contribution is to propose in Ch. 3 a new *twisted* version of the PHS algorithm, that we call *Twisted-PHS*, whose core idea consists in identifying a logarithmic embedding preserving the natural algebraic properties of $\mathcal{S}$-units. More precisely, we include the standard number-theoretic weights coming from the *Product Formula* (see e.g., §2.1.2) to the coordinates of the logarithmic embedding, i.e., for any α in a number field K, we start from:

$$\mathrm{Log}_{\mathcal{S}}\,\alpha = \Big(\big\{[K_\sigma : \mathbb{R}] \cdot \ln|\sigma(\alpha)|\big\}_{\sigma\in\mathcal{S}_\infty}, \big\{-v_{\mathfrak{p}}(\alpha)\cdot \ln \mathcal{N}(\mathfrak{p})\big\}_{\mathfrak{p}\in \mathrm{FB}}\Big),$$

where $K_\sigma = \mathbb{R}$ (resp. $\mathbb{C}$) for any real (resp. complex) embedding $\sigma \in \mathcal{S}_\infty$ from K to $\mathbb{R}$ (resp. $\mathbb{C}$), and for any prime ideal $\mathfrak{p} \in$ FB, $\mathcal{N}(\mathfrak{p})$ is its algebraic norm and $v_{\mathfrak{p}}(\alpha)$ is the valuation of α at $\mathfrak{p}$. By contrast, the log-embedding on which is based [PHS19a] does not include the $\ln \mathcal{N}(\mathfrak{p})$ weights.

Using this $\mathcal{S}$-unit formalism, we prove in Th. 3.14 that our Twisted-PHS algorithm reaches the same asymptotic trade-off between runtime and approximation factor than the PHS algorithm, using the same Cvp solver with preprocessing hint due to Laarhoven [Laa16] to efficiently solve Approx-Cvp instances in the log-$\mathcal{S}$-unit lattice. As a secondary contribution, we also propose several improvements of the PHS algorithm, in an optimized version described in §3.2.3.

Intuitively, adding weights $\ln \mathcal{N}(\mathfrak{p})$ to integer valuations at any prime ideal $\mathfrak{p}$ captures the fact that using a relation increasing the valuations at big norm ideals costs more than using a relation involving smaller norm ideals. This also encodes in the log-$\mathcal{S}$-unit lattice the information on the length and algebraic norm of the $\mathcal{S}$-units, unlike the log-embedding used in [PHS19a] involving only the integer valuations. In the end, these rationales indicate that the Cvp solver in the *twisted* log-$\mathcal{S}$-unit lattice combines more efficiently the goal of searching for a principal multiple of small algebraic norm while still minimizing the size of its generator.

Another fundamental consequence of using a properly normalized logarithmic embedding is what we call the *optimal factor base phenomenon*, i.e., we prove that there exists a factor base $\mathcal{S}$ for which the density of the log-$\mathcal{S}$-unit lattice is maximal. Such a basis is computed by Alg. 3.3.

On the practical side, we provide a fully functional end-to-end implementation of the Twisted-PHS algorithm, where Laarhoven's Cvp oracle is replaced by Babai's Nearest Plane algorithm [Bab86]. This implementation is publicly available at GitHub: ob3rnard/Twisted-PHS[7]. For the first time, this allowed to run complete $\mathcal{S}$-unit attacks on a significant range of concrete examples. Our experiments suggested, for prime conductor cyclotomic fields and NTRU Prime fields of small dimensions, namely up to 70, that:

- under the proper number-theoretic normalization, the log-$\mathcal{S}$-unit lattices at hand have a very particular geometric behaviour and seem very easy to reduce (see §§3.4.1 and 3.4.2);
- the obtained *exact* approximation factors increase very slowly with the dimension (see e.g., Fig. 1.2), "*in a way that could reveal subexponential or even better*".

[7] https://github.com/ob3rnard/Twisted-PHS

To our knowledge, these were the first experimental evidence of the geometric peculiarity of properly normalized log-$\mathcal{S}$-unit lattices and of the practical potential of $\mathcal{S}$-unit attacks. Unfortunately, due to the classical complexity of computing $\mathcal{S}$-units, the attained dimensions are not sufficient to conjecture the asymptotic behaviour of the Twisted-PHS algorithm.

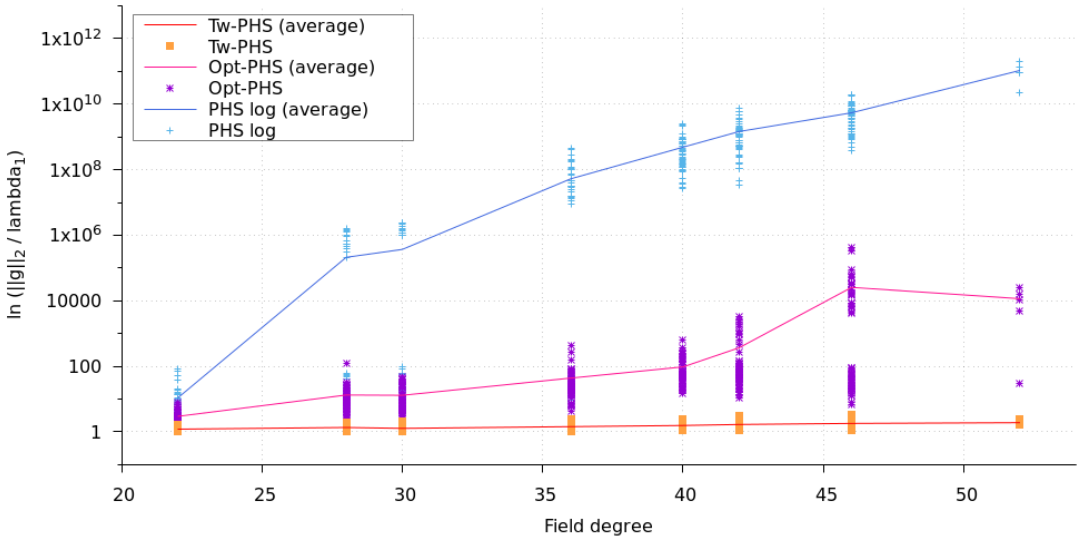

FIGURE 1.2 – Approximation factors reached by Twisted-PHS, Opt-PHS and PHS for cyclotomic fields of conductors 23, 29, 31, 37, 41, 43, 47 and 53 (in log scale).

[BR20] Twisted-PHS: Using the Product Formula to Solve Approx-SVP in Ideal Lattices.
Olivier BERNARD and Adeline ROUX-LANGLOIS.

Published in the proceedings of Asiacrypt 2020, Part II, vol. 12492 of *Lecture Notes in Computer Science* (LNCS) Series, pp.349–380, Springer.

Keywords: Ideal lattices, Approx-SVP, S-unit attacks, Twisted-PHS algorithm.
Links: [ePrint: 2020/1081[9] | GitHub: ob3rnard/Twisted-PHS[7]]

A short basis of the Stickelberger ideal

In the next contribution, we shall see that explicitly computing Stickelberger generators is useful in at least two situations: the first one occurs for removing the last PIP quantum step in the CDW algorithm [CDW21], the second one occurs when approximating the log-$\mathcal{S}$-unit lattice used in the Twisted-PHS algorithm. In the latter, some of the computational steps, notably the 2-*saturation* procedure to obtain denser lattices, become quickly intractable as the bit size of the elements coefficients grows. This motivates us to constrain both the number of Stickelberger generators we use and their size.

This is where our second contribution, given in Ch. 4, reveals extremely useful. Our main result (see Th. 4.29) is to provide the first explicit *short basis* of the Stickelberger ideal $\mathcal{S}_m$ of

[9] https://eprint.iacr.org/2020/1081
[7] https://github.com/ob3rnard/Twisted-PHS

cyclotomic fields of any conductor m, i.e., a basis containing *only* short elements. By definition, an element of $\mathbb{Z}[G_m]$, where G_m denotes the Galois group of the m-th cyclotomic field, is called short whenever it writes as:

$$\sum_{\sigma \in G_m} \varepsilon_\sigma \cdot \sigma \in \mathcal{S}_m \subset \mathbb{Z}\big[G_m\big], \qquad \text{where } \varepsilon_\sigma \in \{0,1\} \text{ for all } \sigma \in G_m.$$

In the prime conductor case, our short basis coincides with the basis given in [Sch08, Th. 9.3(i)]. One ingredient of independent interest in the proof is Pr. 4.15, which describes a large family of short elements of $\mathcal{S}_m$ that encompasses the set from [CDW21, §4.2]. This description uses a very simple arithmetic criterion in the spirit of [Was97, Lem. 16.3] when m is an odd prime power. Picking wisely some elements $\alpha_m(b)$ in this large family yields our proposed short basis.

We also show how to explicitly compute algebraic integers generating $\mathfrak{L}^{\alpha_m(b)}$, for any unramified prime ideal $\mathfrak{L}$ and any element $\alpha_m(b)$ of our short basis. These generators can be expressed as Jacobi sums that turn out to be drastically more efficient to compute than the generators given e.g., in [Was97, §6.2].

Finally, a nice theoretical consequence of our result is to derive an explicit upper bound on the relative part of the class number of the m-th cyclotomic field, given in Cor. 4.32. The proof of our bound also gives an algorithm to compute the relative class number by computing the determinant of some scaled Hadamard matrix: incidentally, this method seems to be significantly more efficient than when using the traditional analytic formula (see e.g., Eq. (2.10)), when the number t of prime factors of m is small.

> **[BK21] A short basis of the Stickelberger ideal of a cyclotomic field.**
> Olivier Bernard and Radan Kučera.

Submitted to AMS :: Mathematics of Computation (*American Mathematical Society*).

2010 MSC classes: 11R18 (Primary), 11R29, 11Y40 (Secondary).
Keywords: Cyclotomic fields, Stickelberger ideal, short basis, relative class number.
Links: [arXiv: 2109.13329 [math.NT][10]]

Log-$\mathcal{S}$-unit lattices using explicit Stickelberger generators

In our last contribution, given in Ch. 5, we extend the experiments of Ch. 3 to cyclotomic fields of any conductor m and of degree up to 210. This effectively breaks the small dimension barrier and reaches ranges of parameters where asymptotic phenomena, e.g., the exponential growth of the class number, start to express.

This breakthrough is obtained as the result of both theoretical and implementational improvements. First, we prove in Th. 5.14 that a full-rank family of independent $\mathcal{S}$-units can be lifted from a set of fundamental $\mathcal{S}^+$-units of the maximal real subfield by adjoining the explicit generators corresponding to a basis of the Stickelberger ideal. Using results from Ch. 4, it turns out these generators are always expressed by Jacobi sums, which are particularly small and easy to compute. This full-rank family generates an $\mathcal{S}$-unit subgroup of explicitly computable index, as we also prove in Th. 5.14. This index contains a large power of 2 that can be removed using classical 2-*saturation* techniques, for which using the results of Ch. 4 is essential. At the end, we obtain sublattices of the full log-$\mathcal{S}$-unit lattice on which to test the Twisted-PHS algorithm. We also provide a full implementation, publicly available at GitHub: ob3rnard/Tw-Sti[8].

[10] https://arxiv.org/abs/2109.13329
[8] https://github.com/ob3rnard/Tw-Sti

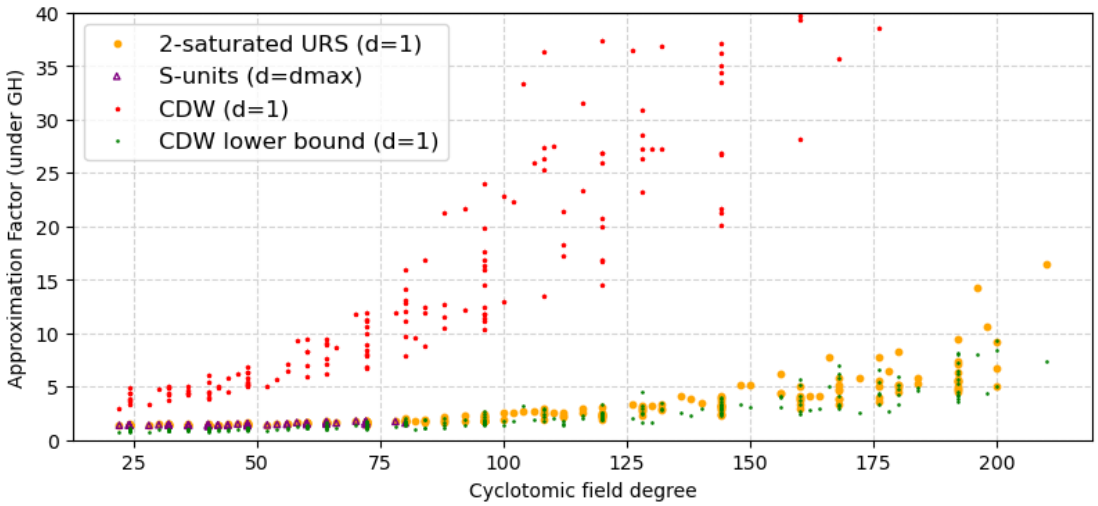

FIGURE 1.3 – Approximation factors comparison for cyclotomic fields K_m of degree $\varphi(m) \leq 210$ with $h_m^+ = 1$, under the Gaussian Heuristic. Our results, labelled as "2-saturated URS", bound Twisted-PHS from above.

The approximation factors obtained in our experiments are detailed in Fig. 1.3. We stress that this graph does neither show a catastrophic impact of $\mathcal{S}$-unit attacks, nor does it clear the threat. Indeed, the approximated mode used beyond dimension 80 only gives a practical upper bound on the performance of the Twisted-PHS algorithm. Nevertheless, we observe a strong correlation between the density of the used log-$\mathcal{S}$-unit sublattice and the approximation factor obtained by the Twisted-PHS algorithm: the denser, the better. We are also able to confirm the peculiar geometric nature of the log-$\mathcal{S}$-unit lattice already observed in Ch. 3, across *all* considered cyclotomic fields, log-$\mathcal{S}$-unit sublattices and factor bases. These recurrent observations in very different regimes suggest that this phenomenon has a possibly deep explanation, an observation that has been recently developed by BERNSTEIN and LANGE in [BL21].

Anyhow, gathering these extensive data in meaningful dimensions is of utmost importance to better understand the performance of $\mathcal{S}$-unit attacks, and should be seen as a first step towards getting a sound estimation of the asymptotic behaviour of the Twisted-PHS algorithm or $\mathcal{S}$-unit attacks with any kind of parameters.

As a side result, we use the knowledge of these explicit Stickelberger generators, as well as the full lattice of real class group relations, to remove almost all quantum steps in the CDW algorithm, under the mild restriction that the plus part of the class number verifies $h_m^+ \leq O(\sqrt{m})$.

[BLNR21] Log-$\mathcal{S}$-unit lattices using Explicit Stickelberger Generators to solve Approx Ideal-SVP.
Olivier BERNARD, Andrea LESAVOUREY, Tuong-Huy NGUYEN and Adeline ROUX-LANGLOIS.

Accepted for publication in the proceedings of Asiacrypt 2022, in *Lecture Notes in Computer Science* (LNCS) Series, Springer..

Keywords: Ideal lattices, Approx-SVP, Stickelberger, S-unit attacks, Twisted-PHS algorithm.
Links: [ePrint: 2021/1384[11] | GitHub: ob3rnard/Tw-Sti[8] | Blog: H2020 Prometheus[12]]

[11] https://eprint.iacr.org/2021/1384
[8] https://github.com/ob3rnard/Tw-Sti
[12] https://www.h2020prometheus.eu/dissemination/blog

Chapter 2

Preliminaries

Number-theoretic objects and properties required within this thesis are recalled in this chapter. The first section introduces $\mathcal{S}$-unit groups and the properties of their associated log-$\mathcal{S}$-unit lattices; in particular, the *Product Formula* plays a central role in our cryptanalyses. Then, the special case of cyclotomic fields, for which many remarkable properties are known, is detailed. The third section deals with algorithmic number theory, including several number-theoretic bounds that are needed in the complexity proofs, and we finish by a piece of Euclidean lattices theory, notably on how to evaluate the quality of a lattice basis w.r.t. the *Closest Vector Problem* (Cvp).

Contents

Notations. Let $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$ and $\mathbb{C}$ denote the integers, rational, real and complex numbers respectively. For any $i, j \in \mathbb{Z}$ with $i \leq j$, let $[\![i, j]\!]$ denote the set $\{k \in \mathbb{Z};\ i \leq k \leq j\}$ of all integers between i and j. For any $x \in \mathbb{Q}$, let $\{x\}$ (resp. $[x]$) denote its fractional (resp. integral) part, i.e., such that $0 \leq \{x\} < 1$ and $[x] = x - \{x\} \in \mathbb{Z}$.

Any vector is designated by a bold letter $\mathbf{v}$, its i-th coordinate by v_i and its ℓ_p-norm, for $p \in \mathbb{N}^* \cup \{\infty\}$, by $\|\mathbf{v}\|_p$. As a special case, the n-dimensional vector whose coefficients are all 1's is written $\mathbf{1}_n$. All matrices will be given using *row* vectors, $\mathcal{D}_{\mathbf{v}}$ is the diagonal matrix with coefficients v_i on the diagonal, I_n is the identity and $\mathbf{1}_{n\times n}$ denotes the square matrix of dimension n filled with 1's.

2.1 On $\mathcal{S}$-unit Groups

2.1.1 Number fields, ideals and class groups

In this thesis, K always denotes a number field of degree n over $\mathbb{Q}$ and $\mathcal{O}_K$ its maximal order.The algebraic trace and norm of $\alpha \in K$, resp. denoted by $\mathrm{Tr}(\alpha)$ and $\mathcal{N}(\alpha)$, are defined as the trace and determinant of the endomorphism $x \mapsto \alpha x$ of K, viewed as a $\mathbb{Q}$-vector space. The discriminant of K is written Δ_K and can be defined, for any $\mathbb{Z}$-basis $\omega_1, \ldots, \omega_n$ of $\mathcal{O}_K$, as $\det\big(\mathrm{Tr}(\omega_i\omega_j)\big)_{i,j}$. Most complexities of number-theoretic algorithms depend on $\ln|\Delta_K|$.

Class groups.

The fractional ideals of K are designated by gothic letters, like $\mathfrak{b}$, and form a multiplicative group $\mathcal{I}_K$ containing the normal subgroup $\mathcal{P}_K := \big\{\langle\alpha\rangle;\ \alpha \in K\big\}$ of principal ideals. The quotient group $\mathcal{I}_K/\mathcal{P}_K$ is called the *class group* of K and denoted by Cl_K. The class group is a finite group, whose order h_K is called the *class number* of K. For any ideal $\mathfrak{b} \in \mathcal{I}_K$, the class of $\mathfrak{b}$ in Cl_K is denoted by $\big[\mathfrak{b}\big]$.

Finally, for any set of prime ideals $\big\{\mathfrak{L}_i;\ i \in [\![1, k]\!]\big\}$, we denote by $h_{K,(\mathfrak{L}_1,\ldots,\mathfrak{L}_k)}$ the cardinal of the subgroup of Cl_K generated by the k classes $\big[\mathfrak{L}_i\big]$, i.e., the determinant of the kernel of:

$$\mathfrak{f}_{\mathfrak{L}_1,\ldots,\mathfrak{L}_k} : (e_1, \ldots, e_k) \in \mathbb{Z}^k \longmapsto \prod_{1\leq i\leq k} \big[\mathfrak{L}_i\big]^{e_i} \in \mathrm{Cl}_K\,.$$

Specific families of number fields.

We will specifically target two families of number fields, widely used in cryptography [Pei16]: cyclotomic fields $\mathbb{Q}(\zeta_m)$, where $\zeta_m := e^{2i\pi/m}$ is a primitive m-th root of unity, and NTRU Prime [BCLV17] fields $\mathbb{Q}(z_q)$, where z_q is a root of $x^q - x - 1$ for q prime. Both families have discriminants of order n^n.

The case of cyclotomic fields is developed in §2.2; in particular, $\mathcal{O}_{\mathbb{Q}(\zeta_m)} = \mathbb{Z}[\zeta_m]$ and their discriminant is explicitly known from the factorization of m. For NTRU Prime fields, the situation is marginally more involved, as $\mathbb{Z}[z_q]$ is maximal if and only if its polynomial discriminant $D_0 = q^q - (q-1)^{q-1}$ [Swa62, Th. 2] is squarefree [Kom75, Th. 4] :

$$\Delta_{\mathbb{Q}(z_q)} = \prod_{p|D_0} p^{v_p(D_0) \bmod 2}, \quad \text{where } p^{v_p(D_0)} \text{ divides exactly } D_0.$$

Note however that there is strong evidence that such D_0's are generically squarefree, say with probability roughly 0.99 [BMT15, Conj. 1.1]. Actually, we checked that the conductor of $\mathbb{Z}[z_q]$ is not divisible by any of the first 10^6 primes for all $q \leq 1000$ outside the set $\{257, 487\}$, for which $59^2 \mid D_0$.

2.1.2 The Product Formula

Places of the number field K are usually split into two parts: the set $\mathcal{S}_\infty$ of *infinite* places can be identified with the embeddings of K into $\mathbb{R}$ or $\mathbb{C}$, up to conjugation; the set $\mathcal{S}_0$ of *finite* places is specified by the infinite set of prime ideals of K.

Let (r_1, r_2) be the signature of K with $n = r_1 + 2r_2$. The real embeddings of K are numbered from σ_1 to σ_{r_1}, whereas the complex embeddings come in pairs $(\sigma_j, \overline{\sigma}_j)$ for $j \in [\![r_1 + 1, r_2]\!]$. Each embedding σ of K into $\mathbb{C}$ induces an Archimedean absolute value $|\cdot|_\sigma$ on K, such that for $\alpha \in K$, $|\alpha|_\sigma = |\sigma(\alpha)|$; two complex conjugate embeddings yield the same absolute value. Thus, it is common to identify the set $\mathcal{S}_\infty$ of infinite places of K with the embeddings of K into $\mathbb{C}$ up to conjugation, so that $\mathcal{S}_\infty = \{\sigma_1, \ldots, \sigma_{r_1}, \sigma_{r_1+1}, \ldots, \sigma_{r_1+r_2}\}$. The completion of K with respect to the absolute value induced by an infinite place $\sigma \in \mathcal{S}_\infty$ is denoted by K_σ; it is $\mathbb{R}$ (resp. $\mathbb{C}$) for real places (resp. complex places).

Likewise, let $\mathfrak{p}$ be a prime ideal of $\mathcal{O}_K$ above $p \in \mathbb{Z}$ of residue degree f. For $\alpha \in K$, the largest power of $\mathfrak{p}$ that divides $\langle \alpha \rangle$ is called the valuation of α at $\mathfrak{p}$, and denoted by $v_\mathfrak{p}(\alpha)$; this defines a non-Archimedean absolute value $|\cdot|_\mathfrak{p}$ on K such that $|\alpha|_\mathfrak{p} = p^{-v_\mathfrak{p}(\alpha)}$. This absolute value can also be viewed as induced by any of the f embeddings of K into its $\mathfrak{p}$-adic completion $K_\mathfrak{p} \subseteq \mathbb{C}_p$, which is an extension of $\mathbb{Q}_p$ of degree f.

Hence, any place $v \in \mathcal{S}_\infty \cup \mathcal{S}_0$ induces an absolute value $|\cdot|_v$ on K, and Ostrowski's theorem for number fields ([Con, Th. 3], [Nar04, Th. 3.3]) shows that all possible absolute values on K are obtained in this way. A remarkable fact is that all these absolute values are tied together by the *Product Formula* ([Con, Th. 4], [Nar04, Th. 3.5]):

$$\prod_{\sigma \in \mathcal{S}_\infty} |\alpha|_\sigma^{[K_\sigma:\mathbb{R}]} \cdot \prod_{\mathfrak{p} \in \mathcal{S}_0 \supset p\mathbb{Z}} |\alpha|_\mathfrak{p}^{[K_\mathfrak{p}:\mathbb{Q}_p]} = \left(|\mathcal{N}(\alpha)| \cdot \prod_{\mathfrak{p} \in \mathcal{S}_0} \mathcal{N}(\mathfrak{p})^{-v_\mathfrak{p}(\alpha)} = \right) 1. \tag{2.1}$$

As all but finitely many of the $|\alpha|_v$'s, for $v \in \mathcal{S}_\infty \cup \mathcal{S}_0$, are 1, their product is really a finite product. Note that the $\mathcal{S}_\infty$ part of this product is $|\mathcal{N}(\alpha)|$, and each term of the $\mathcal{S}_0$ part can be written as $\mathcal{N}(\mathfrak{p})^{-v_\mathfrak{p}(\alpha)}$. This formula is actually a natural generalization to number fields of the innocuous looking product formula for $r \in \mathbb{Q}$, written as: $|r| \cdot \prod_{p \text{ prime}} p^{-v_p(r)} = 1$.

2.1.3 Logarithmic $\mathcal{S}$-embeddings

The idea of using $\mathcal{S}$-units for the cryptanalysis of id-Svp is implicitly underlying the work of [PHS19a], and is formalized in [BR20] (see Ch. 3). We introduce log-$\mathcal{S}$-unit lattices and discuss the proper normalization induced by the Product Formula that was at the heart of the practical improvements presented in Ch. 3.

$\mathcal{S}$-unit groups structure.

Fix a finite set $\mathcal{S}$ of places; in this thesis we shall consider that $\mathcal{S}$ *always* contains $\mathcal{S}_\infty$, hence $\mathcal{S}$ can be written as $\mathcal{S} = \mathcal{S}_\infty \cup \{\mathfrak{p}_1, \ldots, \mathfrak{p}_k\}$, where each $\mathfrak{p}_i \in \mathcal{S}_0$ corresponds to a prime ideal of K. For convenience, we sometimes call the finite places of $\mathcal{S}$ the *factor base*, denoted by FB, i.e., $\text{FB} = \mathcal{S} \cap \mathcal{S}_0 = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_k\}$. Note that we allow $k = 0$, in which case $\mathcal{S} = \mathcal{S}_\infty$ is omitted.

The so-called $\mathcal{S}$-unit group of K, denoted by $\mathcal{O}_{K,\mathcal{S}}^\times$, is the multiplicative subgroup of K generated by all elements whose valuations are non zero only at the finite places of $\mathcal{S}$. Formally:

$$\mathcal{O}_{K,\mathcal{S}}^\times = \left\{ \alpha \in K;\ \langle \alpha \rangle = \prod_{\mathfrak{p} \in \mathcal{S} \cap \mathcal{S}_0} \mathfrak{p}^{v_\mathfrak{p}(\alpha)} \right\}$$

Note that when $\mathcal{S} = \mathcal{S}_\infty$, we obtain the definition of the unit group $\mathcal{O}_K^\times$ as the multiplicative subgroup of invertible algebraic integers of $\mathcal{O}_K$. Both ($\mathcal{S}$-)unit groups always contain the finite torsion subgroup of roots of unity of K, denoted by $\mu(\mathcal{O}_K^\times)$.

Theorem 2.2 (Dirichlet-Chevalley-Hasse [Nar04, Th. III.3.12, Cor. 1])**.** *The $\mathcal{S}$-unit group is the direct product of the group of roots of unity $\mu(\mathcal{O}_K^\times)$ and a free abelian group with $\sharp\mathcal{S}-1$ generators. There exists a fundamental system of $\mathcal{S}$-units $\varepsilon_1, \ldots, \varepsilon_{\sharp\mathcal{S}-1}$ s.t. any $\mathcal{S}$-unit $\varepsilon \in \mathcal{O}_{K,\mathcal{S}}^\times$ uniquely writes as $\varepsilon = \mu \cdot \prod_{i=1}^{\sharp\mathcal{S}-1} \varepsilon_i^{k_i}$, where $\mu \in \mu(\mathcal{O}_K^\times)$ is a root of unity and $k_i \in \mathbb{Z}$.*

In particular, using $\mathcal{S} = \mathcal{S}_\infty$, we recover Dirichlet's unit theorem [Nar04, Th. 3.13], which states that $\mathcal{O}_K^\times$ is a finitely generated abelian group of rank $\nu := r_1 + r_2 - 1$. We shall assume that the fundamental elements $\varepsilon_1, \ldots, \varepsilon_{\sharp\mathcal{S}-1}$ of Th. 2.2 are ordered so that:

$$\mathcal{O}_K^\times \simeq \mu(\mathcal{O}_K^\times) \times \varepsilon_1^{\mathbb{Z}} \times \cdots \times \varepsilon_\nu^{\mathbb{Z}} \qquad \text{and} \qquad \mathcal{O}_{K,\mathcal{S}}^\times \simeq \mathcal{O}_K^\times \times \varepsilon_{\nu+1}^{\mathbb{Z}} \times \cdots \times \varepsilon_{\nu+k}^{\mathbb{Z}}.$$

Log-$\mathcal{S}$-unit lattices.

A fundamental ingredient of the proof of this theorem is to build an embedding of $\mathcal{O}_{K,\mathcal{S}}^\times$ into $\mathbb{R}^{\sharp\mathcal{S}}$, whose kernel is $\mu(\mathcal{O}_K^\times)$ and whose image is a lattice of dimension $(\sharp\mathcal{S} - 1)$. This embedding is called the *logarithmic $\mathcal{S}$-embedding*, and its image is called the log-$\mathcal{S}$-unit lattice.

Several equivalent definitions of this logarithmic $\mathcal{S}$-embedding are acceptable for the proof. However, for cryptanalytic purposes, experimental evidence given in Ch. 3 suggests that it is crucial to use a properly normalized embedding for the decodability of the log-$\mathcal{S}$-unit lattice. Thus, we define [Nar04, §3, p.98] the following log-$\mathcal{S}$-embedding from $K^\times$ to $\mathbb{R}^{r_1+r_2+k}$:

$$\mathrm{Log}_{\mathcal{S}}\, \alpha = \big([K_v : \mathbb{Q}_v] \cdot \ln|\alpha|_v\big)_{v\in\mathcal{S}} = \Big(\big\{[K_\sigma : \mathbb{R}] \cdot \ln|\sigma(\alpha)|\big\}_{\sigma\in\mathcal{S}_\infty}, \big\{-v_{\mathfrak{p}}(\alpha) \cdot \ln \mathcal{N}(\mathfrak{p})\big\}_{\mathfrak{p}\in\mathrm{FB}}\Big).$$

For $\mathcal{S} = \mathcal{S}_\infty$, this corresponds to the classical definition of the logarithmic embedding Log (see e.g., [Coh93, Def. 4.9.6]) from K to $\mathbb{R}^{r_1+r_2}$.

From the definition of $\mathcal{O}_{K,\mathcal{S}}^\times$ and Eq. (2.1), it is easy to see that $\mathrm{Log}_{\mathcal{S}}\, \mathcal{O}_{K,\mathcal{S}}^\times$ lies in the trace zero hyperplane orthogonal to $\mathbf{1}_{\sharp\mathcal{S}}$, i.e.:

$$\mathrm{Log}_{\mathcal{S}}\, \mathcal{O}_{K,\mathcal{S}}^\times \subset \mathbb{R}_0^{\sharp\mathcal{S}} := \big\{\mathbf{y} \in \mathbb{R}^{\sharp\mathcal{S}};\ \textstyle\sum_i y_i = 0\big\}.$$

Showing that its dimension is at least $\sharp\mathcal{S} - 1$ is more involved. Likewise, for any $\alpha \in K$, the sum of the coordinates of $\mathrm{Log}\, \alpha$ is precisely $\ln|\mathcal{N}(\alpha)|$, so that $\mathrm{Log}\, \mathcal{O}_K^\times$ lies in the trace zero hyperplane orthogonal to $\mathbf{1}_{r_1+r_2}$, i.e., $\mathrm{Log}\, \mathcal{O}_K^\times \subset \mathbb{R}_0^{r_1+r_2} = \big\{\mathbf{y} \in \mathbb{R}^{r_1+r_2};\ \sum_i y_i = 0\big\}$.

A row basis $\Lambda_{K,\mathcal{S}}$ of the log-$\mathcal{S}$-unit lattice $\mathrm{Log}_{\mathcal{S}}\, \mathcal{O}_{K,\mathcal{S}}^\times$ is given by the images of the fundamental system of $\mathcal{S}$-units of Th. 2.2 under the log-$\mathcal{S}$-embedding, i.e., $\Lambda_{K,\mathcal{S}} = \big(\mathrm{Log}_{\mathcal{S}}\, \varepsilon_i\big)_{1\le i\le \sharp\mathcal{S}-1}$. In particular, let $\Lambda_K = (\mathrm{Log}\, \varepsilon_i)_{1\le i\le \nu}$ be any $\mathbb{Z}$-basis of $\mathrm{Log}\, \mathcal{O}_K^\times$. Since for any $\varepsilon \in \mathcal{O}_K^\times$, $\mathrm{Log}_{\mathcal{S}}\, \varepsilon$ is uniformly zero on coordinates corresponding to finite places, the shape of $\Lambda_{K,\mathcal{S}}$ is:

$$\Lambda_{K,\mathcal{S}} := \left[\begin{array}{c|c} \Lambda_K & 0 \\ \hline \mathrm{Log}\, \varepsilon_{\nu+1} & \\ \vdots & \Big(-v_{\mathfrak{p}_j}(\varepsilon_{\nu+i}) \ln \mathcal{N}(\mathfrak{p}_j)\Big)_{1\le i,j\le k} \\ \mathrm{Log}\, \varepsilon_{\nu+k} & \end{array}\right]. \tag{2.3}$$

Actually, we shall use that for any maximal set of independent $\mathcal{S}$-units, their images under any logarithmic $\mathcal{S}$-embedding form a full-rank sublattice of the corresponding log-$\mathcal{S}$-unit lattice.

Expanded log-$\mathcal{S}$-embeddings.

As mentioned in [PHS19a, BDPW20], a convenient trick in the context of the cryptanalysis of id-Svp is to consider an *expanded* version of the log-$\mathcal{S}$-embedding, halving and repeating twice $\mathcal{S}_\infty$-coordinates corresponding to complex embeddings, namely, for any $\alpha \in K^\times$:

$$\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha = \Big(\big\{\ln|\sigma_i(\alpha)|\big\}_{i\in[\![1,r_1]\!]}, \big\{\ln|\sigma_{r_1+j}(\alpha)|, \ln|\overline{\sigma}_{r_1+j}(\alpha)|\big\}_{j\in[\![1,r_2]\!]}, \big\{-v_{\mathfrak{p}}(\alpha)\cdot\ln\mathcal{N}(\mathfrak{p})\big\}_{\mathfrak{p}\in\mathrm{FB}}\Big).$$

As for $j \in [\![1, r_2]\!]$, $|\alpha|_{\sigma_{r_1+j}} = |\alpha|_{\overline{\sigma}_{r_1+j}}$, the image $\overline{\mathrm{Log}}\,K^\times$ in $\mathbb{R}^n$ spans the (r_1+r_2)-dimensional space $\mathcal{L}_0 = \big\{\mathbf{y}\in\mathbb{R}^n;\ y_{r_1+2j-1} = y_{r_1+2j}, j\in[\![1,r_2]\!]\big\}$. Similarly, the image $\overline{\mathrm{Log}}_{\mathcal{S}}\,K^\times$ in $\mathbb{R}^{n+k}$ spans the (r_1+r_2+k)-dimensional space $\mathcal{L} = \mathcal{L}_0\times\mathbb{R}^k$. For convenience, we denote by H_0 (resp. H) the span of the log-unit (resp. log-$\mathcal{S}$-unit) lattice under these expanded embeddings, i.e., $H_0 = \mathcal{L}_0\cap\mathbb{R}_0^n$ and $H = \mathcal{L}\cap\mathbb{R}_0^{n+k}$.

In particular, we shall see in Pr. 2.8 that using these expanded log-$\mathcal{S}$-embeddings reduces the volume of the log-$\mathcal{S}$-unit lattice. In practice though, we did not observe any significant difference between the approximation factors obtained using $\mathrm{Log}_{\mathcal{S}}$ or $\overline{\mathrm{Log}}_{\mathcal{S}}$.

2.1.4 Regulators

The ($\mathcal{S}$-)regulator of K quantifies the density of the ($\mathcal{S}$-)unit group in K. We begin by a technical linear algebra lemma, whose result reveals particularly useful for the volume computations of non-square matrices involved in this thesis, e.g., of $\Lambda_{K,\mathcal{S}}$.

Lemma 2.4. *Let $n \geq 1$ and $a_1,\ldots,a_n \in \mathbb{R}^*$. Then, with $\mathbf{1}_{n\times n}$ being the square matrix of dimension n filled with 1's, and $\mathcal{D}_{a_1,\ldots,a_n}$ the diagonal matrix with coefficients a_i:*

$$\det\big(\mathbf{1}_{n\times n} + \mathcal{D}_{a_1,\ldots,a_n}\big) = \Big(1 + \sum_{i=1}^{n}\frac{1}{a_i}\Big)\cdot\prod_{k=1}^{n} a_k.$$

Note that the result is also valid if any of the a_i's is zero by expanding the formula and using the formal simplification $a_i/a_i = 1$. Writing it down in this form would only be much more noisy.

Proof. We prove the result for any $a_1,\ldots,a_n\in\mathbb{R}$ by induction using the minor expansion formula on the last column for the determinant. Let $M[a_1,\ldots,a_n] := \mathbf{1}_{n\times n} + \mathcal{D}_{a_1,\ldots,a_n}$, and let $\delta_{j,n}$ be its (j,n)-minor. The result is obviously true for $n=1$ using $\det M[a_1] = 1+a_1 = a_1(1+1/a_1)$, the last equality being valid for $a_1\neq 0$.

Suppose the result true for matrices of dimension $(n-1)$. The minors $\delta_{j,n}$, for $j\in[\![1,n-1]\!]$ are determinants of matrices $M[a_1,\ldots,a_{j-1},a_{j+1},\ldots,a_{n-1},0]$ whose columns are permuted by a permutation of sign $(-1)^{n-j+1}$. Using the induction hypothesis (with $\prod_\emptyset = 1$ for $n=2$):

$$\forall j\in[\![1,n-1]\!],\qquad \delta_{j,n} = (-1)^{n-j+1}\prod_{\substack{1\leq k\leq n-1\\ k\neq j}} a_k.$$

Meanwhile, the last minor $\delta_{n,n}$ is $\det M[a_1,\ldots,a_{n-1}]$, which we expand to avoid divisions by 0:

$$\delta_{n,n} = \prod_{1\leq k\leq n-1} a_k + \sum_{j=1}^{n-1}\prod_{\substack{1\leq k\leq n-1\\ k\neq j}} a_k.$$

Finally, the determinant of $M[a_1,\dots,a_n]$ is $\big((1+a_n)\delta_{n,n}+\sum_{j=1}^{n-1}(-1)^{n-j}\delta_{j,n}\big)$. A bit of calculation yields the following equation, which is the developed form of the lemma's formula:

$$\det M[a_1,\dots,a_n] = \prod_{1\le k\le n} a_k + \sum_{1\le i\le n}\prod_{\substack{1\le k\le n\\ k\neq i}} a_k. \qquad \square$$

Definition 2.5 ($\mathcal{S}$-regulator)**.** The $\mathcal{S}$-regulator of K with respect to $\mathcal{S}$, written $R_{K,\mathcal{S}}$, is defined as the absolute value of any of the (r_1+r_2+k) minors of $\Lambda_{K,\mathcal{S}}$, i.e., as the absolute value of the determinant of $\Lambda_{K,\mathcal{S}}^{(j)}$ for any $j\in[\![1,\sharp\mathcal{S}]\!]$, where $\Lambda_{K,\mathcal{S}}^{(j)}$ is the submatrix of $\Lambda_{K,\mathcal{S}}$ without the j-th coordinate.

We stress that the $\mathcal{S}$-regulator could not be consistently defined anymore if the twistings by the $\ln\mathcal{N}(\mathfrak{p})$'s were removed from the log-$\mathcal{S}$-embedding definition, as in this case, the property that all columns sum to 0 disappears.

The value of the $\mathcal{S}$-regulator $R_{K,\mathcal{S}}$ is linked to the classical regulator R_K of K (obtained for $\mathcal{S}=\mathcal{S}_\infty$) according to the following proposition:

Proposition 2.6. *Let $h_{K,(\mathrm{FB})}$ the cardinal of the subgroup $\mathrm{Cl}_K^{(\mathrm{FB})}$ of Cl_K generated by classes of ideals in $\mathrm{FB}=\mathcal{S}\cap\mathcal{S}_0$. Then, the $\mathcal{S}$-regulator $R_{K,\mathcal{S}}$ verifies:*

$$R_{K,\mathcal{S}} = h_{K,(\mathrm{FB})}R_K\cdot\prod_{\mathfrak{p}\in\mathrm{FB}}\ln\mathcal{N}(\mathfrak{p}).$$

Proof. Note that $R_{K,\mathcal{S}}$ is the determinant, e.g., of $\Lambda_{K,\mathcal{S}}^{(r_1+r_2)}$ where the (r_1+r_2)-th column is removed, so is the product of $\det\Lambda_K^{(r_1+r_2)}=R_K$ and of the determinant of the (unchanged) square bottom right part of $\Lambda_{K,\mathcal{S}}$. By definition of $\mathcal{O}_{K,\mathcal{S}}^\times$, the matrix $\big(-v_{\mathfrak{p}_j}(\varepsilon_{\nu+i})\big)_{i,j}$ generates the lattice of all relations in Cl_K between ideals of FB, i.e., is the kernel of the following map:

$$\mathfrak{f}_{\mathrm{FB}}:(e_1,\dots,e_k)\in\mathbb{Z}^k\longmapsto\textstyle\prod_j\big[\mathfrak{p}_j\big]^{e_j}\in\mathrm{Cl}_K,$$

whose image is precisely $\mathrm{Cl}_K^{(\mathrm{FB})}$. Thus, $\det(\ker\mathfrak{f}_{\mathrm{FB}})$ is $h_{K,(\mathrm{FB})}=\sharp\big(\mathbb{Z}^k/\ker\mathfrak{f}_{\mathrm{FB}}\big)$, and twisting each column by $\ln\mathcal{N}(\mathfrak{p})$ for $\mathfrak{p}\in\mathrm{FB}$ yields the result. $\square$

Log-$\mathcal{S}$-unit lattice volumes.

The volume of the log-$\mathcal{S}$-unit lattice is tied to the $\mathcal{S}$-regulator $R_{K,\mathcal{S}}$ by the following proposition, which generalizes the classical formula (see e.g., [Neu99, Pr. I.7.5]) linking R_K to $\mathrm{Vol}\big(\mathrm{Log}\,\mathcal{O}_K^\times\big)$:

Proposition 2.7.

$$\mathrm{Vol}\big(\mathrm{Log}_{\mathcal{S}}\,\mathcal{O}_{K,\mathcal{S}}^\times\big)=\sqrt{1+\nu+k}\cdot R_{K,\mathcal{S}}.$$

Proof. By definition, $\mathrm{Vol}\big(\mathrm{Log}_{\mathcal{S}}\,\mathcal{O}_{K,\mathcal{S}}^\times\big)=\sqrt{\det(\Lambda_{K,\mathcal{S}}\Lambda_{K,\mathcal{S}}^{\mathrm{T}})}$. Consider $\Lambda_{K,\mathcal{S}}^{(r_1+r_2+k)}$, removing the last coordinate, whose determinant is $R_{K,\mathcal{S}}$. The concatenated matrix $P=\big(I_{\nu+k}\big\|-\mathbf{1}_{\nu+k}\big)$ verifies $\Lambda_{K,\mathcal{S}}=\Lambda_{K,\mathcal{S}}^{(r_1+r_2+k)}\cdot P$, and a simple induction shows that $\det(PP^{\mathrm{T}})=1+\nu+k$ (use Lem. 2.4 with all a_i's equal to 1). $\square$

Using expanded log-$\mathcal{S}$-embeddings impacts the volume of the log-$\mathcal{S}$-unit lattices given in Pr. 2.7. It is given in following proposition, which generalizes [BDPW20, Lem. A.1]:

Proposition 2.8. *Under the expanded log-$\mathcal{S}$-embedding, the log-$\mathcal{S}$-unit lattice has volume:*

$$\mathrm{Vol}\big(\overline{\mathrm{Log}}_{\mathcal{S}}\,\mathcal{O}_{K,\mathcal{S}}^{\times}\big) = \sqrt{n+k}\cdot 2^{-r_2/2}\cdot R_{K,\mathcal{S}}.$$

Using an empty factor basis, this implies $\mathrm{Vol}\big(\overline{\mathrm{Log}}\,\mathcal{O}_K^{\times}\big) = \sqrt{n}\cdot 2^{-r_2/2}\cdot R_K$.

Proof. Let $\widetilde{\Lambda}_{K,\mathcal{S}}$ be a row basis of $\overline{\mathrm{Log}}_{\mathcal{S}}\,\mathcal{O}_{K,\mathcal{S}}^{\times}$, whose shape is the same as $\Lambda_{K,\mathcal{S}}$ in Eq. (2.3) except that $\overline{\mathrm{Log}}$ is systematically used instead of Log. The proof explicits the transition matrix from the truncated matrix $\Lambda_{K,\mathcal{S}}^{(\nu+1+k)}$, whose determinant is $R_{K,\mathcal{S}}$, to $\widetilde{\Lambda}_{K,\mathcal{S}}$, and computes its volume.

Let $P = \big(I_{\nu+k}\big\|-\mathbf{1}_{\nu+k}\big)$ be such that $\Lambda_{K,\mathcal{S}} = \Lambda_{K,\mathcal{S}}^{(r_1+r_2+k)}\cdot P$. Obtaining $\widetilde{\Lambda}_{K,\mathcal{S}}$ from $\Lambda_{K,\mathcal{S}}$ requires to halve and expand the coordinates corresponding to complex places, all other coordinates staying identical. Let F be the transition matrix verifying $\widetilde{\Lambda}_{K,\mathcal{S}} = \Lambda_{K,\mathcal{S}}\cdot F$, i.e., the block diagonal matrix with three blocks: I_{r_1}, the $(r_2\times 2r_2)$ block of vectors $(\dots,1/2,1/2,\dots)$, and I_k. Then $\widetilde{\Lambda}_{K,\mathcal{S}} = \Lambda_{K,\mathcal{S}}^{(r_1+r_2+k)}\cdot(PF)$. For $k\geq 1$, or $k=0$ and $r_2=0$, (PF) writes as $\big(F_{-1}\big\|-\mathbf{1}_{\nu+k}\big)$, where F_{-1} is F without its last column *and* its last row. We compute:

$$(PF)(PF)^{\mathrm{T}} = \mathbf{1}_{(\nu+k)\times(\nu+k)} + \mathcal{D}_{(\mathbf{1}_{r_1}\|(1/2)\cdot\mathbf{1}_{r_2}\|\mathbf{1}_{k-1})}.$$

Using Lem. 2.4 to obtain that the determinant of this matrix is $(n+k)2^{-r_2}$ completes the proof, except in the case $k=0$, $r_2>0$. In this specific case, (PF) writes as the first $(n-2)$ columns of F_{-1}, concatenated twice with $(-1/2)\cdot\mathbf{1}_\nu$, so that $(PF)(PF)^{\mathrm{T}} = \frac{1}{2}\cdot\big(\mathbf{1}_{\nu\times\nu} + \mathcal{D}_{(2\cdot\mathbf{1}_{r_1}\|\mathbf{1}_{r_2-1})}\big)$. This last matrix has volume $n\cdot 2^{-r_2}$ as expected. □

2.2 Cyclotomic Fields

An important special case of number fields is the family of *cyclotomic fields*, for which many additional properties are known.

For any positive integer $m>1$, we denote the cyclotomic field of conductor m, or the m-th cyclotomic field, by $K_m = \mathbb{Q}(\zeta_m)$, where $\zeta_m = e^{2i\pi/m}$ is a primitive m-th root of unity. It has degree $n=\varphi(m)$, its maximal order is $\mathcal{O}_{K_m} = \mathbb{Z}[\zeta_m]$ ([Was97, Th. 2.6]), and its discriminant, which has the same order of magnitude as n^n, is given precisely by ([Was97, Pr. 2.7]):

$$\Delta_{K_m} = (-1)^{\varphi(m)/2}\frac{m^{\varphi(m)}}{\prod_{p\mid m}p^{\varphi(m)/(p-1)}}$$

For convenience, when it comes to cyclotomic fields, we will index ideal and class groups as well as class numbers by the cyclotomic field conductor, i.e., by m instead of K_m. Hence, the multiplicative group of fractional ideals of K_m is denoted by $\mathcal{I}_m$ instead of $\mathcal{I}_{K_m}$; likewise, the normal subgroup of principal ideals is written $\mathcal{P}_m := \big\{\langle\alpha\rangle;\ \alpha\in K_m\big\}$. The class group of K_m is written Cl_m, and the class number is simply denoted by h_m instead of h_{K_m}.

Note that if m is odd, we have that $K_m = K_{2m}$, so we can further assume $m\not\equiv 2 \mod 4$ without any loss of generality. We shall also write the prime factorization of m as $m = p_1^{e_1}p_2^{e_2}\cdots p_t^{e_t}$ and let $q_i = p_i^{e_i}$ for all $i\in[\![1,t]\!]$. In particular, m has exactly t distinct prime divisors.

Remark 2.9. Note that we implicitly fix an ordering on the factors q_i of m. All our results hold true for any ordering as long as it stays consistent through all subsets of the q_i's. However, if this ambiguity was a problem in an application, we could simply fix an ordering by $p_1<\cdots<p_t$.

2.2.1 Two special arithmetic subsets of $[\![1, m]\!]$

We recall here from [Kuč92, p.293] the definition of two subsets M_m^+ and M_m^- of $[\![1, m]\!]$ that are useful to describe resp. a fundamental family of circular units and a short $\mathbb{Z}$-basis of the Stickelberger ideal of K_m.

Let X_m be the set of all positive integers $a < m$ that are either divisible by q_i or relatively prime to q_i for each $i \in [\![1, t]\!]$, i.e.:

$$X_m = \Big\{ a \in \mathbb{Z};\ 0 < a < m, \big(a, \tfrac{m}{(a,m)}\big) = 1 \Big\}.$$

Let $M_m^{\pm} \subseteq X_m$ be the sets of all $a \in X_m$ satisfying ([Kuč92, p.293]):[15]

- for all $i \in [\![1, t]\!]$, if $q_i \nmid a$ then $a \not\equiv -(a, m) \mod q_i$,
- if $a \nmid m$, let $k = \max\big\{ i \in [\![1, t]\!];\ a \not\equiv (a, m) \mod q_i \big\}$, then $\big\{ \frac{a}{(a,m)q_k} \big\} < \frac{1}{2}$,
- if $a \mid m$ then the set $\big\{ i \in [\![1, t]\!];\ q_i \nmid a \big\}$ has an even (resp. odd) number of elements when defining M_m^+ (resp. when defining M_m^-).

Note that M_m^+ (resp. M_m^-) contains $\frac{\varphi(m)}{2} - 1$ elements (resp. $\frac{\varphi(m)}{2}$ elements). Both sets are obviously easy to compute, using only simple arithmetic criteria.

2.2.2 Galois group and maximal real subfield

Let G_m denote the Galois group of K_m, which can be written explicitly as ([Was97, Th. 2.5]):

$$G_m = \big\{ \sigma_{m,s} : \zeta_m \longmapsto \zeta_m^s;\ 0 < s < m, (s, m) = 1 \big\} \simeq \big(\mathbb{Z}/m\mathbb{Z}\big)^\times.$$

In particular, we denote by $\sigma_{m,s} \in G_m$ the automorphism sending any m-th root of unity to its s-th power. For convenience, the automorphism induced by complex conjugation is written $\tau = \sigma_{m,-1}$, and we will omit m most of the time, when no ambiguity is possible.

The algebraic norm of $\alpha \in K_m$ is defined by $\mathcal{N}(\alpha) = \prod_{\sigma \in G_m} \sigma(\alpha)$, hence the absolute norm element in the integral group ring $\mathbb{Z}[G_m]$ writes as $N_m = \sum_{\sigma \in G_m} \sigma$.

For any positive integers m, r such that $r \mid m$ we have the usual restriction and corestriction maps between the group rings $\mathbb{Q}[G_m]$ and $\mathbb{Q}[G_r]$:

$$\begin{aligned} \mathrm{Res}_{K_m/K_r} &: \mathbb{Q}[G_m] \to \mathbb{Q}[G_r], \\ \mathrm{Cor}_{K_m/K_r} &: \mathbb{Q}[G_r] \to \mathbb{Q}[G_m]. \end{aligned}$$

The restriction map is the ring homomorphism sending each automorphism $\sigma \in G_m$ to its restriction $\sigma|_{K_r}$; the corestriction map is the linear map determined for any $\rho \in G_r$ by:

$$\mathrm{Cor}_{K_m/K_r}(\rho) = \sum_{\substack{\sigma \in G_m \\ \sigma|_{K_r} = \rho}} \sigma.$$

The integral group ring $\mathbb{Z}[G_m]$ acts naturally on $\mathcal{I}_m$; more precisely, for any $\mathfrak{b} \in \mathcal{I}_m$ and any element $\alpha = \sum_{\sigma \in G_m} a_\sigma \sigma \in \mathbb{Z}[G_m]$, we write $\mathfrak{b}^\alpha := \prod_{\sigma \in G_m} \sigma\big(\mathfrak{b}\big)^{a_\sigma}$.

[15] Actually, the set M_+ defined in [Kuč92, p.293] is $M_+ = M_m^+ \cup \{0\}$.

Maximal real subfield.

The maximal real subfield of K_m, denoted by K_m^+, is the fixed subfield of K_m under complex conjugation, i.e., $K_m^+ := K_m^{\langle\tau\rangle} = \mathbb{Q}\big(\zeta_m+\zeta_m^{-1}\big)$. Its maximal order is given by $\mathcal{O}_{K_m^+} = \mathbb{Z}\big[\zeta_m+\zeta_m^{-1}\big]$ (see e.g., [Was97, Pr. 2.16]).

By Galois theory, since $\langle\tau\rangle$ is a normal subgroup of G_m, the maximal real subfield of K_m is a Galois extension of $\mathbb{Q}$ with Galois group $G_m^+ := \mathrm{Gal}\big(K_m^+/\mathbb{Q}\big)$ isomorphic to $G_m/\langle\tau\rangle$. We will consistently identify G_m^+ with the following system of representatives modulo τ restricted to K_m^+:

$$G_m^+ = \big\{\sigma_s|_{K_m^+};\ 0 < s < \tfrac{m}{2}, (s,m) = 1\big\}.$$

Technically, each $\sigma_s|_{K_m^+} \in G_m^+$ extends in G_m to either σ_s or $\tau\sigma_s = \sigma_{-s}$. For simplicity, we always choose to lift $\sigma_s|_{K_m^+} \in G_m^+$ to $\sigma_s \in G_m$ and drop the restriction to K_m^+ which should be clear from the context. This slight abuse of notation appears to be very practical. For example, the corestriction $\mathrm{Cor}_{K_m/K_m^+}\big(\sigma_s|_{K_m^+}\big)$, defined as the sum of all elements of G_m that restricts to $\sigma_s|_{K_m^+}$, namely $\sigma_s + \tau\sigma_s$, is written using the much simpler expression $(1+\tau)\cdot\sigma_s$.

The class group and class number of the maximal real subfield K_m^+ are denoted respectively by Cl_m^+ and h_m^+.

2.2.3 Real and relative class groups

One important specificity of cyclotomic fields is that the real class group Cl_m^+ embeds into Cl_m *via* the natural inclusion map, which to each ideal class $\big[\mathfrak{b}\big] \in \mathrm{Cl}_m^+$ associates the lifted ideal class $\big[\mathfrak{b}\cdot\mathcal{O}_{K_m}\big] \in \mathrm{Cl}_m$ [Was97, Th. 4.14]. The relative norm map $\mathcal{N}_{K_m/K_m^+}$ induces a homomorphism from Cl_m to Cl_m^+, whose kernel is hence isomorphic to the so-called *relative class group*, written Cl_m^- and of cardinal the *relative class number* h_m^-. Thus, by construction, for any $\mathfrak{b}$ s.t. $\big[\mathfrak{b}\big] \in \mathrm{Cl}_m^-$, $\mathfrak{b}^{1+\tau} \cap K_m^+$ is principal. Concretely, it implies that $h_m = h_m^+ \cdot h_m^-$ is the product of the so-called *plus part* and *relative part* of the class number.

Relative part of the class number.

As mentioned earlier, not much is generally known about the class number of a number field, and the analytic class number formula recalled in Eq. (2.34) only allows to obtain a rough upper bound $h_m \leq \widetilde{O}\big(\sqrt{|\Delta_{K_m}|}\big)$.

In the case of cyclotomic fields though, the structure of the relative class group is better understood [FGW92]. Using analytic means, the relative class number has the following explicit expression [Was97, Th. 4.17]:

$$h_m^- = Qw \cdot \prod_{\chi \text{ odd}} \Big(-\tfrac{1}{2}B_{1,\chi}\Big), \tag{2.10}$$

where $w = 2m$ if m is odd and $w = m$ if m is even, $Q = 1$ if m is a prime power and $Q = 2$ otherwise, and $B_{1,\chi}$ is defined by $\frac{1}{f}\sum_{a=1}^{f} a\cdot\chi(a)$ for any odd primitive Dirichlet character χ modulo m of conductor f dividing m.

Computing this value is in practice very efficient, using adequate representations of Dirichlet characters. We shall also introduce in Ch. 4 an algorithmically basic way to obtain h_m^- *via* a determinant computation that is especially competitive when m has few distinct prime divisors.

Plus part of the class number.

The really hard part of cyclotomic class numbers computations is to obtain the plus part h_m^+, and few values are known. We will use the values from [Was97, Tab. §4], [Mil14, Th. 1.1 and 1.2]

m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+
225	120	1	213	140	1	205	160	2	203	168	1	460	176	1	416	192	1
231	120	1	219	144	1	352	160	1	215	168	1	552	176	1	448	192	1
244	120	1	285	144	1	400	160	1	245	168	1	209	180	1	576	192	1
248	120	4	296	144	1	440	160	5	261	168	1	217	180	1	612	192	1
308	120	1	304	144	1	492	160	1	392	168	1	279	180	1	672	192	1
372	120	1	380	144	1	528	160	1	516	168	1	297	180	1	275	200	1
396	120	1	432	144	1	600	160	1	588	168	1	235	184	1	375	200	1
384	128	1	444	144	1	660	160	1	267	176	1	564	184	1	500	200	1
201	132	1	540	144	1	243	162	1	345	176	1	291	192	1			
207	132	1	237	156	1	249	164	1	368	176	1	357	192	1			

TABLE 2.1 – Additional (publicly unavailable) values of h_m^+ for some m with $\varphi(m) \leq 200$.

and [BFHP21, Tab. 1], consistently assuming the *Generalized Riemann Hypothesis* (GRH) (see Heur. 2.33). We also provide 58 additional values in Tab. 2.1, easily obtained using SAGEMATH v9.0 [Sag20], each in less than 3 hours on a Intel® Core™ i7-8650U @3.2GHz CPU.

The fact that the plus part of the class number seems so much smaller than the relative part is striking. On the theoretical side, Weber's conjecture claims that $h_{2^e}^+ = 1$ for any $e > 1$, and Buhler, Pomerance and Robertson [BPR04] argue, based on Cohen-Lenstra heuristics, that for all but finitely many pairs (p, e), where p is a prime and e is a positive integer, $h_{p^{e+1}}^+ = h_{p^e}^+$; hence, for prime power conductors, this conjecture claims that the plus part is asymptotically constant.

On the practical side, these conjectures are backed up by Schoof's extensive calculations [Sch03] in the prime conductor case, and by the above explicit values. In particular, under GRH, Miller proved Weber's conjecture up to $m = 512$, and we note that according to Schoof's table, the inequality $h_m^+ \leq \sqrt{m}$ holds for more than 96.6% of all prime conductors $m = p < 10000$.

2.2.4 Circular units

Circular units are sometimes called *cyclotomic units* in the literature, as in [Was97, §8]. We prefer to use the historical terminology from algebraic number theory, see e.g., SINNOTT [Sin78, §4] and KUČERA [Kuč92, §2], in order to avoid any confusion with the whole unit group $\mathcal{O}_{K_m}^\times$ of the m-th cyclotomic field.

Definition 2.11 (Circular units [Was97, §8.1])**.** Let V_m be the multiplicative subgroup of $K_m^\times$ generated by:

$$\{1 - \zeta_m^a;\ 1 \leq a \leq m\}.$$

The group of *circular units* is the intersection $C_m := V_m \cap \mathcal{O}_{K_m}^\times$.

Note that C_m contains the torsion of K_m, since $-\zeta_m = (1 - \zeta_m)/(1 - \zeta_m^{-1})$. The circular units form a subgroup of $\mathcal{O}_{K_m}^\times$ of finite index, more precisely:

Proposition 2.12 ([Sin78, Th. p.107])**.** *The index of C_m in $\mathcal{O}_{K_m}^\times$ is finite:*

$$\left[\mathcal{O}_{K_m}^\times : C_m\right] = 2^b \cdot h_m^+, \qquad \text{with } b = \begin{cases} 0 & \text{if } t = 1, \\ 2^{t-2} + 1 - t & \text{otherwise.} \end{cases}$$

where t is the number of distinct prime factors of m.

Hence, circular units provide a very large subgroup of $\mathcal{O}_{K_m}^\times$: indeed, the real part of the class number is expected to be small (§2.2.3), and the other factor *generically* grows linearly in m (see [HW38, Th. 430 and 431] for a precise statement).

An explicit system of fundamental circular units for any m has been given in [GK89] and independently in [Kuč92, Th. 6.1]. More precisely, for $0 < a < m$, define the following special circular units, where $m_i = m/p_i^{e_i}$ [Kuč92, p.176]:

$$v_a = \begin{cases} 1 - \zeta_m^a & \text{if } \forall i \in [\![1, t]\!],\ m_i \nmid a, \\ \dfrac{1-\zeta_m^a}{1-\zeta_m^{m_i}} & \text{otherwise, for the unique } m_i \mid a. \end{cases} \tag{2.13}$$

Theorem 2.14 ([Kuč92, Th. 6.1])**.** *The set $\{v_a;\ a \in M_m^+\}$ is a system of fundamental circular units of K_m: for any circular unit $\eta \in C_m$, there exist uniquely determined $k(a) \in \mathbb{Z}$ and root of unity $\mu \in \langle \pm\zeta_m \rangle$ s.t. $\eta = \mu \cdot \prod_{a \in M_m^+} v_a^{k(a)}$.*

A crucial point for the cryptanalysis of id-Svp in [CDW21] is that the logarithmic embedding of these elements is short. Namely, computing explicitly the constants that appear in the proof of [CDW21, Lem. 3.5], we have, for any $0 < a < m$, that $\|\mathrm{Log}(1 - \zeta_m^a)\|_2 \leq 1.32 \cdot \sqrt{m}$.

2.2.5 Stickelberger ideal

In this section, we decribe the Stickelberger ideal of a cyclotomic field K_m, that provides free relations in the class group. Following Sinnott [Sin80], for any $a \in \mathbb{Z}$, let:

$$\theta_m(a) = \sum_{s \in (\mathbb{Z}/m\mathbb{Z})^\times} \left\{-\frac{as}{m}\right\} \cdot \sigma_{m,s}^{-1} \in \mathbb{Q}[G_m], \tag{2.15}$$

and let N_m be the absolute norm element $N_m = \sum_{\sigma \in G_m} \sigma$. Hence, an easy observation gives:

$$a \equiv b \pmod m \qquad \Longrightarrow \qquad \theta_m(a) = \theta_m(b). \tag{2.16}$$

Moreover, if $m \mid a$, $\theta_m(a) = 0$, whereas if $m \nmid a$ we get the following relation:

$$\theta_m(a) + \theta_m(-a) = N_m. \tag{2.17}$$

Definition 2.18 (Stickelberger ideal [Sin80, p.189])**.** Let $\mathcal{S}'_m$ be the subgroup of the additive group of $\mathbb{Q}[G_m]$ generated by $\{\theta_r^{(m)}(a);\ a, r \in \mathbb{Z}, r > 0\}$, where:

$$\theta_r^{(m)}(a) = \mathrm{Cor}_{K_m/K_{(m,r)}}\Big(\mathrm{Res}_{K_r/K_{(m,r)}}\big(\theta_r(a)\big)\Big).$$

The *Stickelberger ideal* of K_m is the intersection $\mathcal{S}_m = \mathcal{S}'_m \cap \mathbb{Z}[G_m]$.

In fact, $\mathcal{S}'_m$ is Sinnott's group S' from [Sin80, p.189], for the abelian field k being the cyclotomic field K_m. The following lemma allows us to simplify the previous definition.

Lemma 2.19. *For any positive integer m, the group $\mathcal{S}'_m$ is the subgroup of $\mathbb{Q}[G_m]$ generated by*

$$\{\theta_m(a);\ 0 < a < m\} \cup \{\tfrac{1}{2} N_m\}.$$

Proof. On one hand, $\theta_m(a) = \theta_m^{(m)}(a) \in \mathcal{S}'_m$. On the other hand, let us consider any positive integer $r \neq m$ and let $d = (m, r)$. For any $a \in \mathbb{Z}$, using [Kuč96, Lem. 12],

$$\mathrm{Res}_{K_r/K_d}\big(\theta_r(a)\big) \in \Big\langle \big\{\theta_d(b);\ 0 < b < d\big\} \cup \big\{\tfrac{1}{2}N_d\big\} \Big\rangle.$$

It is easy to see that $\mathrm{Cor}_{K_m/K_d}\big(\frac{1}{2}N_d\big) = \frac{1}{2}N_m$. Considering $\theta_d(b)$, $0 < b < d$,

$$\begin{aligned}\mathrm{Cor}_{K_m/K_d}\big(\theta_d(b)\big) &= \mathrm{Cor}_{K_m/K_d}\Big(\sum_{\substack{0<s\le d\\(s,d)=1}} \Big\{-\frac{bs}{d}\Big\} \cdot \sigma_{d,s}^{-1}\Big) \\ &= \sum_{\substack{0<s\le m\\(s,m)=1}} \Big\{-\frac{bs}{d}\Big\} \cdot \sigma_{m,s}^{-1} = \theta_m\Big(\frac{bm}{d}\Big). \end{aligned} \tag{2.20}$$

As Cor_{K_m/K_d} is a group homomorphism, the lemma follows from

$$\theta_r^{(m)}(a) = \mathrm{Cor}_{K_m/K_d}\Big(\mathrm{Res}_{K_r/K_d}\big(\theta_r(a)\big)\Big) \in \Big\langle \big\{\theta_m(a);\ 0 < a < m\big\} \cup \big\{\tfrac{1}{2}N_m\big\} \Big\rangle. \qquad \square$$

Remark 2.21. For clarity, let us explain that even though $\mathcal{S}'_m$ is slightly different from Sinnott's group S' from [Sin78], the Stickelberger ideal $S = S' \cap \mathbb{Z}[G_m]$ defined in [Sin78] coincides with $\mathcal{S}_m$. Indeed, S' is defined as the subgroup of $\mathbb{Q}[G_m]$ generated by the set $\big\{\theta_m(a);\ 0 < a < m\big\}$, so Lem. 2.19 implies that $\mathcal{S}'_m = S' + \frac{1}{2}N_m \cdot \mathbb{Z}$. If m is even, then $\frac{1}{2}N_m = \theta_m(\frac{m}{2}) \in S'$, which implies $\mathcal{S}'_m = S'$. Let us suppose that m is odd. Then all generators of S' have 2-integral coefficients, so $\frac{1}{2}N_m \notin S'$, but we have that $N_m = \theta_m(1) + \theta_m(-1) \in S'$. Any $\beta \in \mathcal{S}'_m$ can be written as $\beta = \alpha + k \cdot \frac{1}{2}N_m$, for some $\alpha \in S'$ and $k \in \mathbb{Z}$. If $\beta \in \mathbb{Z}[G_m]$, then the fact that the coefficients of α are 2-integral implies that k is even, which means that $\beta \in S'$. Hence, in this case we also have $\mathcal{S}_m = \mathcal{S}'_m \cap \mathbb{Z}[G_m] = S' \cap \mathbb{Z}[G_m] = S$.

As in [CDW21], we shall refer to the *Stickelberger lattice* when $\mathcal{S}_m$ is viewed as a $\mathbb{Z}$-module. Note that in some references, like in [Was97, §6.2], the Stickelberger ideal is defined as the smaller ideal $\mathbb{Z}[G_m] \cap \theta_m(-1)\mathbb{Z}[G_m]$, which coincides with Def. 2.18 if and only if m is a prime power [Kuč86, Pr. 4.3].

One of the most important feature of the Stickelberger ideal is to give free relations in the class group of K_m, as stated by Stickelberger's theorem, given below.

Theorem 2.22 (Stickelberger's theorem [Sin80, Th. 3.1])**.** *The Stickelberger ideal $\mathcal{S}_m$ of K_m annihilates the class group of K_m. Hence, for any ideal $\mathfrak{b}$ of K_m and any $\alpha = \sum_{\sigma\in G_m} a_\sigma \sigma \in \mathcal{S}_m$, the ideal $\mathfrak{b}^\alpha = \prod_{\sigma\in G_m} \sigma(\mathfrak{b})^{a_\sigma}$ is principal.*

An outstanding point is that the proof of this important result is completely explicit, i.e., for any $\alpha \in \mathcal{S}_m$, and any fractional ideal $\mathfrak{b}$ of K_m, an explicit $\gamma \in K_m$ such that $\langle\gamma\rangle = \mathfrak{b}^\alpha$ is constructed. We shall see in §4.5 that when α is a short element of $\mathcal{S}_m$, i.e., when $\alpha = \sum_{\sigma\in G_m} \varepsilon_\sigma \sigma$ with all $\varepsilon_\sigma \in \{0, 1\}$, this explicit generator is very efficiently computable.

On the rank of the Stickelberger lattice.

A consequence of, e.g., [Kuč92, Th. 6.2], is that the rank of $\mathcal{S}_m$ in $\mathbb{Z}[G_m]$, viewed as a $\mathbb{Z}$-module, is only $\varphi(m)/2 + 1$; in particular, it is not full rank, therefore it cannot be directly used as a lattice of class relations.

However, as noted in [CDW21, §4.3], the Stickelberger lattice modulo $(1 + \tau)$ is a lattice of class relations for the relative class group, which we recall is the kernel of the relative norm

map $\mathcal{N}_{K_m/K_m^+} : \mathrm{Cl}_m \to \mathrm{Cl}_m^+$. We shall follow a quite different exposition here, using Sinnott's formalism from [Sin78, Sin80].

Let $\mathcal{R}_m = \mathbb{Z}\big[G_m\big]$. For any submodule $M \subseteq \mathcal{R}_m$, the kernel of the multiplication by $(1+\tau)$ in M is denoted by M^-. In particular:

$$\mathcal{R}_m^- = \big\{\alpha \in \mathcal{R}_m;\ (1+\tau)\alpha = 0\big\} \quad \text{and} \quad \mathcal{S}_m^- = \big\{\alpha \in \mathcal{S}_m;\ (1+\tau)\alpha = 0\big\}.$$

Clearly, we have $\mathcal{R}_m^- = (1-\tau)\mathcal{R}_m$ and $(1-\tau)\mathcal{S}_m \subsetneq \mathcal{S}_m^-$. Let $\pi : \mathcal{R}_m \longrightarrow \mathcal{R}_m^-$ be the natural projection that associates $(1-\tau)\alpha \in \mathcal{R}_m^-$ to any $\alpha \in \mathcal{R}_m$. A basis of $\mathcal{R}_m^-$, as a $\mathbb{Z}$-module, is given by [Kuč86, Th. 3.1]:

$$\big\{\beta_s;\ 0 < s < \tfrac{m}{2}, (s,m) = 1\big\}, \quad \text{where } \beta_s = \pi\big(\sigma_s\big) = \sigma_s - \sigma_{-s}. \tag{2.23}$$

Hence, $\mathcal{R}_m^-$ is isomorphic, as a $\mathbb{Z}$-module, to $\mathbb{Z}^{\varphi(m)/2}$. Note that the map π defined above corresponds to the projection map $\mathcal{R}_m \to \mathcal{R}_m\big/\big\langle 1+\tau\big\rangle$ of [CDW21], as shown by the expression given in the proof of [CDW21, Lem. 4.6].

Theorem 2.24 ([Sin78, Th. p.107])**.** *The index of $\mathcal{S}_m^-$ in $\mathcal{R}_m^-$ is finite:*

$$\big[\mathcal{R}_m^- : \mathcal{S}_m^-\big] = 2^a \cdot h_m^-, \quad \textit{where } a = \begin{cases} 0 & \textit{if } t = 1, \\ 2^{t-2} - 1 & \textit{if } t \geq 2. \end{cases}$$

In particular, $\mathcal{S}_m^-$ has full rank $\frac{\varphi(m)}{2}$ in $\mathcal{R}_m^-$. The restriction to the relative class group means that the action of $(1+\tau)$ factors through the projection in $\mathcal{S}_m^-$, hence $\mathcal{S}_m^-$ can be used as a lattice of class relations for G_m-orbits of Cl_m^-.

Remark 2.25. We note that the projected Stickelberger lattice $(1-\tau)\mathcal{S}_m$ used in [CDW21] is *strictly smaller* than $\mathcal{S}_m^- = \mathcal{S}_m \cap \mathcal{R}_m^-$. In fact, a consequence of the proof of Lem. 5.15 is that its index is $\big[\mathcal{S}_m^- : (1-\tau)\mathcal{S}_m\big] = 2^{\varphi(m)/2-1}$.

Technical lemmata on $\mathcal{S}_m$.

In this paragraph, we recall some technical results that will be useful mainly to explicit and prove the correctness of our short basis of the Stickelberger ideal. First, the index of the Stickelberger ideal in $\mathcal{S}'_m$ is related to torsion units as follows:

Lemma 2.26. *For any integer $m > 1$, $m \not\equiv 2 \pmod 4$, the index $w = [\mathcal{S}'_m : \mathcal{S}_m]$ is equal to the number of roots of unity in the m-th cyclotomic field K_m, i.e., $w = 2m$ if m is odd, and $w = m$ if m is even.*

Proof. This is a part of [Sin80, Pr. 2.1]. □

We now introduce auxiliary elements that allow to write relations that are useful for the proof of Th. 4.2. For any $a \in \mathbb{Z}$, we set

$$\omega_m(a) = \begin{cases} \theta_m(a) - \frac{1}{2}N_m, & \text{if } m \nmid a, \\ 0, & \text{if } m \mid a. \end{cases} \tag{2.27}$$

Adapting Eqs. (2.16), (2.17) and (2.20), we deduce respectively, for $d \mid m$ and $0 < b < d$,

$$\omega_m(a+m) = \omega_m(a) \qquad \text{and} \qquad \omega_m(-a) = -\omega_m(a), \tag{2.28}$$

$$\mathrm{Cor}_{K_m/K_d}\big(\omega_d(b)\big) = \mathrm{Cor}_{K_m/K_d}\big(\theta_d(b) - \tfrac{1}{2}N_d\big) = \omega_m\big(\tfrac{bm}{d}\big). \tag{2.29}$$

The last equality uses that Cor_{K_m/K_d} is a linear map and $\mathrm{Cor}_{K_m/K_d}(N_d) = N_m$. Moreover, by Lem. 2.19, $\mathcal{S}'_m$ is the subgroup of $\mathbb{Q}[G_m]$ generated by

$$\{\omega_m(a);\ 0 < a < m\} \cup \{\tfrac{1}{2}N_m\}. \tag{2.30}$$

Lemma 2.31. *Let d, r be positive integers and $m = rd$. Then for any $k \in \mathbb{Z}$ we have*

$$\sum_{\substack{a=0,\dots,m-1 \\ a \equiv k \pmod r}} \omega_m(a) = \sum_{i=0}^{d-1} \omega_m(k+ir) = \omega_m(kd).$$

Proof. The lemma follows from the following well-known identity

$$\sum_{i=0}^{d-1} \left\{ -\frac{s(k+ir)}{m} \right\} = \left\{ -\frac{skd}{m} \right\} + \frac{d-1}{2},$$

valid for any $s \in \mathbb{Z}$ relatively prime to m. □

Recall that $m = q_1 q_2 \dots q_t$, where $q_i = p_i^{e_i} > 2$ for each $i \in [\![1, t]\!]$, is the prime factorization of m. Let $\ell_i \in \mathbb{Z}$ satisfy $p_i \ell_i \equiv 1 \pmod{\frac{m}{q_i}}$, and $\ell_i \equiv 1 \pmod{q_i}$. Lemma 2.31 implies the following result:

Lemma 2.32. *For the chosen m, for any $i \in [\![1, t]\!]$ and any $a \in X_m$, we have*

$$\sum_{\substack{k \equiv 1 \pmod{m/q_i} \\ 0 < k \le m,\ p_i \nmid k}} \omega_m(ka) = \begin{cases} \varphi(q_i) \cdot \omega_m(a), & \text{if } q_i \mid a, \\ \omega_m(aq_i) - \omega_m(aq_i\ell_i), & \text{if } q_i \nmid a, \end{cases}$$

where $\varphi()$ is Euler's totient function.

2.3 Algorithmic Number Theory

In this thesis, we will consistently assume the *Generalized Riemann Hypothesis* (GRH), on which rely many useful number-theoretic bounds and algorithmic complexities.

Heuristic 2.33 (Generalized Riemann Hypothesis (GRH))**.** *The Dedekind zeta function of K, defined for $s \in \mathbb{C} \setminus \{1\}$ as $\zeta_K(s) = \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{1}{\mathcal{N}(\mathfrak{a})^s}$ when $\Re(s) > 1$ and as its analytic continuation elsewhere, is zero-free in the half plane $\Re(s) > 1/2$.*

2.3.1 Number-theoretic bounds

This section presents several number-theoretic bounds that are useful to control namely the volume of log-$\mathcal{S}$-unit lattices, and the algebraic norm of the factor base prime ideals.

Analytic class number formula.

The residue $\kappa_K = \lim_{s \to 1}(s-1)\zeta_K(s)$ is linked to $h_K R_K$ through the so-called *analytic class number formula* [Neu99, Cor. 5.11(ii)], which states that:

$$\kappa_K = \frac{2^{r_1}(2\pi)^{r_2} R_K h_K}{w_K \sqrt{|\Delta_K|}}, \tag{2.34}$$

where $w_K = \sharp\mu(\mathcal{O}_K^\times)$. Actually, computing κ_K is much easier than computing directly h_K or R_K (see e.g., [BF15]) and is generally performed as a first step towards these quantities.

The best currently known explicit bound is $\kappa_K \leq \left(\frac{e\ln|\Delta_K|}{2(n-1)}\right)^{n-1}$ by [Lou00, Th. 1]. It implies the following upper bound on $h_K R_K$, as precisely shown in [BDPW20, Lem. 2.3], which can then be used to control the volume of the log-$\mathcal{S}$-unit lattice:

$$\ln\left(\sqrt{\tfrac{n}{2^{r_2}}}\cdot h_K R_K\right) \leq \tfrac{1}{2}\ln|\Delta_K| + n\ln\ln|\Delta_K| + n(1-\ln n). \tag{2.35}$$

Class Group Generators.

When picking a set of prime ideals in the algorithms of this thesis, an important feature is that they generate Cl_K. It is hence useful to bound both h_K and the norms of the generating prime ideals. Note that, as for any finite group, any non redundant generating set of Cl_K must have at most $\log h_K$ elements. Not much is generically known about the class number, so that the analytic estimation above is traditionally used to obtain $h_K \leq \widetilde{O}\left(\sqrt{|\Delta_K|}\right)$.

Let $\mathfrak{L}_{\mathsf{max}}$ be any prime ideal of maximum norm inside a generating set of Cl_K which has the smallest possible maximum norm. Bach proved that [Bac90, Th. 4]:

$$\mathcal{N}(\mathfrak{L}_{\mathsf{max}}) \leq 12\ln^2|\Delta_K|. \tag{2.36}$$

In practice though, this upper bound on the ratio $t_K := \mathcal{N}(\mathfrak{L}_{\mathsf{max}})/\ln^2|\Delta_K| \leq 12$ seems very pessimistic. Experimental evidence suggests that $t_K > 0.7$ only occurs in pathological cases [BDF08, §6], and as noted in [BDF08, p.1186], "*it even looks plausible that the* average *value of* $\mathcal{N}(\mathfrak{L}_{max})$ *as the discriminant of* K *increases is* $O(\ln|\Delta_K|)^{1+\varepsilon}$ *for any* $\varepsilon > 0$".

On the other hand, let us consider the relative part Cl_m^- of the class group of a cyclotomic field K_m. In this case, prime ideals belong to Cl_m^- only with probability roughly $1/h_m^+$, so we expect that searching for generators of the *subgroup* Cl_m^- mechanically increases the provable upper bound on generators. More precisely, writing as $\mathfrak{L}_{\mathsf{max}}^-$ the biggest ideal of a generating set of Cl_m^-, Wesolowski proved [Wes18, Rem. 2] that $\mathcal{N}(\mathfrak{L}_{\mathsf{max}}^-) \leq \left(2.71h_m^+\cdot\ln|\Delta_{K_m}| + 4.13\right)^2$.

Prime Ideal Theorem.

In order to constitute sufficiently large sets of prime ideals of polynomially bounded norms, it is useful to know the density of prime ideals in K. This is the object of the *Prime Ideal Theorem*, which states that prime ideals have more or less the same asymptotic behaviour as prime numbers.

Let $\pi_K(x) = \sharp\{\mathfrak{p} : \mathfrak{p} \text{ prime ideal}, \mathcal{N}(\mathfrak{p}) \leq x\}$, and $\vartheta_K(x) = \sum_{\mathcal{N}(\mathfrak{p})\leq x}\ln\mathcal{N}(\mathfrak{p})$. In [Lan03, §II.4–5], Landau proved the following asymptotic equivalences:

$$\pi_K(x) \sim_{x\to\infty} \int_2^x \frac{dt}{\ln t}, \quad \text{and} \quad \vartheta_K(x) \sim_{x\to\infty} x. \tag{2.37}$$

The general rough intuition is that each prime $p \in \mathbb{Z}$ yields *on average* one prime ideal in K of norm p. Of course, this global behaviour is not valid locally: for instance in cyclotomic fields $\mathbb{Q}(\zeta_m)$, ideals of prime norm p come in batches of $\varphi(m)$ elements for primes $p \equiv 1 \bmod m$, whose density is by Dirichlet's arithmetic progression theorem about $1/\varphi(m)$.

Unfortunately, whereas even for reasonably small bounds these asymptotic estimations yield astonishingly good results in practice, only effective versions are rigorously applicable.

Theorem 2.38 (Explicit Prime Ideal Theorem [GM16, Cor. 1.4])**.** *Under GRH,* $\forall x \geq 3$*:*

$$\left|\pi_K(x) - \pi_K(3) - \int_3^x \frac{dt}{\ln t}\right| \leq \sqrt{x}\cdot\Big[c_1(x)\cdot\ln|\Delta_K| + c_2(x)\cdot n\ln x + c_3(x)\Big],$$

with $c_1(x) = \left(\frac{1}{2\pi} - \frac{\ln\ln x}{\pi \ln x} + \frac{5.8}{\ln x}\right)$, $c_2(x) = \left(\frac{1}{8\pi} - \frac{\ln\ln x}{2\pi \ln x} + \frac{3.6}{\ln x}\right)$, $c_3(x) = \left(0.3 + \frac{14}{\log x}\right)$.

This can be used to show that a polynomial bound in $\ln|\Delta_K|$ yields sufficiently many prime ideals, like in [PHS19a, Cor. 2.9]. A precise version of that statement is given in [BDPW20, Lem. A.3]: for $x \geq \max\left\{(12\ln|\Delta_K|+8n+28)^4, 3\cdot 10^{11}\right\}$, $\pi_K(x) \geq \frac{x}{2\ln x}$. Note how this theoretical condition on x seems unnecessarily large in practice.

2.3.2 Hard problems in number theory

For our exposition, the most important problem to be considered is probably the *Class Group Discrete Logarithm Problem* (ClDLP). Solving this problem remains the major bottleneck in the classical query complexity of the Approx-id-SVP algorithms proposed in [CDW17, PHS19a, CDW21] and in this thesis.

Problem 2.39 (Class Group Discrete Logarithm Problem (ClDLP) [BS16])**.** *Given a set of prime ideals* $\{\mathfrak{L}_1, \ldots, \mathfrak{L}_k\}$*, and a challenge ideal* $\mathfrak{b}$*, find, if they exist,* $\alpha \in K$ *and integers* $v_1, \ldots, v_k$ *such that* $\langle\alpha\rangle = \mathfrak{b} \cdot \prod_i \mathfrak{L}_i^{v_i}$.

In this definition, we also require an *explicit* generator $\alpha \in K$, which slightly differs from the definition of e.g., [CDW17, Pr. 2]. Nevertheless, we note that in both quantum and classical worlds, the standard way to solve this problem boils down to computing $\mathcal{S}$-units, for $\mathcal{S}$ containing $\mathfrak{b}$ and the $\mathfrak{L}_i$'s, so that this explicit element is really a byproduct of the resolution. Furthermore, it is worth noting that the *Principal Ideal Problem* (PIP), i.e., that asks for a generator of $\mathfrak{b}$ if it exists, is encompassed in this definition of the ClDLP problem, using an empty set of ideals [BS16, Alg. 2].

Given a principal ideal described by some generator α, the *Shortest Generator Problem* (SGP) asks for the shortest generator α' such that $\langle\alpha\rangle = \langle\alpha'\rangle$. The SGP resolution can be reduced to a closest vector problem in the log-unit lattice, as is folklore in computational number theory. Similarly, we define:

Problem 2.40 (Shortest Class Group Discrete Logarithm (S-ClDLP))**.** *Given* $\langle\alpha\rangle = \mathfrak{b} \cdot \prod_i \mathfrak{L}_i^{v_i}$ *a solution to the Cl*DLP*, find* positive $w_1, \ldots, w_k \in \mathbb{Z}_{\geq 0}$ *and* $\alpha' \in K$ *such that* α' *is the smallest possible element such that* $\langle\alpha'\rangle = \mathfrak{b} \cdot \prod_i \mathfrak{L}_i^{w_i}$.

The condition for the w_i's to be positive is crucial. Note that all recent algorithms for Approx-id-SVP that are not bound to principal ideals eventually output an approximate solution of the S-ClDLP [CDW21, PHS19a, BR20]. If the set of prime ideals is sufficiently large compared to $\mathfrak{b}$, then S-ClDLP is exactly id-SVP.

Finally, we mention the *Close Principal Multiple Problem* (CPMP) which, given an ideal $\mathfrak{b}$, asks to find $\mathfrak{c}$ such that $\mathfrak{bc}$ is principal and $\mathcal{N}(\mathfrak{c})$ is "reasonably small" [CDW17, §2.2]. This specific problem also appears in [CDW21], where the authors prove that under GRH, using a factor base containing all prime ideals of norm up to $m^{4+o(1)}$, guarantees that a solution $\mathfrak{c}$ exists that satisfies $\mathcal{N}(\mathfrak{c}) \leq \exp\left(\widetilde{O}(m^{1+o(1)})\right)$ [CDW21, §1.3.4].

2.3.3 $\mathcal{S}$-unit groups computations

As shown in [BS16], the computation of class groups, unit groups, class group discrete logarithms and principal ideal generators can all be reduced to $\mathcal{S}$-units computations for appropriate sets of places $\mathcal{S}$. Thus, we are mostly interested in the running time of $\mathcal{S}$-unit groups related computations in K, which is denoted by $\mathrm{T}_{\mathsf{Su}}(K)$. Under the GRH:

- in the quantum world, $\mathrm{T}_{\mathsf{Su}}(K) = \widetilde{O}\big(\ln|\Delta_K|\big)$ is polynomial, as shown in [BS16], building upon generalizations of Shor's algorithm from [EHKS14];
- in the classical world, it remains subexponential in $\ln|\Delta_K|$, i.e., $\mathrm{T}_{\mathsf{Su}}(K) = \exp\widetilde{O}(\ln^{\alpha}|\Delta_K|)$ where $\alpha = 1/2$ for cyclotomic fields [BEF+17],[14] and $\alpha = 2/3$ in the general case [BF14], recently lowered to 3/5 by Gélin [Gél17].

Note that by abuse of notations, we omit here polynomial factors in $\sharp\mathcal{S}$ and $\max_{\mathfrak{p}\in\mathcal{S}\cap\mathcal{S}_0}\ln\mathcal{N}(\mathfrak{p})$.

2.4 Euclidean Lattices

Let L be a lattice. For any $p \in \mathbb{N}^*\cup\{\infty\}$ and $1 \le i \le \dim L$, the i-th minimum $\lambda_i^{(p)}(L)$ of L for the ℓ_p-norm is the minimum radius $r > 0$ such that $\{\mathbf{v} \in L : \|\mathbf{v}\|_p \le r\}$ has rank i [NV10, Def. 2.13]. For any $\mathbf{t}$ in the span of L, the distance between $\mathbf{t}$ and L is $\mathrm{dist}_p(\mathbf{t}, L) = \inf_{\mathbf{v}\in L}\|\mathbf{t}-\mathbf{v}\|_p$, and the *covering radius* of L w.r.t. the ℓ_p-norm is $\mu_p(L) = \sup_{\mathbf{t}\in L\otimes\mathbb{R}}\mathrm{dist}_p(\mathbf{t}, L)$. For the Euclidean norm, we occasionnally omit $p = 2$.

2.4.1 Estimating approximation factors

An *ideal lattice* of K is the full-rank image under the Minkowski embedding in $\mathbb{R}^n$ of a fractional ideal $\mathfrak{b}$ of K, where n is the degree of K. Its volume is given by $\mathrm{Vol}(\mathfrak{b}) = \mathcal{N}(\mathfrak{b})\cdot\sqrt{|\Delta_K|}$. Unlike generic lattices, a lower bound of the first minimum is implied by the arithmetic-geometric mean inequality, using that for any $\alpha \in \mathfrak{b}$, $\mathcal{N}(\mathfrak{b})$ divides $|\mathcal{N}(\alpha)|$. Thus, we obtain:

$$\sqrt{n}\cdot\mathcal{N}(\mathfrak{b})^{1/n} \le \lambda_1(\mathfrak{b}) \le \sqrt{n}\cdot\mathcal{N}(\mathfrak{b})^{1/n}\sqrt{|\Delta_K|}^{1/n}, \tag{2.41}$$

where the right inequality is Minkowski's inequality [NV10, Th. 2.4]. More precisely, the first minimum is bounded by $\lambda_1(\mathfrak{b}) \le (1+o(1))\sqrt{\frac{2n}{\pi e}}\cdot\mathrm{Vol}^{1/n}(\mathfrak{b})$, and the Gaussian Heuristic for full-rank random lattices [NV10, Def. 2.8] actually predicts $\lambda_1(\mathfrak{b}) \approx \sqrt{\frac{n}{2\pi e}}\cdot\mathrm{Vol}^{1/n}(\mathfrak{b})$ on average.

Applying the Gaussian Heuristic to ideal lattices yields a pretty good estimation of the shortness of vectors, even though $\lambda_1(\mathfrak{b})$ is not known precisely in general. This hypothesis is commonly used for the analysis of cryptosystems based on structured lattices, and the *exact* solutions found during the Twisted-PHS algorithm experiments in §3.4.3 match this heuristic.

For any $\mathbf{x} \in \mathfrak{b}$, let $\gamma(\mathbf{x}) = \|\mathbf{x}\|_2/\lambda_1(\mathfrak{b})$ denote the approximation factor reached by $\mathbf{x}$ in the ideal lattice $\mathfrak{b}$. As $\lambda_1(\mathfrak{b})$ is not known, the approximation factor $\mathrm{af}(\mathbf{x})$ is not directly accessible, but Eq. (2.41) implies the bounds $\gamma_{\mathsf{inf}}(\mathbf{x}) \le \gamma(\mathbf{x}) \approx \gamma_{\mathsf{gh}}(\mathbf{x}) \le \gamma_{\mathsf{sup}}(\mathbf{x})$, where:

$$\gamma_{\mathsf{inf}}(\mathbf{x}) := \frac{\|\mathbf{x}\|_2}{\sqrt{n}\cdot\mathrm{Vol}^{1/n}(\mathfrak{b})}, \qquad \gamma_{\mathsf{sup}}(\mathbf{x}) := \frac{\|\mathbf{x}\|_2}{\sqrt{n}\cdot\mathcal{N}(\mathfrak{b})^{1/n}}, \tag{2.42}$$
$$\gamma_{\mathsf{gh}}(\mathbf{x}) := \sqrt{2\pi e}\cdot\gamma_{\mathsf{inf}}(\mathbf{x}).$$

2.4.2 Computational problems

We will consider the following algorithmic lattice problems. Both problems can be readily restricted to ideal lattices under the labels Approx-id-Svp and Approx-id-Cvp.

Problem 2.43 (Approximate Shortest Vector Problem (Approx-Svp) [NV10, Pb. 2.2])**.** *Given a lattice L and an approximation factor $\gamma \ge 1$, find a vector $\mathbf{v} \in L$ such that $\|\mathbf{v}\|_2 \le \gamma\cdot\lambda_1(L)$.*

[14] The article [BEF+17] is written for prime-power cyclotomic fields for historical reasons, but readily adapts to the general case for class group computations.

Problem 2.44 (Approximate Closest Vector Problem (Approx-CVP) [NV10, Pb. 2.5])**.** *Given a lattice L, a target $\mathbf{t} \in L \otimes \mathbb{R}$ and an approximation factor $\gamma \geq 1$, find a vector $\mathbf{v} \in L$ such that $\|\mathbf{t} - \mathbf{v}\|_2 \leq \gamma \cdot \mathrm{dist}_2(\mathbf{t}, L)$.*

Actually, it will be more convenient to work with a slightly modified version of Approx-CVP, where the output is required to be at distance absolutely bounded by some B, independently of the target distance to the lattice. By abuse of terminology, we still call this variant Approx-CVP. A practical Approx-CVP oracle is given by Babai's Nearest Plane algorithm [Bab86], [Gal12, §18.1, Alg. 26].

2.4.3 Quality of a lattice basis

Evaluating the quality of a lattice basis is actually a tricky task that depends partly on the targeted problem (see e.g., [Xu13]), and several indicators have been used in the literature to attempt to measure this quality w.r.t. the SVP or the CVP.

Let $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ be a basis of a full-rank n-dimensional lattice L, and let the *Gram-Schmidt Orthogonalization* (GSO) of B be $\mathrm{GSO}(B) = (\mathbf{b}_1^\star, \ldots, \mathbf{b}_n^\star)$. Approximation algorithms usually attempt to compute a *good* basis of the given lattice, i.e., whose vectors are as short and as orthogonal as possible. These lattice reduction algorithms, such as LLL [LLL82] or BKZ [CN11], try to limit the decrease of the Gram-Schmidt norms $\|\mathbf{b}_i^\star\|_2$: intuitively, a wide gap in the sequence $\ln\|\mathbf{b}_i^\star\|_2$ at $i \geq 2$ reveals that $\mathbf{b}_i$ is rather not orthogonal to the previously generated subspace $\langle \mathbf{b}_1, \ldots, \mathbf{b}_{i-1} \rangle$. We will also consider the following standard quantities:

1. the root-Hermite factor δ_0 is widely used to measure the performance of lattice reduction algorithms [NS06, GN08, CN11], especially for solving SVP-like problems:

$$\delta_0^n(B) = \frac{\|\mathbf{b}_1\|_2}{\mathrm{Vol}^{1/n} B}. \tag{2.45}$$

   Experimental evidence suggest that on average, LLL achieves $\delta_0^{\mathsf{LLL}} \approx 1.022$ [NS06, GN08] and BKZ with block size b achieves $\delta_0^{\mathsf{BKZ}_b} \approx \left(\frac{b}{2\pi e}(\pi b)^{1/b}\right)^{1/(2b-2)}$ for $b \geq 50$ [Che13, CN11].
2. the (normalized) orthogonality defect δ [MG02, Def. 7.5] captures the global quality of the basis, not just of the first vector, and is especially useful for CVP-like problems e.g., if the lattice possesses abnormally short vectors:

$$\delta^n(B) = \prod_{i=1}^{n} \frac{\|\mathbf{b}_i\|_2}{\mathrm{Vol}^{1/n} B}. \tag{2.46}$$

   For purely orthogonal bases $\delta = 1$, and by Minkowski's second theorem [NV10, Th. 2.5], its smallest possible value is $\left(\prod_i \lambda_i(L)/\mathrm{Vol}\, L\right)^{1/n} \leq \sqrt{1 + \frac{n}{4}}$.
3. the minimum vector basis angle, defined as [Xu13, Eq. (15)]:

$$\theta_{\mathsf{min}}(B) = \min_{1 \leq i < j \leq n} \min\{\theta_{ij}, \pi - \theta_{ij}\} \text{ for } \theta_{ij} = \frac{\arccos\langle \mathbf{b}_i, \mathbf{b}_j \rangle}{\|\mathbf{b}_i\|_2 \|\mathbf{b}_j\|_2}. \tag{2.47}$$

   We propose to consider also the mean vector basis angle $\theta_{\mathsf{avg}}(B)$, which averages over all $\min\{\theta_{ij}, \pi - \theta_{ij}\}$.

Chapter 3

Twisted-PHS: Using the Product Formula to Solve Approx-SVP in Ideal Lattices

This chapter is based on an extended version of the first contribution of this thesis, which is a joint work with Adeline Roux-Langlois.

> [BR20] **Twisted-PHS: Using the Product Formula to Solve Approx-SVP in Ideal Lattices.**
> Olivier Bernard and Adeline Roux-Langlois.

Published in the proceedings of Asiacrypt 2020, Part II, vol. 12492 of *Lecture Notes in Computer Science* (LNCS) Series, pp.349–380, Springer.

Keywords: Ideal lattices, Approx-SVP, S-unit attacks, Twisted-PHS algorithm.
Links: [ePrint: 2020/1081[9] | GitHub: ob3rnard/Twisted-PHS[7]]

Contents

[9] https://eprint.iacr.org/2020/1081
[7] https://github.com/ob3rnard/Twisted-PHS

3.1 Introduction

In 2019, PELLET-MARY, HANROT and STEHLÉ [PHS19a] proposed an extended version of the cryptanalyses of [CDPR16, CDW17], which is proven for any number field K. The main feature of their algorithm is to use an exponential amount of preprocessing, depending only on K, in order to efficiently combine the two principal resolution steps of the CDW algorithm [CDW17, CDW21], namely the CPMP (*Close Principal Multiple Problem*) and the SGP (*Shortest Generator Problem*). Combining these two steps in a single CVP instance provides some guarantee that the output of the CPMP solver has a generator which is "*not much larger*" than its shortest non-zero vector.

In order to guarantee the output size and the running time of the PHS algorithm, a key ingredient is to use a CVP with preprocessing hint algorithm due to LAARHOVEN [Laa16], which represents the most costly part of the preprocessing phase.

3.1.1 Our contributions

Our main contribution is to propose a new "twisted" version of the PHS [PHS19a] algorithm, that we call Twisted-PHS. As a minor contribution, we also propose several improvements of the PHS algorithm, in a optimized version described in §3.2.3. On the theoretical side, we prove that our Twisted-PHS algorithm reaches the same asymptotic trade-off between runtime and approximation factor as the original PHS algorithm, using the same CVP solver with preprocessing hint by LAARHOVEN [Laa16].

On the practical side though, we provide a full implementation of our algorithm, which suggests that much better approximation factors are achieved and that the given lattice bases are much more orthogonal than the ones used in [PHS19a]. To our knowledge, this is the first time that this type of algorithm is completely implemented and tested for fields of degrees up to 60. As a point of comparison, experiments of [PHS19a] constructed the log-$\mathcal{S}$-unit lattice for cyclotomic fields of degrees at most 24, all but the last two being principal [PHS19a, Fig. 4.1]. We shall also mention the extensive simulations performed by [DPW19] using the Stickelberger lattice in prime power cyclotomic fields. Adapting these results to our construction is not immediate, as we need *explicit* $\mathcal{S}$-units to compute our lattice. This is left for future work.

We explain our experiments in §3.4, where we evaluate three algorithms instantiated with the same practical CVP oracle: the original PHS algorithm with the lattice implemented in [PHS19b]; our optimized version Opt-PHS (§3.2.3), and our new twisted variant Twisted-PHS (§3.3). We target two families of number fields, namely non-principal cyclotomic fields $\mathbb{Q}(\zeta_m)$ of prime conductors $m \in [\![23, 71]\!]$, and NTRU Prime fields $\mathbb{Q}(z_q)$ where z_q is a root of $x^q - x - 1$, for $q \in [\![23, 47]\!]$ prime. These correspond to the range of what is feasible in a reasonable amount of time in a classical setting. For cyclotomic fields, we managed to compute $\mathcal{S}$-units up to $\mathbb{Q}(\zeta_{71})$ for all factor bases, and all log-$\mathcal{S}$-unit lattice variants up to $\mathbb{Q}(\zeta_{61})$. For NTRU Prime fields, we managed all computations up to $\mathbb{Q}(z_{47})$.

3.1.2 Experiments

We chose to perform three experiments to test the performance of our Twisted-PHS algorithm, and to compare it with the two other algorithms:

- We first evaluate the *geometric characteristics* of the lattice output by the preprocessing phase: the root Hermite factor δ_0, the orthogonality defect δ, and the average vector basis angle θ_{avg}, as described in detail in §2.4.2. The last one seems difficult to interpret as it gives similar results in all cases, but the two other seem to show that the lattice output by Twisted-PHS is of better quality than in the two other cases. It shows significantly better root Hermite factor and orthogonality defect than any other lattice.
- For our second experiment, we evaluate *the Gram-Schmidt log norms* of each produced lattice. We propose two comparisons, the first one is before and after BKZ_{40} reduction to see the evolution of the norms in each case: it shows that the two curves are almost identical for Twisted-PHS but not for the other PHS variants. The second one is between the lattices output by the different algorithms, after BKZ_{40} reduction. The experiments emphasises that the decrease of the log norms seems much smaller in the twisted case than in the two other. Those two observations seem to corroborate the fact that the Twisted-PHS lattice is already quite orthogonal.
- Finally, we implemented all three algorithms from end to end and used them on numerous challenges to estimate their practically achieved *approximation factors*. This is to our knowledge the first time that these types of algorithms are completely run on concrete examples. The results of the experiments, shown in Fig. 3.1, suggest that the approximation factor reached by our algorithm increases very slowly with the dimension, in a way that could reveal subexponential or even better. We think that this last feature would be particularly interesting to prove.

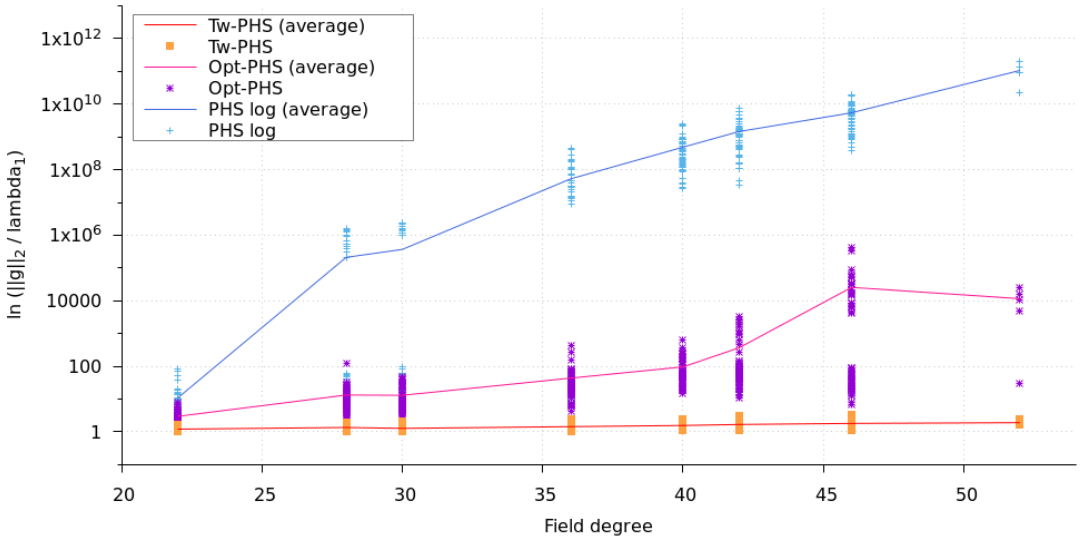

Figure 3.1 – Approximation factors reached by Twisted-PHS, Opt-PHS and PHS for cyclotomic fields of conductors 23, 29, 31, 37, 41, 43, 47 and 53 (in log scale).

3.1.3 Technical overview

We first quickly recall the principle of the PHS algorithm described in [PHS19a], which is split in two phases. The first phase consists in building a lattice that depends only on the number field K and allows to express any Approx-id-Svp instance in K as an Approx-Cvp instance in the lattice. This preprocessing chooses a factor base FB, and builds an associated lattice consisting in the

diagonal concatenation of some log-unit related lattice and the lattice of relations in the class group Cl_K between ideals of FB, with explicit generators. It then computes a hint of constrained size for the lattice to facilitate forthcoming Approx-CVP queries. Concretely, they suggest to use Laarhoven's algorithm [Laa16], which for any $\omega \in [0, 1/2]$ outputs a hint $\mathcal{V}$ of bit-size bounded by $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$ that allows to deliver answers for approximation factors $\widetilde{O}(\log|\Delta_K|^{\omega})$ in time bounded by the bit-size of $\mathcal{V}$ [Laa16, Cor. 1–2]. The second phase reduces the resolution of Approx-id-SVP to a single call to an Approx-CVP oracle in the lattice output by the preprocessing phase, for any challenge ideal $\mathfrak{b}$ in the maximal order of K. The main idea of this reduction is to multiply the principal ideal output by the ClDLP of $\mathfrak{b}$ on FB by ideals in FB until a "better" principal ideal is reached, i.e., having a short generator.

Our first contribution is to propose three improvements of the PHS algorithm. The first one consists in writing an explicit candidate for the isometry used in the computation of the lattice, and using its geometric properties to derive a smaller lattice dimension, while still guaranteeing the same proven approximation factor. The last two respectively modify the composition of the factor base and the definition of the target vector in a way that significantly improves the approximation factor experimentally achieved by the second phase of the algorithm. Although these improvements do not modify the core of PHS algorithm and have no impact on the asymptotics, they nevertheless are of importance in practice, as shown by our experiments in §3.4.

We now explain our main contribution, called Twisted-PHS, which is based on the PHS algorithm. As in PHS algorithm, our algorithm relies on the so-called *log-$\mathcal{S}$-unit lattice* with respect to a collection FB of prime ideals, called the factor base. This lattice captures local informations on FB, not only on (infinite) embeddings, to reduce a close principal multiple of a target ideal $\mathfrak{b}$ to a principal ideal containing $\mathfrak{b}$ which is guaranteed to have a somehow short generator. The main feature of our algorithm is to use the *Product Formula* to describe this log-$\mathcal{S}$-unit lattice. This induces two major changes in PHS algorithm:

1. The first one is twisting the $\mathfrak{p}$-adic valuations by $\ln \mathcal{N}(\mathfrak{p})$, giving weight to the fact that using a relation increasing the valuations at big norm ideals costs more than a relation involving smaller norm ideals.
2. The second one is projecting the target directly inside the log-$\mathcal{S}$-unit lattice and not only into the unit log-lattice corresponding to fundamental units.

Actually, the way our twisted version uses $\mathcal{S}$-units with respect to FB to reduce the solution of the ClDLP problem can be viewed as a natural generalization of the way classical algorithms reduce principal ideal generators using regular units.

Adding weights $\ln \mathcal{N}(\mathfrak{p})$ to integer valuations at any prime ideal $\mathfrak{p}$ intuitively allows to make a more relevant combination of the $\mathcal{S}$-units we use to reduce the output of the ClDLP, quantifying the fact that increasing valuations at big norm prime ideals costs more than increasing valuations at small norm prime ideals. Besides, the product formula induces the possibility to project elements on the whole log-$\mathcal{S}$-unit lattice instead of projecting only on the subspace corresponding to the log-unit lattice. As a consequence, it maintains inside the lattice the size and the algebraic norm logarithm of the $\mathcal{S}$-units. At the end, the CVP solver in this alternative lattice combines more efficiently the goal of minimizing the algebraic norm for the CPMP while still guaranteeing a small size for the SGP solution in the obtained principal multiple.

In §3.3, we describe two versions of our Twisted-PHS algorithm. The first one, composed by $\mathcal{A}^{(\mathrm{Laa})}_{\mathsf{tw\text{-}pcmp}}$ and $\mathcal{A}^{(\mathrm{Laa})}_{\mathsf{tw\text{-}query}}$ is proven to reach the same asymptotic trade-off between runtime and approximation factor as the original PHS algorithm, using the same CVP solver with preprocessing hint by Laarhoven. In practice, we propose two alternative algorithms $\mathcal{A}^{(\mathrm{bkz})}_{\mathsf{tw\text{-}pcmp}}$ and $\mathcal{A}^{(\mathrm{np})}_{\mathsf{tw\text{-}query}}$ with the following differences. Algorithm $\mathcal{A}^{(\mathrm{bkz})}_{\mathsf{tw\text{-}pcmp}}$ performs a minimal reduction step of the lattice as sole lattice preprocessing to smooth the input basis. Algorithm $\mathcal{A}^{(\mathrm{np})}_{\mathsf{tw\text{-}query}}$ resorts to

Babai's Nearest Plane algorithm for the Cvp solver role. Experimental evidence in §3.4 suggest that these algorithms perform remarkably well, because the twisted description of the log-$\mathcal{S}$-unit lattice seems much more orthogonal than expected. Proving this property would remove, in a quantum setting, the only part that is not polynomial in $\ln|\Delta_K|$.

3.2 The PHS Algorithm

This section describes the PHS algorithm, as introduced by Pellet-Mary, Hanrot and Stehlé in [PHS19a] for solving Approx-id-Svp, and discusses several improvements. The PHS algorithm extends the techniques from [CDPR16, CDW17] to any number field K and is split in two phases:

1. the preprocessing phase $\mathcal{A}_{\mathsf{pre\text{-}proc}}$, described in §3.2.1, builds a specific lattice together with some hint allowing to efficiently solve Approx-Cvp instances;
2. the query phase $\mathcal{A}_{\mathsf{query}}$, detailed in §3.2.2, reduces each Approx-id-Svp challenge to an Approx-Cvp instance in this fixed lattice.

More precisely, under the GRH and several heuristic assumptions detailed in [PHS19a, Heur. 1–6], they prove the following theorem:

Theorem 3.1 ([PHS19a, Th. 1.1]). *Let $\omega \in [0, 1/2]$ and K be a number field of degree n and discriminant Δ_K with a known basis of $\mathcal{O}_K$. Under some conjectures and heuristics, there exist two algorithms $\mathcal{A}_{\mathsf{pre\text{-}proc}}$ and $\mathcal{A}_{\mathsf{query}}$ such that:*

- *Algorithm $\mathcal{A}_{\mathsf{pre\text{-}proc}}$ takes as input $\mathcal{O}_K$, runs in time $2^{\widetilde{O}(\log|\Delta_K|)}$ and outputs a hint $\mathcal{V}$ of bit-size $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$;*
- *Algorithm $\mathcal{A}_{\mathsf{query}}$ takes as inputs any ideal $\mathfrak{b}$ of $\mathcal{O}_K$, whose algebraic norm has bit-size bounded by $2^{\mathrm{poly}(\log|\Delta_K|)}$, and the hint $\mathcal{V}$ output by $\mathcal{A}_{\mathsf{pre\text{-}proc}}$, runs in time $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)} + \mathrm{T}_{\mathsf{Su}}(K)$, and outputs a non-zero element $x \in \mathfrak{b}$ such that $\|x\|_2 \leq 2^{\widetilde{O}(\log^{\omega+1}|\Delta_K|/n)} \cdot \lambda_1(\mathfrak{b})$.*

We start by describing the preprocessing phase $\mathcal{A}_{\mathsf{pre\text{-}proc}}$ in §3.2.1, then the query phase $\mathcal{A}_{\mathsf{query}}$ in §3.2.2, and recall the proof of Th. 3.1 in detail. We thereafter discuss several algorithmic and theoretical minor improvements in §3.2.3.

3.2.1 Preprocessing of the number field

From a number field K and a size parameter $\omega \in [0, 1/2]$, the preprocessing phase consists in building and preparing a lattice L_{phs} that depends only on the number field K and allows to express any Approx-id-Svp instance in K as an Approx-Cvp instance in L_{phs}. The most significant part of this preprocessing is devoted to the computation of a hint of constrained size that can be used to facilitate those forthcoming Approx-Cvp queries.

We first define the lattice which is used in [PHS19a], discuss how the authors derive its dimension from volume considerations, and then expose the full preprocessing algorithm.

Definition of the lattice L_{phs}.

Let $\mathrm{FB} = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_k\}$ be a set of prime ideals generating the class group Cl_K. The lattice L_{phs} proposed in [PHS19a, §3.1] consists in the diagonal concatenation of some log-unit related lattice and the lattice of relations in Cl_K between ideals of FB, with explicit generators. Formally, it is

generated by the $(\nu + k)$ rows of the following square matrix:

$$B_{L_{\mathsf{phs}}} := \left[\begin{array}{c|c} c \cdot B_\Lambda & 0 \\ \hline c \cdot f_{H_0}(\mathbf{h}^{(0)}_{\eta_1}) & \\ \vdots & \ker \mathfrak{f}_{\mathrm{FB}} = \left(-v_{\mathfrak{p}_j}(\eta_i)\right)_{1 \le i,j \le k} \\ c \cdot f_{H_0}(\mathbf{h}^{(0)}_{\eta_k}) & \end{array}\right], \tag{3.2}$$

- where f_{H_0} is an isometry from $H_0 \subset \mathbb{R}^n$ to $\mathbb{R}^\nu$, where H_0 is the intersection of the span $\mathcal{L}_0$ of $\overline{\mathrm{Log}}\, \mathcal{O}_K$, i.e., $\mathcal{L}_0 = \{\mathbf{y} \in \mathbb{R}^n : y_{r_1+2i-1} = y_{r_1+2i}, i \in [\![1, r_2]\!]\}$, and of the trace zero hyperplane $\mathbb{R}^n_0 = \mathbf{1}_n^\perp$;
- the matrix B_Λ is a row basis of $f_{H_0}\left(\overline{\mathrm{Log}}\, \mathcal{O}_K^\times\right)$;
- the bottom right part of $B_{L_{\mathsf{phs}}}$ generates the lattice of all relations in Cl_K between ideals of FB, i.e., is the kernel of $\mathfrak{f}_{\mathrm{FB}} : (e_1, \ldots, e_k) \in \mathbb{Z}^k \mapsto \prod_j [\mathfrak{p}_j]^{e_j}$;
- each row vector $\mathbf{v}_i = (v_{i1}, \ldots, v_{ik})$ of $\ker \mathfrak{f}_{\mathrm{FB}}$ is associated to $\eta_i \in K$ s.t. $\langle \eta_i \rangle \cdot \prod_j \mathfrak{p}_j^{v_{ij}} = \mathcal{O}_K$, thus $v_{ij} = -v_{\mathfrak{p}_j}(\eta_i)$, and $\mathbf{h}^{(0)}_{\eta_i} = \pi_{H_0}\left(\overline{\mathrm{Log}}\, \eta_i\right)$, where π_{H_0} is the projection on H_0 in $\mathbb{R}^n$;
- c is a scaling parameter whose value depends on f_{H_0} (set later to $n^{3/2}/k$).

Note that this definition differs from the one given in [PHS19a, §3.1] by a sign change in the last k coordinates. This is a purely editorial detail allowing to use the same convention through the exposition of the algorithm and its proof.

The condition that the factor base generates Cl_K guarantees that for any challenge ideal there exists a solution to the ClDLP on FB. It can be relaxed to some extent to generate only a small index subgroup of Cl_K like in [CDW17].

The isometry f_{H_0} happens to play an important role in the proof of $\mathcal{A}_{\mathsf{query}}$. It forces the introduction of the scaling factor c, whose value is non-negligible and indirectly implies the use of a larger factor base. Note that this isometry is not explicitly defined in [PHS19a], whereas the associated code [PHS19b] uses a pruning strategy which removes the $r_2 + 1$ coordinates corresponding to the conjugates of complex places plus an arbitrary one. We stress that this implemented pruning strategy could negatively impact the quality of the Approx-CVP solver, as it hides potentially huge size variations of the $\mathcal{S}$-units on the removed coordinate. That's the reason why we thoroughly study in §3.2.3 a candidate isometry for f_{H_0} that also induces lower values for c. Furthermore, note that the projection on H_0 removes out of the picture the logarithm of the algebraic norms of the (non-regular) $\mathcal{S}$-units; hence, it seems that this partial information prevents L_{phs} from optimally achieving its initial goal of minimizing the algebraic norm for the CPMP while guaranteeing a SGP solution of small length. Our new algorithm, detailed in §3.3, aims in particular at fixing these flaws.

Finally, we aggregate the material present in [PHS19a, fn. 3, Lem. 3.1] to propose a simpler and more concise way to define L_{phs}; using the same notations as above, let φ_{phs} be the following map from K to $\mathbb{R}^\nu \times \mathbb{Z}^k$:

$$\varphi_{\mathsf{phs}}(\alpha) = \left(c \cdot f_{H_0} \circ \pi_{H_0}\left(\overline{\mathrm{Log}}\, \alpha\right), \left\{-v_{\mathfrak{p}_i}(\alpha)\right\}_{1 \le i \le k}\right). \tag{3.3}$$

Then, L_{phs} can be seen as the full-rank lattice generated by the images under φ_{phs} of the fundamental elements generating the $\mathcal{S}$-unit group $\mathcal{O}_{K,\mathcal{S}}^\times$, as given by Th. 2.2, with $\mathcal{S} = \mathcal{S}_\infty \cup \mathrm{FB}$ and for each $i \in [\![1, k]\!]$, $\varepsilon_{\nu+i} = \eta_i$. It is easy to see that both definitions coincide: for regular units $\varepsilon \in \mathcal{O}_K^\times$, all finite valuations are zero, so is the last part of $\varphi_{\mathsf{phs}}(\varepsilon)$, and $\pi_{H_0}\left(\overline{\mathrm{Log}}\, \varepsilon\right) = \overline{\mathrm{Log}}\, \varepsilon$.

Using the homomorphism properties of φ_{phs} on K, namely $\varphi_{\mathsf{phs}}(\alpha\alpha') = \varphi_{\mathsf{phs}}(\alpha) + \varphi_{\mathsf{phs}}(\alpha')$ and $\forall \lambda \in \mathbb{Z}$, $\varphi_{\mathsf{phs}}(\alpha^\lambda) = \lambda \cdot \varphi_{\mathsf{phs}}(\alpha)$, proving that each element of L_{phs} corresponds to an element of $\mathcal{O}_{K,\mathcal{S}}^\times$ [PHS19a, Lem. 3.1] becomes tautological. Further, we stress that φ_{phs} is injective on $\mathcal{O}_{K,\mathcal{S}}^\times/\mu(\mathcal{O}_K^\times)$ and therefore defines an isomorphism between $\mathcal{O}_{K,\mathcal{S}}^\times/\mu(\mathcal{O}_K^\times)$ and L_{phs}.

Volume of L_{phs} and cardinality of FB.

It remains to derive an explicit value for the cardinality k of the factor base FB; in [PHS19a, §4.1], this is done by considering the smallest k such that the root volume $\mathrm{Vol}^{1/(\nu+k)} L_{\mathsf{phs}}$ is at most constant. By Minkowski's inequality, this quantity bounds the first minimum in ℓ_∞-norm, and under the heuristic that L_{phs} behaves like a random lattice [PHS19a, Heur. 4], it also controls the ℓ_∞-norm covering radius $\mu_\infty(L_{\mathsf{phs}})$.

First, we evaluate the volume of L_{phs}, which writes as $c^\nu \cdot \det B_\Lambda \cdot \det(\ker \mathfrak{f}_{\mathrm{FB}})$ by definition of $B_{L_{\mathsf{phs}}}$. The determinant of $\ker \mathfrak{f}_{\mathrm{FB}}$ is $h_K = \sharp\big(\mathbb{Z}^k/\ker \mathfrak{f}_{\mathrm{FB}}\big)$. On the other hand, remark that B_Λ is the image under f_{H_0} of a basis of $\overline{\mathrm{Log}}\,\mathcal{O}_K^\times$, whose volume is $\sqrt{n}\cdot 2^{-r_2/2}\cdot R_K$ by Pr. 2.8. Finally, the isometry f_{H_0} stabilizes $\mathcal{L}_0 \cap \mathbb{R}_0^n$, thus preserves the volume of B_Λ; hence, we get:

$$\mathrm{Vol}\, L_{\mathsf{phs}} = c^\nu \cdot \frac{\sqrt{n}}{2^{r_2/2}} \cdot h_K R_K. \tag{3.4}$$

Note that [PHS19a] only gives an asymptotic bound on $\mathrm{Vol}\, L_{\mathsf{phs}}$, whereas Eq. (3.4) is exact. The idea is then to choose k such that $\mathrm{Vol}^{1/(\nu+k)} = O(1)$, e.g., taking $(\nu + k) = \ln \mathrm{Vol}\, L_{\mathsf{phs}}$. Using the number-theoretic bound given by Eq. (2.35), and using the fact that c will be later set to $n^{3/2}/k$, $\mathrm{Vol}\, L_{\mathsf{phs}}$ is asymptotically bounded by $\exp \widetilde{O}\big(\ln|\Delta_K| + n \ln\ln|\Delta_K|\big)$; therefore, $(\nu + k)$ can be set to:

$$\nu + k = \max\big\{\nu + \log h_K, \ln|\Delta_K| + n \ln\ln|\Delta_K|\big\}. \tag{3.5}$$

The $\log h_K$ part is there as a sufficient but not necessary condition ensuring that Cl_K can be generated by $k \geq \log h_K$ ideals [PHS19a, Lem. 2.7]. As $h_K \leq \widetilde{O}(\sqrt{|\Delta_K|})$, we remark that the second term dominates, so the maximum in the above formula can be ignored; in the associated code [PHS19b], $(k + \nu)$ is explicitly set to $\lfloor \ln|\Delta_K| \rfloor$. We stress that in practice the dimension of L_{phs} is quite sensitive to small changes in the value of c or the targeted root volume. We refer to §3.2.3 for more details and examples.

Preprocessing algorithm.

Algorithm 3.1 details the complete preprocessing procedure that, from a number field and some precomputation size parameter, chooses a factor base FB, builds the associated matrix $B_{L_{\mathsf{phs}}}$, and processes L_{phs} in order to facilitate Approx-Cvp queries.

The dimension k of the factor base and the scaling factor c are set in step 1 as in the published code [PHS19b]. Steps 2 and 3 are a concise version of [PHS19a, Alg. 3.1, steps 1–5]; it basically enlarges a generating set of Cl_K of size $k' \leq \log h_K$ by picking $(k - k')$ random prime ideals of bounded norms. The crucial point is to invoke the prime ideal theorem to show that taking a bound which is polynomial in k and $\log|\Delta_K|$ [PHS19a, Cor. 2.10] is actually sufficient.

The last step consists in preprocessing L_{phs} in order to solve Approx-Cvp instances efficiently. As noted in [PHS19a, p.6], the problem is easy without any constraint on the size of the output hint. To guarantee a hint size that is not exceeding the query phase time, they suggest to use Laarhoven's algorithm [Laa16], which outputs a hint $\mathcal{V}$ of bit-size bounded by $2^{\widetilde{O}((\nu+k)^{1-2\omega})}$, i.e., $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$ using $(\nu+k) = \widetilde{O}(\log|\Delta_K|)$, allowing to deliver the answer for approximation factors $(\nu + k)^\omega$ in time bounded by the bit-size of $\mathcal{V}$ [Laa16, Cor. 1–2].

Algorithm 3.1 PHS Preprocessing $\mathcal{A}_{\mathsf{pre\text{-}proc}}$

Input: A number field K of degree n and a parameter $\omega \in [0, 1/2]$.
Output: The basis $B_{L_{\mathsf{phs}}}$ with the preimages of its rows in $\mathcal{O}_{K,\mathcal{S}}^{\times}$, and Laarhoven's hint $\mathcal{V}(L_{\mathsf{phs}})$.

1: Set $k = \left(\lfloor \ln|\Delta_K| \rfloor - \nu\right)$ and $c = \left(n^{3/2}/k\right)$.
2: Compute $\mathrm{Cl}_K = \left\langle \left[\mathfrak{p}_1\right], \ldots, \left[\mathfrak{p}_{k'}\right] \right\rangle$, with $k' \leq \log h_K$.
3: Randomly extend $\left\{\mathfrak{p}_1, \ldots, \mathfrak{p}_{k'}\right\}$ by prime ideals of bounded norm to get $\mathrm{FB} = \left\{\mathfrak{p}_1, \ldots, \mathfrak{p}_k\right\}$.
4: Compute fundamental elements $\varepsilon_1, \ldots, \varepsilon_{\nu+k}$ of $\mathcal{O}_{K,\mathcal{S}}^{\times}$ as in Th. 2.2.
5: Create the matrix $B_{L_{\mathsf{phs}}}$ whose rows are $\varphi_{\mathsf{phs}}(\varepsilon_1), \ldots, \varphi_{\mathsf{phs}}(\varepsilon_{\nu+k})$ as defined in Eq. (3.2).
6: Use Laarhoven's algorithm to compute a hint $\mathcal{V} = \mathcal{V}(L_{\mathsf{phs}})$ of size $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$.
7: **return** $\left(\mathcal{O}_{K,\mathcal{S}}^{\times}, B_{L_{\mathsf{phs}}}, \mathcal{V}(L_{\mathsf{phs}})\right)$.

Proof of the first part of Th. 3.1. Costly steps of Alg. 3.1 are steps 2, 4 and 6 that compute the class group Cl_K, the $\mathcal{S}$-unit group $\mathcal{O}_{K,\mathcal{S}}^{\times}$ and the hint $\mathcal{V}(L_{\mathsf{phs}})$. The former two are $\mathcal{S}$-unit group related computations that cost $\mathrm{T}_{\mathsf{Su}}(K) \leq 2^{\widetilde{O}(\log^{2/3}|\Delta_K|)}$ each; the latter runs independently of ω in time $2^{O(\nu+k)} = 2^{\widetilde{O}(\log|\Delta_K|)}$. Note that in a quantum setting, only Laarhoven's algorithm is not polynomial in n; in a classical setting, it remains the dominating exponential part. □

3.2.2 Query phase: solving id-SVP using the preprocessing

This section describes the query phase $\mathcal{A}_{\mathsf{query}}$ of PHS algorithm; for any challenge ideal $\mathfrak{b} \subseteq K$ having a polynomial description in $\log|\Delta_K|$, it reduces the resolution of Approx-id-SVP in $\mathfrak{b}$ to a single call to an Approx-CVP oracle in L_{phs} as output by the preprocessing phase.

The main idea of this reduction is to multiply the principal ideal output by the ClDLP of $\mathfrak{b}$ on FB by ideals in FB until a "better" principal ideal is reached, i.e., having a short generator. In L_{phs}, it translates into adding vectors of L_{phs} to some target vector derived from $\mathfrak{b}$ until the result is short, hence into solving a CVP instance. This is formalized in Alg. 3.2, which rewrites [PHS19a, Alg. 3.2] to take into account our change of conventions in the definition of L_{phs} and the choice of Laarhoven's algorithm as the Approx-CVP oracle [Laa16, §4.2].

Algorithm 3.2 PHS Query $\mathcal{A}_{\mathsf{query}}$

Input: A challenge $\mathfrak{b}$, $\mathcal{A}_{\mathsf{pre\text{-}proc}}(K, \omega) = \left(\mathcal{O}_{K,\mathcal{S}}^{\times}, B_{L_{\mathsf{phs}}}, \mathcal{V}\right)$, and $\beta > 0$ s.t. for any $\mathbf{t}$, the Approx-CVP oracle using $\mathcal{V}(L_{\mathsf{phs}})$ outputs $\mathbf{w} \in L_{\mathsf{phs}}$ with $\|\mathbf{t} - \mathbf{w}\|_\infty \leq \beta$.
Output: A short element $x \in \mathfrak{b} \setminus \{0\}$.

1: Solve the ClDLP for $\mathfrak{b}$ on FB, i.e., find $\alpha \in K$ s.t. $\langle \alpha \rangle = \mathfrak{b} \cdot \prod_{\mathfrak{p}_i \in \mathrm{FB}} \mathfrak{p}_i^{v_i}$.
2: Define the target as $\mathbf{t} = \left(c \cdot f_{H_0} \circ \pi_{H_0}\left(\overline{\mathrm{Log}}\, \alpha\right), \left\{-v_i + \beta\right\}_{1 \leq i \leq k}\right)$.
3: Use the Approx-CVP solver with $\mathcal{V}(L_{\mathsf{phs}})$ to output $\mathbf{w} \in L_{\mathsf{phs}}$ s.t. $\|\mathbf{t} - \mathbf{w}\|_\infty \leq \beta$.
4: Compute $s = \varphi_{\mathsf{phs}}^{-1}(\mathbf{w}) \in \mathcal{O}_{K,\mathcal{S}}^{\times}$, using the preimages of $B_{L_{\mathsf{phs}}}$ rows.
5: **return** α/s.

Note that the output of the ClDLP in step 1 is an $\mathcal{S}$-unit if and only if $\mathfrak{b}$ is only divisible by prime ideals in the factor base. Each exponent v_i can be expressed as $v_i = v_{\mathfrak{p}_i}(\alpha) - v_{\mathfrak{p}_i}(\mathfrak{b})$. Then, the target defined in step 2 can be viewed as a drifted by β image of α in L_{phs}; using the formalism we introduced in Eq. (3.3), it writes simply as $\mathbf{t} = \varphi_{\mathsf{phs}}(\alpha) + \mathbf{b}_{\mathsf{phs}}$, where $\mathbf{b}_{\mathsf{phs}} = (0, \ldots, 0, \beta, \ldots, \beta)$ is non zero only on the k last coordinates. We stress that the role of $\mathbf{b}_{\mathsf{phs}}$ in the definition of the target serves a unique purpose: guarantee that $\alpha/s \in \mathfrak{b}$. In practice, this is not an anecdotic condition, and choosing β carefully has a significant impact on the length of the output, as we will see in §3.2.3.

The rest of this section is devoted to recall the proof of correctness, quality and running time of Alg. 3.2. These make an extensive use of the following log-unit structure lemma, which is classical and freely used e.g., in [CDPR16, §6.1]:

Lemma 3.6 ([PHS19a, Lem. 2.11–2.12])**.** *Define* $\mathbf{h}_\alpha^{(0)} := \pi_{H_0}\big(\overline{\mathrm{Log}}\,\alpha\big)$*, for* $\alpha \in K$*. Then we have* $\overline{\mathrm{Log}}\,\alpha = \mathbf{h}_\alpha^{(0)} + \frac{\ln|\mathcal{N}(\alpha)|}{n}\cdot\mathbf{1}_n$*. Further, the length of* α *is bounded by:*

$$\|\alpha\|_2 \leq \sqrt{n}\cdot|\mathcal{N}(\alpha)|^{1/n}\cdot\exp\|\mathbf{h}_\alpha^{(0)}\|_\infty \leq \sqrt{n}\cdot|\mathcal{N}(\alpha)|^{1/n}\cdot\exp\|\mathbf{h}_\alpha^{(0)}\|_2.$$

Proof. Recall that $\mathbb{R}_0^n = \mathbf{1}_n^\perp$ and $\overline{\mathrm{Log}}\,\alpha \in \mathcal{L}_0$, hence $\overline{\mathrm{Log}}\,\alpha$ decomposes as $\pi_{H_0}\big(\overline{\mathrm{Log}}\,\alpha\big) + a\cdot\mathbf{1}_n$, with $a = \big\langle\overline{\mathrm{Log}}\,\alpha, \mathbf{1}_n\big\rangle/\|\mathbf{1}_n\|_2^2 = \ln|\mathcal{N}(\alpha)|/n$, by definition of the projection on $\mathbb{R}_0^n$. Moreover, generically we have $\|\alpha\|_2 \leq \sqrt{n}\cdot\|\alpha\|_\infty$; using the above decomposition coordinate-wise, the j-th coordinate of $\overline{\mathrm{Log}}\,\alpha$ writes $\big(\overline{\mathrm{Log}}\,\alpha\big)_j = \big(\mathbf{h}_\alpha^{(0)}\big)_j + \frac{\ln|\mathcal{N}(\alpha)|}{n}$ and therefore:

$$\|\alpha\|_\infty = \max_{\sigma\in\mathcal{S}_\infty}|\sigma(\alpha)| = \exp\max_{\sigma\in\mathcal{S}_\infty}\ln|\sigma(\alpha)| \leq \exp\Big[\tfrac{\ln|\mathcal{N}(\alpha)|}{n} + \max_{1\leq j\leq n}\big(\mathbf{h}_\alpha^{(0)}\big)_j\Big].$$

Using $\max_j\big(\mathbf{h}_\alpha^{(0)}\big)_j \leq \|\mathbf{h}_\alpha^{(0)}\|_\infty$ and $\|\mathbf{h}_\alpha^{(0)}\|_\infty \leq \|\mathbf{h}_\alpha^{(0)}\|_2$ concludes. □

Notice how well the ℓ_∞-norm apparently behaves with respect to the logarithm embedding. We stress however that logarithms of small infinite valuations can become large negatives, so $\|\mathbf{h}_\alpha^{(0)}\|_\infty$ could be really far from $\max_{1\leq j\leq n}\big(\mathbf{h}_\alpha^{(0)}\big)_j$. This bounding method also somehow hides the fact that complex valuations count twice in the final Euclidean norm.

Theorem 3.7 ([PHS19a, Th. 3.3])**.** *Given access to an Approx-*Cvp *oracle that, on any input, outputs* $\mathbf{w} \in L_{\mathsf{phs}}$ *at infinity distance at most* β*, algorithm* $\mathcal{A}_{\mathsf{query}}$ *computes* $x \in \mathfrak{b}\setminus\{0\}$ *such that:*

$$\|x\|_2 \leq \sqrt{n}\cdot\mathcal{N}(\mathfrak{b})^{1/n}\cdot\exp\left[O\left(\frac{\beta\cdot k\cdot\ln\ln|\Delta_K|}{n}\right)\right].$$

Proof. Let $w_i = v_{\mathfrak{p}_i}(s)$, so that $\mathbf{w} = \varphi_{\mathsf{phs}}(s) = \big(c\cdot f_{H_0}(\mathbf{h}_s^{(0)}), \{-w_i\}_{1\leq i\leq k}\big)$. The first step is to prove correctness, i.e., that $x = (\alpha/s)$ is indeed in $\mathfrak{b}\setminus\{0\}$. By definition, we have $\langle s\rangle = \prod_{\mathfrak{p}_i\in\mathrm{FB}}\mathfrak{p}_i^{w_i}$, thus: $\langle\alpha/s\rangle = \mathfrak{b}\cdot\prod_{\mathfrak{p}_i\in\mathrm{FB}}\mathfrak{p}_i^{v_i-w_i}$. As $\|\mathbf{t}-\mathbf{w}\|_\infty \leq \beta$, for each i we have $|w_i - v_i + \beta| \leq \beta$, hence $0 \leq v_i - w_i \leq 2\beta$.

The second step is to bound the ℓ_2-norm of the output using Lem. 3.6. Hence, it is necessary to bound $|\mathcal{N}(\alpha/s)|$ and $\|\mathbf{h}_{\alpha/s}^{(0)}\|_\infty$. Bounding the former uses again that $0 \leq v_i - w_i \leq 2\beta$, as well as the fact that the maximal norm $\mathcal{N}(\mathfrak{L}_{\mathsf{max}})$ of FB is bounded by Bach's bound $O(\ln^2|\Delta_K|)$:

$$|\mathcal{N}(\alpha/s)|^{1/n} \leq \mathcal{N}(\mathfrak{b})^{1/n}\cdot\mathcal{N}(\mathfrak{L}_{\mathsf{max}})^{\sum_i(v_i-w_i)/n} \leq \mathcal{N}(\mathfrak{b})^{1/n}\cdot\exp\left[O\left(\frac{\beta\cdot k\cdot\ln\ln|\Delta_K|}{n}\right)\right].$$

As for the latter, $\|\mathbf{h}_{\alpha/s}^{(0)}\|_\infty \leq \|\mathbf{h}_{\alpha/s}^{(0)}\|_2 = \|f_{H_0}(\mathbf{h}_\alpha^{(0)} - \mathbf{h}_s^{(0)})\|_2 \leq \sqrt{n}/c\cdot\|\mathbf{t}-\mathbf{w}\|_\infty \leq \sqrt{n}\beta/c$. The value of c should then be set so that this bound is not greater than the previous $\frac{\beta\cdot k\cdot\ln\ln|\Delta_K|}{n}$. Taking $c = \frac{n^{3/2}}{k}$ as in [PHS19a] is sufficient. □

Before proving the second part of Th. 3.1, we remark that, taking the least possible values derived in §3.2.1 for $k = \ln|\Delta_K| \gtrsim n\ln n$ and $\mu_\infty(L_{\mathsf{phs}}) \approx 1$, and also assuming a perfect CVP solver in infinity norm for $\beta = \mu_\infty(L_{\mathsf{phs}})$, Th. 3.7 can at best only assess for a subexponential $n^{\ln n}$ approximation factor; polynomial approximation factors are not provably reached.

Proof of the second part of Th. 3.1. It breaks down to plugging $k = \widetilde{O}(\ln|\Delta_K|)$ and a value for β into Th. 3.7. In [PHS19a, §4.2], deriving this β relies on several heuristics [PHS19a, Heur. 4–6] implying that $\mu_2(L_{\mathsf{phs}}) = O(\sqrt{\nu+k})$, and that on average $\|\mathbf{v}\|_\infty \leq \frac{\ln \nu+k}{\sqrt{\nu+k}} \cdot \|\mathbf{v}\|_2$. The Approx-CVP solver from Laarhoven's algorithm using $\mathcal{V}(L_{\mathsf{phs}})$ outputs a lattice vector at Euclidean distance which is at most $O\big((\nu+k)^\omega \cdot \mu_2(L_{\mathsf{phs}})\big)$. Using the above heuristics, the infinity distance of the output is therefore $\widetilde{O}\big((\nu+k)^\omega\big) = \widetilde{O}(\ln^\omega|\Delta_K|)$, giving the claimed bound.

As for the running time of Alg. 3.2, it is essentially determined by those of steps 1 and 3. Solving the ClDLP problem requires to compute $\mathcal{S}$-units for an extended factor basis containing FB and prime factors of $\mathfrak{b}$, hence costs $\mathrm{T}_{\mathsf{Su}}(K)$. Note that in a quantum setting, $\mathrm{T}_{\mathsf{Su}}(K)$ is polynomial, but in a classical world it remains subexponential in the discriminant; furthermore, since it depends on the challenge, this cost cannot be mitigated by some preprocessing effort. On the other hand, solving Approx-CVP with Laarhoven's algorithm runs in time bounded by $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$, the size of V. Finally, the total run time of $\mathcal{A}_{\mathsf{query}}$ is bounded by $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)} + \mathrm{T}_{\mathsf{Su}}(K)$. □

3.2.3 Optimizing PHS parameters

In this section, we propose three improvements of the PHS algorithm. The first one consists in writing an explicit candidate for f_{H_0} and using its geometric properties to derive a smaller lattice dimension, while still guaranteeing the same proven approximation factor. The last two respectively modify the composition of the factor base and the definition of the target vector in a way that drastically improves the approximation factor experimentally achieved by $\mathcal{A}_{\mathsf{query}}$.

Although these improvements do not modify the core of PHS algorithm and have no impact on the asymptotics, they nevertheless are of importance in practice, as we will see in §3.4.

Expliciting the isometry: towards smaller factor bases.

We exhibit an explicit candidate for the isometry f_{H_0} going from $H_0 = \mathbb{R}_0^n \cap \mathcal{L}_0 \subseteq \mathbb{R}^n$ to $\mathbb{R}^\nu$ and evaluate its effect on the infinity norm; it allows to lower the value of c in the proof of Th. 3.7 from $n\sqrt{n}/k$ to $n(1+\ln n)/k$, which in turn implies using a smaller factor base for the same proven approximation factor. We define the isometry f_{H_0} as the linear map represented by $\overline{\mathrm{GSO}}^{\mathrm{T}}(M_{H_0})$, with:

$$M_{H_0} := \begin{pmatrix} n & \overbrace{\begin{matrix} -1 & 1 & & \\ & -1 & 1 & \\ & & \ddots & \ddots \\ & & & -1 \quad 1 \end{matrix}}^{\nu+1} \end{pmatrix} \cdot \begin{pmatrix} \overbrace{\begin{matrix} I_{r_1} & \begin{matrix} \frac{1}{2} & \frac{1}{2} \\ & \ddots \end{matrix} \end{matrix}}^{r_1+2r_2} \\ \begin{matrix} \frac{1}{2} & \frac{1}{2} \\ & \ddots \end{matrix} \end{pmatrix} \begin{matrix} \} r_1 \\ \} r_2 \end{matrix} \quad (3.8)$$

Actually, M_{H_0} is simply a basis of $\mathbb{R}_0^n \cap \mathcal{L}_0$ in $\mathbb{R}^n$, constituted of vectors that are orthogonal to $\mathbf{1}_n$ and to each of the r_2 independent vectors $\mathbf{v}_j$, $j \in [\![1, r_2]\!]$, that sends any $\mathbf{y} \in \mathcal{L}_0$ to $\mathbf{0}$ by substracting y_{r_1+2j} from its copy y_{r_1+2j-1} and forgetting every other coordinate.

Proposition 3.9. *Let f_{H_0} be the isometry represented by $\overline{\mathrm{GSO}}^{\mathrm{T}}(M_{H_0})$. Then:*

$$\forall \mathbf{h} \in H_0, \qquad \|\mathbf{h}\|_\infty \leq (1+\ln n)\cdot \|f_{H_0}(\mathbf{h})\|_\infty,$$
$$\|f_{H_0}(\mathbf{h})\|_\infty \leq 2\sqrt{2}\cdot \|\mathbf{h}\|_\infty.$$

Proof. Let $\mathbf{h} \in \mathbb{R}_0^n \cap \mathcal{L}_0$, and $\mathbf{v} = f_{H_0}(\mathbf{h}) \in \mathbb{R}^\nu$. We prove $\|f_{H_0}^{-1}(\mathbf{v})\|_\infty \leq (1+\ln n)\cdot\|\mathbf{v}\|_\infty$, which is trivially equivalent. By definition, $f_{H_0}^{-1}(\mathbf{v}) = \mathbf{v}\cdot\overline{\mathrm{GSO}}(M_{H_0})$, hence bounding the ℓ_1-norm of each column of $\overline{\mathrm{GSO}}(M_{H_0})$ by $(1+\ln n)$ yields the first inequality. Similarly, bounding the ℓ_1-norm of each row of $\overline{\mathrm{GSO}}(M_{H_0})$ by $2\sqrt{2}$ proves the second.

Let $\mathbf{b}_1, \ldots, \mathbf{b}_\nu$ be the row vectors of M_{H_0}; the Gram-Schmidt orthogonalization (resp. orthonormalization) vectors of M_{H_0} are denoted by $\mathbf{b}_i^\star$ (resp. $\overline{\mathbf{b}}_i^\star$). Because of the particular structure of M_{H_0}, $\overline{\mathbf{b}}_{i+1}^\star$ only depends on $\mathbf{b}_{i+1}$ and $\mathbf{b}_i^\star$. Then, a simple induction shows that:

$$\begin{cases} \forall i \in [\![1, r_1 - 1]\!]: & \overline{\mathbf{b}}_i^\star = \left(-\frac{1}{\sqrt{i(i+1)}}, \ldots, \sqrt{\frac{i}{i+1}}, 0, \ldots\right), \\ \forall j \in [\![0, r_2 - 1]\!],\ i = r_1 + 2j: & \overline{\mathbf{b}}_{r_1+j}^\star = \left(-\frac{\sqrt{2}}{\sqrt{i(i+2)}}, \ldots, \frac{\sqrt{i}}{\sqrt{2(i+2)}}, \frac{\sqrt{i}}{\sqrt{2(i+2)}}, 0, \ldots\right), \end{cases}$$

where in each configuration the first i coordinates are equal, and zeroes pad to dimension n. Bounding each $\|\overline{\mathbf{b}}_i^\star\|_1$ by $2\sqrt{2}$ is trivial from these formulas, proving the second inequality. Let $\mathbf{c}_1, \ldots, \mathbf{c}_n$ be the columns of $\overline{\mathrm{GSO}}(M_{H_0})$. We claim that $\|\mathbf{c}_n\|_1 \leq \|\mathbf{c}_{n-1}\|_1 \leq \cdots \leq \|\mathbf{c}_1\|_1$. Indeed, $\|\mathbf{c}_1\|_1 = \|\mathbf{c}_2\|_1$, and for all $i \geq 2$, $\|\mathbf{c}_i\|_1 - \|\mathbf{c}_{i+1}\|_1 = |(\overline{\mathbf{b}}_{i-1}^\star)_i| + |(\overline{\mathbf{b}}_i^\star)_i| - |(\overline{\mathbf{b}}_i^\star)_{i+1}| \geq 0$. Using $\sqrt{\frac{1}{i(i+1)}} < \frac{1}{i}$ and $\sqrt{\frac{2}{i(i+2)}} \leq \frac{1}{\sqrt{2}}\left(\frac{1}{i} + \frac{1}{i+1}\right)$ yields $\|\mathbf{c}_1\|_1 \leq \sum_{i=1}^{n-1} \frac{1}{i} \leq 1 + \ln(n-1)$. □

As a consequence, we can directly inject this result into the proof of Th. 3.7 to bound the ℓ_∞-norm $\|\mathbf{h}_{\alpha/s}^{(0)}\|_\infty$ by $(1+\ln n)/c \cdot \|\mathbf{t} - \mathbf{w}\|_\infty \leq (1+\ln n)\beta/c$ instead of $\sqrt{n}\beta/c$. We also use the following refined practical bound on the algebraic norm of α/s. Indeed, when conducting experiments, FB is known and there is no need to suffer from Bach's generic bound for $\mathcal{N}(\mathfrak{L}_{\mathsf{max}})$:

$$|\mathcal{N}(\alpha/s)|^{1/n} \leq \mathcal{N}(\mathfrak{b})^{1/n} \cdot \prod_{\mathfrak{p}_i \in \mathrm{FB}} \mathcal{N}(\mathfrak{p}_i)^{(v_i - w_i)/n} \leq \mathcal{N}(\mathfrak{b})^{1/n} \cdot \exp\left[\frac{2\beta \cdot \sum_{\mathfrak{p}\in\mathrm{FB}} \ln \mathcal{N}(\mathfrak{p})}{n}\right]. \quad (3.10)$$

Then, as a smaller value of c implies a smaller volume of L_{phs} hence a smaller factor base, it should be chosen as the smallest s.t. the former bound $(1+\ln n)\beta/c$ on $\|\mathbf{h}_{\alpha/s}^{(0)}\|_\infty$ is below the above $\frac{2\beta\cdot\sum_{\mathfrak{p}\in\mathrm{FB}}\ln\mathcal{N}(\mathfrak{p})}{n}$, which implies $c \geq \frac{(1+\ln n)n}{\sum_{\mathfrak{p}\in\mathrm{FB}}\ln\mathcal{N}(\mathfrak{p})}$. Nevertheless, as there is no reason to artificially *increase* the bound on $\|\mathbf{h}_{\alpha/s}^{(0)}\|_\infty$ using $c < 1$ when the other already dominates, we should also ensure $c \geq 1$. This finally leads us to choose:

$$c = \max\left(1, \frac{(1+\ln n)n}{\sum_{\mathfrak{p}\in\mathrm{FB}}\ln\mathcal{N}(\mathfrak{p})}\right). \quad (3.11)$$

To quantify the gain obtained by this new value of c, we computed factor base dimensions in different settings for two families of number fields: Tab. 3.1 deals with non-principal cyclotomic fields $\mathbb{Q}(\zeta_m)$ of prime conductors $m \in [\![23, 71]\!]$; Tab. 3.2 handles NTRU Prime fields $\mathbb{Q}(z_q)$, where z_q is a root of $x^q - x - 1$, for q prime in $[\![23, 61]\!]$. These correspond to the range of explicit computations feasible within a limited amount of time. By contrast, experiments reported in [PHS19a, Fig. 4.1] were limited to cyclotomic fields of degree at most 24, most of them being principal. For each field, we compare the expected factor base dimensions in four situations:

m	$\ln V^{1/(\nu+k)}$ [PHS19b]	Eq. (3.5) [PHS19a]	$c = n^{3/2}/k$ [PHS19b]	$c = \frac{(1+\ln n)n}{k}$	$c = \max\left(1, \frac{(1+\ln n)n}{\sum \ln \mathcal{N}(\mathfrak{p})}\right)$
23	0.292	147	55	53	34
29	0.305	204	77	72	50
31	0.304	223	85	79	55
37	0.314	283	109	100	72
41	0.323	324	125	114	84
43	0.323	345	134	121	91
47	0.327	388	151	136	103
53	0.336	453	177	158	122
59	0.341	520	204	181	141
61	0.343	543	213	189	148
67	0.348	611	241	212	168
71	0.350	658	260	228	182

TABLE 3.1 – Values of k for $K = \mathbb{Q}(\zeta_m)$: using Eq. (3.5); using Eq. (3.4) with same root volume target $V^{1/(\nu+k)}$ as in [PHS19b] and given values of c.

q	$\ln V^{1/(\nu+k)}$ [PHS19b]	Eq. (3.5) [PHS19a]	$c = n^{3/2}/k$ [PHS19b]	$c = \frac{(1+\ln n)n}{k}$	$c = \max\left(1, \frac{(1+\ln n)n}{\sum \ln \mathcal{N}(\mathfrak{p})}\right)$
23	0.264	159	61	58	37
29	0.285	216	83	77	52
31	0.289	236	91	84	58
37	0.299	296	115	105	75
41	0.306	338	132	119	88
43	0.313	359	140	126	93
47	0.320	402	157	141	106
53	0.325	467	184	164	125
59	0.335	535	211	187	145
61	0.334	557	220	194	151

TABLE 3.2 – Values of k for $K = \mathbb{Q}(z_q)$: using Eq. (3.5); using Eq. (3.4) with same root volume target $V^{1/(\nu+k)}$ as in [PHS19b] and given values of c.

first, for completeness we use Eq. (3.5), taken from [PHS19a, §4.1]; then we report the value used by [PHS19b], i.e., $k = \lfloor \ln|\Delta_K| \rfloor - \nu$, and provide the resulting root volume $\mathrm{Vol}^{1/(\nu+k)} L_{\mathsf{phs}}$ corresponding to $c = \frac{n^{3/2}}{k}$ for reference. Finally, we target this reference root volume using on one hand $c = \frac{(1+\ln n)n}{k}$, hence mimicking the proof of Th. 3.7, and on the other hand using our recommended value given by Eq. (3.11).

The last experiment, dealing with Eq. (3.11), simulates all factor bases of cardinality k by taking the k prime ideals of smallest norms. This choice might not be directly suitable for a factor base, as it gives no theoretical insurance to generate Cl_K. Nevertheless, in all experiments the obtained k is well above $\log h_K$, the maximum number of generators of Cl_K [PHS19a, Lem. 2.7], so that replacing some of these ideals by bigger norm representatives of missing classes until the set generates Cl_K would only *reduce* the value of c by increasing $\sum_{\mathfrak{p}\in \mathrm{FB}} \ln \mathcal{N}(\mathfrak{p})$. Thus, the given factor base dimensions remain in any case an upper bound of the correct dimension.

To end this section, we remark that there might exist better ℓ_∞-norm preserving isometries than $\overline{\mathrm{GSO}}(M_{H_0})^{\mathrm{T}}$; nevertheless, as the value of c derived from Eq. (3.11) is already equal to 1

most of the time, we cannot expect a substantial gain from this. Furthermore, it should be stressed that the complexity of known lattice reduction algorithms only depends on the rank of the lattice, and *not* on the ambient space dimension, so that this isometry can be removed in practice. It however serves the theoretical purpose of being able to transpose Minkowski's inequalities and heuristics on covering radii that are valid only for full-rank lattices.

Lowering the factor base weight.

Second, we suggest choosing the k elements of the factor base as the k prime ideals of least possible norm, instead of randomly picking them up to some polynomial bound. As shown by Eq. (3.10), this incidentally lowers the approximation factor, which depends on $\prod_{\mathfrak{p}\in\mathrm{FB}}\mathcal{N}(\mathfrak{p})$.

Formally, this only modifies step 3 of Alg. 3.1 as follows. Let $\{\mathfrak{p}_1,\dots,\mathfrak{p}_{k'}\}$ be a generating set of Cl_K, with $k'\leq \log h_K$, as obtained by the previous step 2. As in Alg. 3.1, using the prime ideal theorem yields that we can choose some bound B polynomial in k and $\log|\Delta_K|$ such that the set of prime ideals of norm bounded by B contains at least k elements. Then, we order this set by increasing norms, choosing an arbitrary permutation for isonorm ideals, and remove ideals that were already present in $\{\mathfrak{p}_1,\dots,\mathfrak{p}_{k'}\}$. It remains to extract the first $(k-k')$ elements to obtain our factor base.

There is one issue to consider, namely adapting the justification of [PHS19a, Heur. 4], relying on L_{phs} being a "somehow random" lattice to derive that $\mu_\infty(L_{\mathsf{phs}})$ is close to $\lambda_1^{(\infty)}(L_{\mathsf{phs}})$. We argue that in practice (as discussed with more details for Heur. 3.28 in §3.3.2), it is always possible to empirically upper bound the infinity covering radius of L_{phs} to verify that this heuristic holds. For example, as described in [PHS19a, §4.1]: take sufficiently many random samples $\mathbf{t}_i$ in the span of L_{phs} from a continuous Gaussian distribution of sufficiently large deviation; solve Approx-CVP for the ℓ_2-norm for each of them to obtain vectors $\mathbf{w}_i\in L_{\mathsf{phs}}$ close to $\mathbf{t}_i$; finally, majorate $\mu_\infty(L_{\mathsf{phs}})$ by $\max_i\|\mathbf{t}_i-\mathbf{w}_i\|_\infty$. Then, if the expected heuristic behaviour is too far from this estimate, we could still replace one ideal of FB by an ideal of bigger norm and iterate the process.

Minimizing the target drift.

Our last suggested improvement modifies the definition of the target vector to take into account the fact that valuations at prime ideals are integers. Hence, the condition enforcing $\alpha/s\in\mathfrak{b}$, which was written as $\forall\mathfrak{p}\in\mathrm{FB}$, $v_\mathfrak{p}(\alpha)-v_\mathfrak{p}(s)\geq 0$, can be replaced by the equivalent requirement that $\forall\mathfrak{p}\in\mathrm{FB}$, $v_\mathfrak{p}(\alpha)-v_\mathfrak{p}(s)>-1$. Intuitively, this reduces the valuations at prime ideals of the output element by one on average, hence lowering the approximation factor bound in Eq. (3.10). Formally, using the notations of Alg. 3.2, we only modify the definition of the target $\mathbf{t}$ in step 2 of Alg. 3.2. For any $0<\varepsilon<1$, let $\widetilde{\beta}=(\beta-1+\varepsilon)$ and let $\widetilde{\mathbf{b}}_{\mathsf{phs}}=(0,\dots,0,\widetilde{\beta},\dots,\widetilde{\beta})$ with non zero values only on the k last coordinates. The modified target is defined as:

$$\widetilde{\mathbf{t}}=\varphi_{\mathsf{phs}}(\alpha)+\widetilde{\mathbf{b}}_{\mathsf{phs}}=\Big(c\cdot f_{H_0}\circ\pi_{H_0}\big(\overline{\mathrm{Log}}\,\alpha\big),\big\{-v_i+\widetilde{\beta}\big\}_{1\leq i\leq k}\Big). \tag{3.12}$$

The remaining steps of Alg. 3.2 stay unchanged. We have to prove that the output is still correct, i.e., that $\alpha/s\in\mathfrak{b}$, where $\mathbf{w}=\varphi_{\mathsf{phs}}(s)\in L_{\mathsf{phs}}$ verifies $\|\widetilde{\mathbf{t}}-\mathbf{w}\|_\infty\leq\beta$. This is done in the following Pr. 3.13, which adapts Th. 3.7 to benefit from all the improvements of this section.

Though this adjustment might seem insignificant at first sight, we stress that the induced gain is of order $\prod_{\mathfrak{p}\in\mathrm{FB}}\mathcal{N}(\mathfrak{p})^{1/n}$, which is roughly subexponential in n, and that its impact is very noticeable experimentally. In fact, the quality of the output is so sensitive to this $\widetilde{\beta}$ that we implemented a dichotomic strategy to find, for each challenge $\mathfrak{b}$, the smallest possible translation $\widetilde{\beta}$ that must be applied to $\varphi_{\mathsf{phs}}(\alpha)$ to ensure $(\alpha/s)\in\mathfrak{b}$.

Proposition 3.13. *Given access to an Approx-*CVP *oracle that, on any input, output* $\mathbf{w} \in L_{phs}$ *at infinity distance at most* β*, the modified algorithm* $\mathcal{A}_{query}$ *using the isometry* f_{H_0} *defined in Eq. (3.8), the value* c *defined in Eq. (3.11), and for any* $0 < \varepsilon < 1$*, the modified target* $\widetilde{\mathbf{t}}$ *defined in Eq. (3.12), computes* $x \in \mathfrak{b} \setminus \{0\}$ *such that:*

$$\|x\|_2 \leq \sqrt{n} \cdot \mathcal{N}(\mathfrak{b})^{1/n} \cdot \exp\left[\frac{(\beta + \lfloor 2\beta - 1 \rfloor) \cdot \sum_{\mathfrak{p} \in \mathrm{FB}} \ln \mathcal{N}(\mathfrak{p})}{n}\right].$$

Proof. As in the proof of Th. 3.7, let $\mathbf{w} = \varphi_{\mathsf{phs}}(s) = \left(c \cdot f_{H_0}(\mathbf{h}_s^{(0)}), \{-w_i\}_{1 \leq i \leq k}\right)$, with $w_i = v_{\mathfrak{p}_i}(s)$, be such that $\|\widetilde{\mathbf{t}} - \mathbf{w}\|_\infty \leq \beta$. The main point is proving that $x = (\alpha/s) \in \mathfrak{b}$. Recall that $\langle \alpha/s \rangle = \mathfrak{b} \cdot \prod_{\mathfrak{p}_i \in \mathrm{FB}} \mathfrak{p}_i^{v_i - w_i}$. As $\|\widetilde{\mathbf{t}} - \mathbf{w}\|_\infty \leq \beta$, for each i we have $-1 + \varepsilon \leq v_i - w_i \leq 2\beta - 1 + \varepsilon$. Using that v_i, w_i are in $\mathbb{Z}$ and $\varepsilon > 0$ implies $0 \leq v_i - w_i \leq \lfloor 2\beta - 1 \rfloor$, hence $x \in \mathfrak{b} \setminus \{0\}$.

The ℓ_2-norm of x is upper bounded using again Lem. 3.6. The previous discussion also shows $|\mathcal{N}(\alpha/s)|^{1/n} \leq \mathcal{N}(\mathfrak{b})^{1/n} \cdot \exp\left(\frac{\lfloor 2\beta - 1 \rfloor \cdot \sum_{\mathfrak{p} \in \mathrm{FB}} \ln \mathcal{N}(\mathfrak{p})}{n}\right)$. Using the isometry properties given by Pr. 3.9, we obtain $\|\mathbf{h}_{\alpha/s}^{(0)}\|_\infty \leq (1 + \ln n)\beta/c$, and using $c \geq \frac{(1+\ln n)n}{\sum_{\mathfrak{p} \in \mathrm{FB}} \ln \mathcal{N}(\mathfrak{p})}$ as implied by Eq. (3.11) finally yields the result. □

3.3 The Twisted-PHS Algorithm

Our main contribution is to propose a twisted version of the PHS algorithm. The main idea consists in using the natural description of the log-$\mathcal{S}$-unit lattice given in Eq. (2.3) and deduced from the product formula in Eq. (2.1). This basically adds weights to each $\mathfrak{p}$-adic valuation, which has several valuable consequences.

On the theoretical side, we prove that our Twisted-PHS algorithm reaches the same asymptotic trade-off between runtime and approximation factor as the original PHS algorithm, using the same CVP solver with preprocessing hint by Laarhoven. Formally, under the GRH and heuristics:

Theorem 3.14. *Let* $\omega \in [0, 1/2]$ *and* K *be a number field of degree* n *and discriminant* Δ_K*. Assume that a basis of* $\mathcal{O}_K$ *is known. Under GRH (Heur. 2.33) and Heur. 3.28 and 3.29, there exist two algorithms* $\mathcal{A}^{(Laa)}_{\mathsf{tw\text{-}pcmp}}$ *and* $\mathcal{A}^{(Laa)}_{\mathsf{tw\text{-}query}}$ *such that:*

- *Algorithm* $\mathcal{A}^{(Laa)}_{\mathsf{tw\text{-}pcmp}}$ *takes as input* $\mathcal{O}_K$*, runs in time* $2^{\widetilde{O}(\log|\Delta_K|)}$ *and outputs a hint* $\mathcal{V}$ *of bit-size* $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$*;*
- *Algorithm* $\mathcal{A}^{(Laa)}_{\mathsf{tw\text{-}query}}$ *takes as inputs any ideal* $\mathfrak{b}$ *of* $\mathcal{O}_K$*, whose algebraic norm has bit-size bounded by* $2^{\mathrm{poly}(\log|\Delta_K|)}$*, and the hint* $\mathcal{V}$ *output by* $\mathcal{A}^{(Laa)}_{\mathsf{tw\text{-}pcmp}}$*, runs in time* $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)} + \mathrm{T}_{\mathsf{Su}}(K)$*, and outputs a non-zero element* $x \in \mathfrak{b}$ *such that* $\|x\|_2 \leq 2^{\widetilde{O}(\log^{\omega+1}|\Delta_K|/n)} \cdot \lambda_1(\mathfrak{b})$*.*

On the practical side though, experimental evidence given in §3.4 suggest that we achieve much better approximation factors than expected, and that the given lattice bases are a lot more orthogonal than the ones used in [PHS19a].

3.3.1 Preprocessing of the number field

As for the PHS algorithm, the preprocessing phase consists, from a number field K and a size parameter $\omega \in [0, 1/2]$, in building and preparing a lattice L_{tw} that depends only on the number field and allows to express any Approx-id-SVP instance in K as an Approx-CVP instance in L_{tw}.

Theoretically, the only difference between the original PHS preprocessing and ours resides in the lattice definition and in the factor base elaboration. Its most significant part still consists in

computing a hint of constrained size to facilitate forthcoming Approx-Cvp queries. In practice though, we replace this hint computation by merely a few rounds of BKZ with small block size (see §3.4). In a quantum setting this removes the only part that is not polynomial in $\ln|\Delta_K|$, and in a classical setting avoids the dominating exponential part.

Defining the lattice L_{tw}: a full-rank version of the log-$\mathcal{S}$-unit lattice.

Let $\mathrm{FB} = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_k\}$ be a set of prime ideals generating the class group Cl_K. The lattice L_{tw} used by our Twisted-PHS algorithm is basically the log-$\mathcal{S}$-unit lattice $\overline{\mathrm{Log}}_{\mathcal{S}}\,\mathcal{O}_{K,\mathcal{S}}^{\times}$ w.r.t. $\mathcal{S}$, where $\mathcal{S} = \mathcal{S}_\infty \cup \mathrm{FB}$, under the expanded log-$\mathcal{S}$-embedding, to which we apply an isometric transformation to obtain a full-rank lattice in $\mathbb{R}^{\nu+k}$.

Formally, L_{tw} is defined as the lattice generated by the images of the fundamental elements generating the $\mathcal{S}$-unit group $\mathcal{O}_{K,\mathcal{S}}^{\times}$, as given by Th. 2.2, under the following map φ_{tw} from K to $\mathbb{R}^{\nu+k}$:

$$\varphi_{\mathsf{tw}}(\alpha) = f_H \circ \pi_H\big(\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha\big), \tag{3.15}$$

- where f_H is an isometry from $H \subset \mathbb{R}^{n+k}$ to $\mathbb{R}^{\nu+k}$, with H the intersection of the trace zero hyperplane $\mathbb{R}_0^{n+k} = \mathbf{1}_{n+k}^{\perp}$, and $\mathcal{L} = \big\{\mathbf{y} \in \mathbb{R}^{n+k} : y_{r_1+2i-1} = y_{r_1+2i}, i \in [\![1, r_2]\!]\big\}$ the span of $\overline{\mathrm{Log}}_{\mathcal{S}}\,K$;
- π_H is the projection on H, in particular it is the identity on the $\mathcal{S}$-unit group.

This map naturally inherits from the homomorphism properties of $\overline{\mathrm{Log}}_{\mathcal{S}}$, i.e., $\varphi_{\mathsf{tw}}(\alpha\alpha') = \varphi_{\mathsf{tw}}(\alpha) + \varphi_{\mathsf{tw}}(\alpha')$ and $\forall\lambda \in \mathbb{Z}$, $\varphi_{\mathsf{tw}}(\alpha^\lambda) = \lambda \cdot \varphi_{\mathsf{tw}}(\alpha)$, and also defines an isomorphism between $\mathcal{O}_{K,\mathcal{S}}^{\times}/\mu\big(\mathcal{O}_K^{\times}\big)$ and L_{tw}.

The isometry f_H must be carefully chosen in order to control its effect on the ℓ_∞-norm. Nevertheless, it should be seen as a technicality allowing to work with tools designed for full-rank lattices. Formally, let f_H be the linear map represented by $\overline{\mathrm{GSO}}^{\mathrm{T}}(M_H)$, which denotes the transpose of the Gram-Schmidt orthonormalization of the following matrix:

$$M_H := \begin{pmatrix} \overbrace{\begin{matrix} -1 & 1 & & & \\ & -1 & 1 & & \\ & & \ddots & \ddots & \\ & & & -1 & 1 \end{matrix}}^{\nu+1+k} \\ \nu+k \end{pmatrix} \cdot \begin{pmatrix} \overbrace{I_{r_1}}^{r_1} & \overbrace{\begin{matrix} \frac{1}{2} & \frac{1}{2} & & \\ & \frac{1}{2} & \frac{1}{2} & \\ & & \ddots & \ddots \\ & & & \frac{1}{2} \end{matrix}}^{2r_2} & \overbrace{\begin{matrix} \\ \\ \\ \frac{1}{2} \end{matrix}}^{k} \\ r_1 & r_2 & k \\ & & I_k \end{pmatrix}. \tag{3.16}$$

Actually, M_H is a basis of $H = \mathbb{R}_0^{n+k} \cap \mathcal{L}$ in $\mathbb{R}^{n+k}$, constituted of vectors that are orthogonal to $\mathbf{1}_{n+k}$ and to each of the r_2 independent vectors $\mathbf{v}_j, j \in [\![1, r_2]\!]$ that sends any $\mathbf{y} \in \mathcal{L}$ to $\mathbf{0}$ by substracting y_{r_1+2j} from its copy y_{r_1+2j-1} and forgetting every other coordinate. Hence,

graphically, a row basis of L_{tw} is:

$$B_{L\mathsf{tw}} := \left[\begin{array}{c|c} \widetilde{\Lambda}_K & 0 \\ \hline \overline{\mathrm{Log}}\,\varepsilon_{\nu+1} & \\ \vdots & \left(-v_{\mathfrak{p}_j}(\varepsilon_{\nu+i})\ln\mathcal{N}(\mathfrak{p}_j)\right)_{1\le i,j\le k} \\ \overline{\mathrm{Log}}\,\varepsilon_{\nu+k} & \end{array}\right] \cdot \overline{\mathrm{GSO}}^{\mathrm{T}}(M_H), \tag{3.17}$$

where the first part is the basis $\widetilde{\Lambda}_{K,\mathcal{S}}$ of $\overline{\mathrm{Log}}_{\mathcal{S}}\,\mathcal{O}_{K,\mathcal{S}}^{\times}$ defined in Pr. 2.8.

Volume of L_{tw} and optimal factor base choice.

First, we evaluate the volume of $L_{\mathsf{tw}} = f_H\big(\overline{\mathrm{Log}}_{\mathcal{S}}\,\mathcal{O}_{K,\mathcal{S}}^{\times}\big)$. As the isometry f_H stabilizes the span of the log-$\mathcal{S}$-unit lattice, it preserves its volume, which is given by Pr. 2.8. Using that ideal classes of FB generate the class group, hence $h_{K,(\mathrm{FB})} = h_K$, yields:

$$\mathrm{Vol}\,L_{\mathsf{tw}} = \sqrt{n+k}\cdot 2^{-r_2/2}\cdot h_K R_K \prod_{1\le i\le k}\ln\mathcal{N}(\mathfrak{p}_i). \tag{3.18}$$

Certainly, the volume of L_{tw} is growing with the log norms of the factor base prime ideals, but a remarkable property is that this growth is at first slower than the lattice density increase induced by the bigger dimension. The meaning of this is that we can enlarge the factor base to densify our lattice up to an optimal point, after which including new ideals becomes counter-productive.

Formally, let $V_{k'}$ denote the *reduced* volume $\mathrm{Vol}^{1/(\nu+k')}\,L_{\mathsf{tw}}$ for a factor base of size $k' \ge k_0$, where k_0 is the number of generators of Cl_K. We have:

$$V_{k'+1} = V_{k'}\cdot\left(\sqrt{1+\tfrac{1}{n+k'}}\cdot\frac{\ln\mathcal{N}(\mathfrak{p}_{k'+1})}{V_{k'}}\right)^{1/(\nu+k'+1)}. \tag{3.19}$$

This shows that $V_{k'+1} < V_{k'}$ is equivalent to $\ln\mathcal{N}(\mathfrak{p}_{k'+1}) < V_{k'}\Big/\sqrt{1+\frac{1}{n+k'}}$. Using this property, Alg. 3.3 outputs a factor base maximizing the density of L_{tw}.

Algorithm 3.3 Twisted-PHS Factor Base Choice $\mathcal{A}_{\mathsf{tw\text{-}FB}}$

Input: A number field K of degree n.
Output: An optimal factor base FB generating Cl_K that minimizes $\mathrm{Vol}^{1/(\nu+k)}\,L_{\mathsf{tw}}$.

1: Compute $\mathrm{Cl}_K = \langle[\mathfrak{q}_1],\ldots,[\mathfrak{q}_{k_0}]\rangle$, with $k_0 \le \log h_K$.
2: Compute $\mathcal{P}(B) = \{\mathfrak{p}_i : \mathcal{N}(\mathfrak{p}_i) \le B\} \setminus \{\mathfrak{q}_1,\ldots,\mathfrak{q}_{k_0}\}$ ordered by increasing norms, where B is chosen s.t. $\pi_K(B) = \mathrm{poly}(\ln|\Delta_K|) \ge k_0$.
3: FB $\leftarrow \{\mathfrak{q}_1,\ldots,\mathfrak{q}_{k_0}\}$.
4: $i \leftarrow 0$.
5: **while** $\ln\mathcal{N}(\mathfrak{p}_{i+1}) < V_{k_0+i}\Big/\sqrt{1+\frac{1}{n+k_0+i}}$ **do**
6: Add $\mathfrak{p}_{i+1}$ to FB.
7: $i \leftarrow i+1$.
8: **return** FB.

First, for a fixed factor base of size k, we compare the reduced volume V_k of L_{tw} with the reduced volume of L_{phs}, denoted $V_{\mathsf{phs}} := \left(\sqrt{\frac{n}{2^{r_2}}}\cdot h_K R_K\right)^{1/(\nu+k)}$.

Lemma 3.20.

$$\frac{V_k}{V_{phs}} \leq \frac{e^{1/ne}}{k} \cdot \sum_{\mathfrak{p}\in \mathrm{FB}} \ln \mathcal{N}(\mathfrak{p}).$$

This means that the gap between the reduced volume of the twisted lattice and the reduced volume of the untwisted lattice evolves roughly as the arithmetic mean of the $\ln \mathcal{N}(\mathfrak{p})$. We stress that this bound is valid for *any* k, and remark that $e^{1/ne} \leq e^{1/2e} \approx 1.202$.

Proof. The quotient V_k/V_{phs} is $\left(\sqrt{\frac{n+k}{n}} \prod \ln \mathcal{N}(\mathfrak{p})\right)^{1/(\nu+k)}$. The square root power is bounded by $\left(\frac{n+k}{n}\right)^{1/(n+k)}$, as $\frac{1}{\nu+k} < \frac{2}{n+k}$, which reaches when $k+n = ne$ its maximum value $e^{1/ne}$. On the other hand, $\frac{1}{\nu+k} < \frac{1}{k}$, thus by Jensen's inequality:

$$\Big(\prod_{\mathfrak{p}\in \mathrm{FB}} \ln \mathcal{N}(\mathfrak{p})\Big)^{1/(\nu+k)} \leq \Big(\prod_{\mathfrak{p}\in \mathrm{FB}} \ln \mathcal{N}(\mathfrak{p})\Big)^{1/k} \leq \frac{1}{k} \cdot \sum_{\mathfrak{p}\in \mathrm{FB}} \ln \mathcal{N}(\mathfrak{p}). \qquad \square$$

Although the reduced volume significantly decreases in the first loop iterations, reaching precisely the minimum value can be very gradual, so that it might be clever to early abort the loop in Alg. 3.3 when the gradient is too low, or truncate the output to at most $k' = \widetilde{O}(\ln|\Delta_K|)$. We quantify the fact that the density loss is at most constant in the worst case in the following result.

Lemma 3.21. *Let $k' = C\big(\ln|\Delta_K| + n \ln\ln|\Delta_K|\big)$. Let V_{min} be the minimum reduced volume output by $\mathcal{A}_{\mathsf{tw\text{-}FB}}$, and suppose V_{min} is reached for $k > k'$, then:*

$$V_{k'} \leq e^{1/C+1/ne} \cdot V_{min}.$$

Proof. By Eq. (2.35), this choice of k' implies $\left(\sqrt{\frac{n}{2^{r_2}}} \cdot h_K R_K\right)^{1/(\nu+k')} \leq e^{1/C}$. Lemma 3.20 thus gives $V_{k'} \leq e^{1/C+1/ne} \ln \mathcal{N}(\mathfrak{p}_{k'})$. The result follows from the fact that by design, $\ln \mathcal{N}(\mathfrak{p}_{k'}) \leq V_{\mathsf{min}} \leq V_{k'}$. $\square$

In practice, experiments of §3.4 report that the factor bases output by $\mathcal{A}_{\mathsf{tw\text{-}FB}}$ have significantly smaller dimensions than the dimensions showed in Tab. 3.1 and 3.2 for the (optimized) PHS algorithm, so that Lem. 3.21 is never triggered.

Proposition 3.22. *Algorithm $\mathcal{A}_{\mathsf{tw\text{-}FB}}$ terminates in time $\mathrm{T}_{\mathsf{Su}}(K) + \mathrm{poly}(\ln|\Delta_K|)$ and outputs a factor base of size $k = \mathrm{poly}(\ln|\Delta_K|)$ using $B = \mathrm{poly}(\ln|\Delta_K|)$.*

Proof. We first show termination. If $\ln \mathcal{N}(\mathfrak{p}_1) \geq V_{k_0} \Big/ \sqrt{1 + \frac{1}{n+k_0}}$, the algorithm stops. Otherwise, by Eq. (3.19), $V_{k_0+i+1} < V_{k_0+i}$ at best until $\ln \mathcal{N}(\mathfrak{p}_{i+1}) \geq V_{k_0+i}$. Since there are at most n prime ideals of a given norm, $\ln \mathcal{N}(\mathfrak{p}_i)$ must increase, so that at some point $V_{k_0+i+1} > V_{k_0+i}$, where the density of L_{tw} decreases.

We now bound B and k. For $C > 0$, let $k' = C\big(\ln|\Delta_K| + n \ln\ln|\Delta_K|\big)$, and let $B' = \mathcal{N}(\mathfrak{p}_{k'})$. By the Prime Ideal Theorem (Th. 2.38), $B' \leq \mathrm{poly}(\ln|\Delta_K|)$. Using the same arguments as in the proof of Lem. 3.21, we obtain:

$$V_{k'} \leq e^{1/C} \cdot \left(\sqrt{\tfrac{n+k'}{n}}\right)^{1/(\nu+k')} \cdot \ln^{k'/(\nu+k')} \mathcal{N}(\mathfrak{p}_{k'}) \leq e^{1/C+1/ne} \cdot \ln B'.$$

If $\ln \mathcal{N}(\mathfrak{p}_{k'+1}) \geq V_{k'}$, we take $B = B'$ and $k = k'$. Note that this is generically the case in practice. Otherwise, it is necessary to increase B' to at most $B = \ell B'$, with $\ell = \exp\big(e^{1/C+1/ne}\big)$. This value of ℓ verifies that if $k > k'$ is such that $\mathcal{N}(\mathfrak{p}_{k+1}) \geq B \geq \mathcal{N}(\mathfrak{p}_k)$, then $\ln \mathcal{N}(\mathfrak{p}_{k+1}) \geq V_{k'} > V_k$,

and by definition $\sharp\mathrm{FB} \leq k$. Note that this scaling value ℓ is small, e.g., for $C \geq 4$ and $n \geq 3$ we have $\ell \leq 4$. The key is now to show that this new $k = \pi_K(\ell B')$ is not much larger than $k' = \pi_K(B')$. Actually, provided B' is (polynomially in $\ln|\Delta_K|$) large enough, invoking again the Prime Ideal Theorem yields $k' = \pi_K(B') \geq \frac{B'}{2\ln B'}$ [BDPW20, Lem. A.3] and:

$$k \leq \pi_K(\ell B') \leq \tfrac{2n(\ell B')}{\ln \ell B'} = (4\ell n) \cdot \tfrac{B'}{2\ln B'} \leq (4\ell n) \cdot \pi_K(B') = \mathrm{poly}(\ln|\Delta_K|).$$

Note that Bach's bound (Eq. (2.36)) is $\mathrm{poly}(\ln|\Delta_K|)$, as B and k. Therefore, steps 2–7 run in time $\mathrm{poly}(\ln|\Delta_K|)$, and step 1 computes Cl_K in time $\mathrm{T}_{\mathsf{Su}}(K)$. □

Preprocessing algorithm.

Algorithm 3.4 details the complete preprocessing procedure that, from a number field and some precomputation size parameter, chooses a factor base FB, builds the associated matrix $B_{L_{\mathsf{tw}}}$, and processes L_{tw} in order to facilitate Approx-CVP queries.

Algorithm 3.4 Twisted-PHS Preprocessing $\mathcal{A}_{\mathsf{tw\text{-}pcmp}}$

Input: A number field K of degree n and a parameter $\omega \in [0, 1/2]$ or b.
Output: The basis $B_{L_{\mathsf{tw}}}$ with the preimages of its rows in $\mathcal{O}_{K,\mathcal{S}}^{\times}$, and Laarhoven's hint $\mathcal{V}(L_{\mathsf{tw}})$.
1: Get an optimal factor base $\mathrm{FB} = \mathcal{A}_{\mathsf{tw\text{-}FB}}(K)$ of size $k = \sharp\mathrm{FB}$. If needed, truncate the output to $k = \widetilde{O}(\ln|\Delta_K|)$ as in Lem. 3.21.
2: Compute fundamental elements $\varepsilon_1, \ldots, \varepsilon_{\nu+k}$ of $\mathcal{O}_{K,\mathcal{S}}^{\times}$ as in Th. 2.2.
3: Create $B_{L_{\mathsf{tw}}}$, whose rows are $\varphi_{\mathsf{tw}}(\varepsilon_1), \ldots, \varphi_{\mathsf{tw}}(\varepsilon_{\nu+k})$ as defined in Eq. (3.17).
4: Use Laarhoven's algorithm to compute a hint $\mathcal{V} = \mathcal{V}(L_{\mathsf{tw}})$ of size $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$.
5: (or) Use a BKZ of small block size to reduce the basis of L_{tw}.
6: **return** $\big(\mathcal{O}_{K,\mathcal{S}}^{\times}, B_{L_{\mathsf{tw}}}, \mathcal{V}(L_{\mathsf{tw}})\big)$.

This Twisted-PHS preprocessing differs from the original PHS preprocessing given in Alg. 3.1 on two aspects: the factor base, output by $\mathcal{A}_{\mathsf{tw\text{-}FB}}$ in step 1 and which is essentially much smaller in practice, and the new twisted lattice in step 3.

The last two alternative steps consists in preprocessing L_{tw} in order to solve Approx-CVP instances efficiently. Theoretically, we retain in step 4 the same approach as in step 6 of the original PHS preprocessing Alg. 3.1, that guarantees a hint size not exceeding the query phase time using Laarhoven's algorithm [Laa16]. This outputs a hint $\mathcal{V}$ of bit size bounded by $2^{\widetilde{O}(\nu+k)^{1-2\omega}}$, i.e., $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$ using $(\nu+k) = \widetilde{O}(\log|\Delta_K|)$, allowing to deliver the answer for approximation factors $(\nu+k)^{\omega}$ in time bounded by the bit size of $\mathcal{V}$ [Laa16, Cor. 1–2]. This theoretical version will be denoted by $\mathcal{A}^{(\mathrm{Laa})}_{\mathsf{tw\text{-}pcmp}}$.

Nevertheless, in practice the twisted lattice output by Alg. 3.4 incidentally appears to be a lot more orthogonal than expected. That's the reason why we suggest to replace the exponential step 4 of Alg. 3.4 by step 5, which performs some polynomial lattice reduction using a small block size BKZ. In a quantum setting this removes the only part that is not polynomial in $\ln|\Delta_K|$, and in a classical setting avoids the dominating exponential part. This practical version will be denoted by $\mathcal{A}^{(\mathrm{bkz})}_{\mathsf{tw\text{-}pcmp}}$.

Proof of the first part of Th. 3.14. The complexity of step 1 is given by Pr. 3.22. Neglecting terms in $\mathrm{poly}(\ln|\Delta_K|)$, the other costly steps are steps 2 and 4. The former costs $\mathrm{T}_{\mathsf{Su}}(K) \leq 2^{\widetilde{O}(\log^{2/3}|\Delta_K|)}$ by §2.3.3; the latter, independently of ω, runs in $2^{O(\nu+k)} = 2^{\widetilde{O}(\log|\Delta_K|)}$ by the bound on k. Hence, Alg. 3.4 has the same complexity as the original PHS preprocessing, i.e., at

most $2^{\widetilde{O}(\log|\Delta_K|)}$. Note that in practice, the dimension of L_{tw} is much smaller than the one of L_{phs}, which directly lowers the practical complexity of $\mathcal{A}^{(\mathrm{Laa})}_{\mathsf{tw\text{-}pcmp}}$ and $\mathcal{A}^{(\mathrm{bkz})}_{\mathsf{tw\text{-}pcmp}}$. □

3.3.2 Query phase

This section describes the query phase $\mathcal{A}_{\mathsf{tw\text{-}query}}$ of the Twisted-PHS algorithm. As for the query phase of the original PHS algorithm, it reduces the resolution of Approx-id-Svp in $\mathfrak{b}$, for any challenge ideal $\mathfrak{b} \subseteq K$ having a polynomial description in $\log|\Delta_K|$, to a single call to an Approx-Cvp oracle in L_{tw} as output by the preprocessing phase. The main idea of this reduction remains to multiply the principal ideal generator output by the ClDlp of $\mathfrak{b}$ on FB by elements of $\mathcal{O}^{\times}_{K,\mathcal{S}}$ until we reach a principal ideal having a short generator. This translates into adding vectors of L_{tw} to some target vector derived from $\mathfrak{b}$ until the result is short, hence into solving a Cvp instance in the log-$\mathcal{S}$-unit lattice L_{tw}.

The essential difference of the Twisted-PHS version lies in the definition of this target, which is adapted in order to benefit from the twisted description of the log-$\mathcal{S}$-unit lattice. This is formalized in Alg. 3.5.

Algorithm 3.5 Twisted-PHS Query $\mathcal{A}_{\mathsf{tw\text{-}query}}$

Input: Challenge $\mathfrak{b}$, $\mathcal{A}_{\mathsf{tw\text{-}pcmp}}(K,\omega) = \left(\mathcal{O}^{\times}_{K,\mathcal{S}}, B_{L\mathsf{tw}}, \mathcal{V}\right)$, and $\widetilde{\beta} > 0$ s.t. for any $\mathbf{t}$, the Approx-Cvp oracle using $\mathcal{V}(L_{\mathsf{tw}})$ outputs $\mathbf{w} \in L_{\mathsf{tw}}$ with $\|f_H^{-1}(\mathbf{t}-\mathbf{w})\|_\infty \leq \widetilde{\beta}$.

Output: A short element $x \in \mathfrak{b} \setminus \{0\}$.

1: Solve the ClDlp for $\mathfrak{b}$ on FB, i.e., find $\alpha \in K$ s.t. $\langle\alpha\rangle = \mathfrak{b} \cdot \prod_{\mathfrak{p}_i \in \mathrm{FB}} \mathfrak{p}_i^{v_i}$, for $v_i \in \mathbb{Z}$.

2: Define the target $\mathbf{t}$ as $f_H^{-1}(\mathbf{t}) = \pi_H\left(\overline{\mathrm{Log}}\,\alpha, \left\{-v_i \ln \mathcal{N}(\mathfrak{p}_i)\right\}_{1\leq i\leq k}\right) + \mathbf{b}_{\mathsf{tw}}$, where the drift vector $\mathbf{b}_{\mathsf{tw}} \in H$ will be defined in Eq. (3.23).

3: Solve Approx-Cvp with $\mathcal{V}(L_{\mathsf{tw}})$ to get $\mathbf{w} \in L_{\mathsf{tw}}$ s.t. $\|f_H^{-1}(\mathbf{t}-\mathbf{w})\|_\infty \leq \widetilde{\beta}$.

4: (or) Use Babai's Nearest Plane to get $\mathbf{w} \in L_{\mathsf{tw}}$ s.t. $\|f_H^{-1}(\mathbf{t}-\mathbf{w})\|_\infty$ is small.

5: Compute $s = \varphi_{\mathsf{tw}}^{-1}(\mathbf{w}) \in \mathcal{O}^{\times}_{K,\mathcal{S}}$, using the preimages of the rows of $B_{L\mathsf{tw}}$.

6: **return** α/s.

Note that the output of the ClDlp in step 1 is not an $\mathcal{S}$-unit unless $\mathfrak{b}$ is divisible only by prime ideals of FB; for each i, $v_i = v_{\mathfrak{p}_i}(\alpha) - v_{\mathfrak{p}_i}(\mathfrak{b})$. For convenience and without any loss of generality we shall assume that $\mathfrak{b}$ is coprime with *all* elements of the factor base, i.e., $\forall \mathfrak{p} \in \mathrm{FB}$, $v_{\mathfrak{p}}(\mathfrak{b}) = 0$. In that case, the target in step 2 writes naturally as $\mathbf{t} = \varphi_{\mathsf{tw}}(\alpha) + f_H\left(\mathbf{b}_{\mathsf{tw}}\right)$. This target definition calls for a few comments. First, the output of the ClDlp is projected on the whole log-$\mathcal{S}$-unit lattice instead of only on the log-unit sublattice, hence maintaining its length and algebraic norm logarithms in the instance scope. Thus, the way our algorithm uses $\mathcal{S}$-units to reduce the solution of the ClDlp problem can be seen as a smooth generalization of the way traditional Sgp solvers use regular units to reduce the solution of the Pip as in [CDPR16]. Second, the sole purpose of the drift by $\mathbf{b}_{\mathsf{tw}}$ is to ensure that $\alpha/s \in \mathfrak{b}$. Adapting its definition to the twisted setting is slightly tedious and deferred to the next paragraph. The most notable novelty is that we force the use of a drift that is *inside* the log-$\mathcal{S}$-unit lattice span. This somehow captures and compensates for the perturbation induced on infinite places for correcting negative valuations on finite places using $\mathcal{S}$-units.

Finally, as already mentioned, L_{tw} seems much more orthogonal *in practice* than expected, so that we advise to resort to Babai's Nearest Plane algorithm for solving Approx-Cvp in L_{tw}, instead of using Laarhoven's query phase with the precomputed hint. We only keep Laarhoven's algorithm to theoretically prove the correctness and complexity of our new algorithm. The theoretical and practical versions of $\mathcal{A}_{\mathsf{tw\text{-}query}}$ are respectively denoted by $\mathcal{A}^{(\mathrm{Laa})}_{\mathsf{tw\text{-}query}}$ and $\mathcal{A}^{(\mathrm{np})}_{\mathsf{tw\text{-}query}}$.

We now detail explicitly our target choice and prove the correctness and the output quality of Alg. 3.5.

Definition of the target vector.

Recall that we assumed that $\mathfrak{b}$ is coprime with FB, hence $f_H^{-1}(\mathbf{t}) = \pi_H\big(\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha\big) + \mathbf{b}_{\mathsf{tw}}$, for some $\mathbf{b}_{\mathsf{tw}} \in H$ that must ensure $\alpha/s \in \mathfrak{b}$, for $s = \varphi_{\mathsf{tw}}^{-1}(\mathbf{w})$ and when $\|f_H^{-1}(\mathbf{t}-\mathbf{w})\|_\infty \leq \widetilde{\beta}$. Indexing coordinates by places, we exhibit $\mathbf{b}_{\mathsf{tw}} = \big(\{b_\sigma\}_{\sigma\in\mathcal{S}_\infty\cup\overline{\mathcal{S}}_\infty}, \{b_\mathfrak{p}\}_{\mathfrak{p}\in\mathrm{FB}}\big)$, where:

$$\begin{cases} b_\sigma = -\frac{k}{n}\left(\frac{\ln\mathcal{N}(\mathfrak{b})}{n+k} + \widetilde{\beta}\right) + \frac{1}{n}\sum_{\mathfrak{p}\in\mathrm{FB}}\ln\mathcal{N}(\mathfrak{p}) & \text{for } \sigma \in \mathcal{S}_\infty \cup \overline{\mathcal{S}}_\infty, \\ b_\mathfrak{p} = \widetilde{\beta} - \ln\mathcal{N}(\mathfrak{p}) + \frac{\ln\mathcal{N}(\mathfrak{b})}{n+k} & \text{for } \mathfrak{p} \in \mathrm{FB}. \end{cases} \tag{3.23}$$

It is easy to verify that all coordinates sum to 0, i.e., $\mathbf{b}_{\mathsf{tw}} \in H$. We now explain this choice, first showing that under the above hypotheses, Alg. 3.5 is correct.

Proposition 3.24. *Given access to an Approx-*CVP *oracle that on any input* $\mathbf{t}$*, outputs* $\mathbf{w} \in L_{tw}$ *s.t.* $\|f_H^{-1}(\mathbf{t}-\mathbf{w})\|_\infty \leq \widetilde{\beta}$*,* $\mathcal{A}_{\textit{tw-query}}$ *outputs* $x \in \mathfrak{b}\setminus\{0\}$*.*

Proof. Recall that $x = \alpha/s$, where $s = \varphi_{\mathsf{tw}}^{-1}(\mathbf{w}) \in \mathcal{O}_{K,\mathcal{S}}^\times$ and that for the sake of clarity, $\mathfrak{b}$ is taken coprime to FB. Therefore, it is sufficient to show that for any fixed $\mathfrak{p} \in \mathrm{FB}$, $v_\mathfrak{p}(\alpha/s) \geq v_\mathfrak{p}(\mathfrak{b}) = 0$. Indexing coordinates of $\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha$ by places and using the simplified notation $\alpha_v := \big(\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha\big)_v$, we have that for $\mathbf{h}_\alpha = \pi_H\big(\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha\big)$, $(\mathbf{h}_\alpha)_\mathfrak{p} = \alpha_\mathfrak{p} - \frac{\ln\mathcal{N}(\mathfrak{b})}{n+k}$. By hypothesis:

$$\left|\alpha_\mathfrak{p} - \tfrac{\ln\mathcal{N}(\mathfrak{b})}{n+k} - s_\mathfrak{p} + b_\mathfrak{p}\right| = \left|-\big(v_\mathfrak{p}(\alpha) - v_\mathfrak{p}(s) + 1\big)\ln\mathcal{N}(\mathfrak{p}) + \widetilde{\beta}\right| \leq \widetilde{\beta}.$$

Rearranging terms, and using that $v_\mathfrak{p}(\cdot) \in \mathbb{Z}$ to round integers towards 0:

$$0 \leq v_\mathfrak{p}(\alpha/s) \leq \left\lfloor \frac{2\widetilde{\beta}}{\ln\mathcal{N}(\mathfrak{p})} - 1 \right\rfloor.$$

This concludes the correctness proof. □

The proof of Pr. 3.24 quantifies the intuition that the output element has smaller valuations at big norm prime ideals. In particular, strictly positive valuations occur only for ideals s.t. $\ln\mathcal{N}(\mathfrak{p}) \leq \widetilde{\beta}$. This has a very valuable consequence: estimating the ℓ_∞-norm covering radius of L_{tw} allows to control the prime ideal support of any optimal solution. Hence, even if the Approx-CVP cannot reach $\mu_\infty(L_{\mathsf{tw}})$, it is possible to confine the algebraic norm of each query output by *not* including in FB the prime ideals whose log-norm would *in fine* exceed $\mu_\infty(L_{\mathsf{tw}})$, and at which the optimal solution provably has a null valuation. Roughly speaking, this is what $\mathcal{A}_{\mathsf{tw\text{-}FB}}$ tends to achieve in Alg. 3.3.

Translating infinite coordinates. As already mentioned, one important novelty consists in forcing the drift used to ensure $\alpha/s \in \mathfrak{b}$ to be *inside* the log-$\mathcal{S}$-unit span. The underlying intuition is that "correcting" negative valuations at finite primes should only involve $\mathcal{S}$-units. We modelize this by splitting the weight of the $b_\mathfrak{p}$'s evenly across the infinite places coordinates, hence obtaining Eq. (3.23). This heuristically presumes that $\mathcal{S}$-units absolute value logarithms are generically balanced on infinite places. Let us summarize our target definition:

$$\mathbf{t} = f_H\left(\left\{\alpha_\sigma - \tfrac{1}{n}\left[k\widetilde{\beta} + \ln\mathcal{N}(\mathfrak{b}) - \textstyle\sum_{\mathfrak{p}\in\mathrm{FB}}\ln\mathcal{N}(\mathfrak{p})\right]\right\}_\sigma, \left\{\alpha_\mathfrak{p} + \widetilde{\beta} - \ln\mathcal{N}(\mathfrak{p})\right\}_{\mathfrak{p}\in\mathrm{FB}}\right). \tag{3.25}$$

Quality of the output of $\mathcal{A}^{(\mathbf{Laa})}_{\mathsf{tw\text{-}query}}$.

To bound the quality of the output of Alg. 3.5, the general idea is that minimizing the distance of our target to the twisted lattice directly minimizes the p-adic absolute values $-v_{\mathfrak{p}}(\alpha)\ln\mathcal{N}(\mathfrak{p})$ instead of minimizing the valuations $v_{\mathfrak{p}}(\alpha)$ independently of $\ln\mathcal{N}(\mathfrak{p})$.

This makes use of the following log-$\mathcal{S}$-unit lattice structure lemma, adapting its log-unit lattice classical equivalent [PHS19a, Lem. 2.11–2.12], [CDPR16, §6.1]:

Lemma 3.26. *For $\alpha \in K$, let $\mathbf{h}_\alpha := \pi_H\big(\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha\big)$. Decompose $\langle\alpha\rangle$ on* FB *as $\mathfrak{b}\cdot\prod_{\mathfrak{p}\in\mathrm{FB}}\mathfrak{p}^{v_{\mathfrak{p}}(\alpha)}$, with $\mathfrak{b}$ coprime to* FB. *Then $\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha = \mathbf{h}_\alpha + \frac{\ln\mathcal{N}(\mathfrak{b})}{n+k}\cdot\mathbf{1}_{n+k}$. Furthermore, the length of α is bounded by:*

$$\|\alpha\|_2 \le \sqrt{n}\cdot\mathcal{N}(\mathfrak{b})^{1/(n+k)}\cdot\exp\Big[\max_{1\le j\le n}(\mathbf{h}_\alpha)_j\Big].$$

Note that using the max of the coordinates of $\mathbf{h}_\alpha$ instead of its ℓ_∞-norm norm acknowledges for the fact that logarithms of small infinite valuations can become large negatives that should be ignored when evaluating the length of α.

Proof. By definition of the orthogonal projection on H, $\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha$ decomposes as $\mathbf{h}_\alpha + a\cdot\mathbf{1}_{n+k}$, with $a = \big\langle\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha, \mathbf{1}_{n+k}\big\rangle/\|\mathbf{1}_{n+k}\|_2^2$. The scalar product is:

$$\sum_{\sigma\in\mathcal{S}_\infty\cup\overline{\mathcal{S}}_\infty}\ln|\sigma(\alpha)| - \sum_{\mathfrak{p}\in\mathrm{FB}} v_{\mathfrak{p}}(\alpha)\cdot\ln\mathcal{N}(\mathfrak{p}) = \ln\mathcal{N}\Big(\langle\alpha\rangle\Big/\prod\nolimits_{\mathfrak{p}\in\mathrm{FB}}\mathfrak{p}^{v_{\mathfrak{p}}(\alpha)}\Big) = \ln\mathcal{N}(\mathfrak{b}).$$

Therefore, $a = \frac{\ln\mathcal{N}(\mathfrak{b})}{n+k}$. Moreover, generically we have $\|\alpha\|_2 \le \sqrt{n}\cdot\|\alpha\|_\infty$; using the above decomposition coordinate-wise, the j-th-coordinate of $\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha$ writes $\big(\overline{\mathrm{Log}}_{\mathcal{S}}\,\alpha\big)_j = (\mathbf{h}_\alpha)_j + \frac{\ln\mathcal{N}(\mathfrak{b})}{n+k}$ and thus:

$$\|\alpha\|_\infty = \exp\max_{\sigma\in\mathcal{S}_\infty}\ln|\sigma(\alpha)| \le \exp\Big[\tfrac{\ln\mathcal{N}(\mathfrak{b})}{n+k} + \max_{1\le j\le n}(\mathbf{h}_\alpha)_j\Big].$$ □

Theorem 3.27. *Given access to an Approx-*Cvp *oracle that on any input $\mathbf{t}$, outputs $\mathbf{w}\in L_{tw}$ s.t. $\|f_H^{-1}(\mathbf{t}-\mathbf{w})\|_\infty \le \widetilde{\beta}$, $\mathcal{A}_{\mathsf{tw\text{-}query}}$ computes $x\in\mathfrak{b}\setminus\{0\}$ such that:*

$$\|x\|_2 \le \sqrt{n}\cdot\mathcal{N}(\mathfrak{b})^{1/n}\cdot\exp\left[\frac{(n+k)\widetilde{\beta} - \sum_{\mathfrak{p}\in\mathrm{FB}}\ln\mathcal{N}(\mathfrak{p})}{n}\right].$$

This outperforms the bound of Pr. 3.13 if $(n+k)\cdot\widetilde{\beta} \le 2\beta\cdot\sum_{\mathfrak{p}\in\mathrm{FB}}\ln\mathcal{N}(\mathfrak{p})$. In particular, this is implied by Lem. 3.20 if $\frac{\widetilde{\beta}}{\beta} \approx \frac{V_k}{V_{\mathrm{phs}}}$ for $k \ge n$. We will see that under some reasonable heuristics, this is indeed the case when using the *same* factor base, and that experiments suggest a much broader gap. One intuitive reason for this behaviour is that the covering radius of our twisted lattice grows at a slower pace than the log-norm of the prime ideals of FB.

Proof. The correctness comes from Pr. 3.24. As before, let $s = \varphi_{\mathsf{tw}}^{-1}(\mathbf{w})$, where $\mathbf{w}$ verifies $\|f_H^{-1}(\mathbf{t}-\mathbf{w})\|_\infty \le \widetilde{\beta}$. It is necessary to bound $\max_{\sigma\in\mathcal{S}_\infty}(\mathbf{h}_{\alpha/s})_\sigma$ in order to invoke Lem. 3.26. Note that $\mathbf{h}_{\alpha/s} = \mathbf{h}_\alpha - \mathbf{h}_s$, hence:

$$(\mathbf{h}_{\alpha/s})_\sigma = \alpha_\sigma - \tfrac{\ln\mathcal{N}(\mathfrak{b})}{n+k} - s_\sigma.$$

Recalling the target definition given in Eq. (3.23), the σ-coordinate of $f_H^{-1}(\mathbf{t}-\mathbf{w})$ writes $\big(\alpha_\sigma - \frac{\ln\mathcal{N}(\mathfrak{b})}{n+k} + b_\sigma\big) - s_\sigma = (\mathbf{h}_{\alpha/s})_\sigma + b_\sigma$, and the promise on $\mathbf{w}$ yields:

$$(\mathbf{h}_{\alpha/s})_\sigma \le \widetilde{\beta} - b_\sigma = \frac{(n+k)\widetilde{\beta} - \sum_{\mathfrak{p}\in\mathrm{FB}}\ln\mathcal{N}(\mathfrak{p})}{n} + \frac{k}{n(n+k)}\cdot\ln\mathcal{N}(\mathfrak{b}).$$

Injecting this bound in Lem. 3.26 using $\frac{1}{n+k} + \frac{k}{n(n+k)} = \frac{1}{n}$ ends the proof. □

Heuristic evaluation of $\widetilde{\beta}$.

Proving the second part of Th. 3.14 necessitates to evaluate $\widetilde{\beta}$. This evaluation rely on several heuristics that adapt heuristics [PHS19a, Heur. 4–6]. We argue that the arguments developed in [PHS19a, §4] to support these heuristics can be transposed to our setting, and both heuristics are validated by experiments in §3.4.

Heuristic 3.28 (Adapted from [PHS19a, Heur. 4]). *The ℓ_∞-norm covering radius of L_{tw} is bounded by $O\big(\mathrm{Vol}^{1/(\nu+k)} L_{\mathsf{tw}}\big)$. Likewise, $\mu_2(L_{\mathsf{tw}}) = O\big(\sqrt{\nu+k} \cdot \mathrm{Vol}^{1/(\nu+k)} L_{\mathsf{tw}}\big)$.*

This assumption relies on L_{tw} to behave like a random lattice, implying its successive minima and covering radius to be even. In [PHS19a], the randomness essentially comes from the choice of the factor base, while for L_{tw}, this choice is deterministic. We argue that heuristically, prime ideals of FB represent uniformly random classes in Cl_K,[16] and $\mathcal{S}$-units Archimedean absolute value logarithms are likely to be uniform in $\mathbb{R}^n/\overline{\mathrm{Log}}\,\mathcal{O}_K^\times$. The volumetric arguments of [PHS19a, §4.1] can also be readily adapted, using $\ln \mathcal{N}(\mathfrak{p}) \leq \mathrm{Vol}^{1/(\nu+k)} L_{\mathsf{tw}}$ by construction.

Heuristic 3.29 (Adapted from [PHS19a, Heur. 5–6]). *With non-negligible probability over the input target vector $\mathbf{t}$, the vector $\mathbf{w}$ output by Laarhoven's algorithm satisfies $\|f_H^{-1}(\mathbf{t}-\mathbf{w})\|_\infty \leq O\big(\ln(n+k)/\sqrt{n+k}\big) \cdot \|\mathbf{t}-\mathbf{w}\|_2$.*

This heuristic conveys the idea that coefficients of the output of Laarhoven's algorithm are somehow balanced, so that $\|\mathbf{w}\|_2 \approx \sqrt{n+k} \cdot \|f_H^{-1}(\mathbf{w})\|_\infty$. Typically, continuous Gaussian vectors $\mathbf{y}$ of dimension d verify $\|\mathbf{y}\|_\infty/\|\mathbf{y}\|_2 = O(\ln d/\sqrt{d})$ with good probability, as shown by [PHS19a, Lem. 4.1]. In our setting, this is justified by assuming $\mathbf{t}$ is uniformly distributed in $\big(\mathbb{R} \otimes L_{\mathsf{tw}}\big)/L_{\mathsf{tw}}$, and can be randomized by multiplying $\mathfrak{b}$ by small ideals coprime to FB.

Proof of the second part of Th. 3.14. It breaks down to plugging into Th. 3.27 a value for k and $\widetilde{\beta}$. Using Lem. 3.21, we take $k = \widetilde{O}(\ln|\Delta_K|)$, so that $V_k = O(\ln \mathcal{N}(\mathfrak{L}_{\mathsf{max}})) = O(\ln\ln|\Delta_K|)$ by Lem. 3.20 and Pr. 3.22. We stress that if $\mathcal{A}_{\mathsf{tw\text{-}FB}}$ terminates with a smaller k, this can by definition only yield a smaller V_k. By Heur. 3.28, it implies $\mu_2(L_{\mathsf{tw}}) = O(\sqrt{\nu+k} \cdot \ln\ln|\Delta_K|)$, and Heur. 3.29 yield on average $\|f_H^{-1}(\mathbf{v})\|_\infty \leq \frac{\ln n+k}{\sqrt{n+k}} \cdot \|\mathbf{v}\|_2$. The Approx-CVP solver from Laarhoven's algorithm using $\mathcal{V}(L_{\mathsf{tw}})$ outputs a lattice vector at Euclidean distance which is at most $O\big((\nu+k)^\omega \cdot \mu_2(L_{\mathsf{tw}})\big)$. Hence, its infinity distance is $\widetilde{O}\big((\nu+k)^\omega \cdot \ln\ln|\Delta_K|\big)$, and $(k+n)\widetilde{\beta} = \widetilde{O}\big((\nu+k)^{\omega+1} \cdot \ln\ln|\Delta_K|\big) = \widetilde{O}\big(\ln^{\omega+1}|\Delta_K|\big)$, as claimed.

As for the running time of Alg. 3.5, it is essentially determined by those of steps 1 and 3. Solving the ClDLP problem requires to compute $\mathcal{S}$-units for an extended factor basis containing FB and prime factors of $\mathfrak{b}$, hence costs $\mathrm{T}_{\mathsf{Su}}(K)$. Note that since it depends on the challenge, this cost cannot be mitigated by some preprocessing effort. On the other hand, solving Approx-CVP with Laarhoven's algorithm runs in time bounded by $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)}$, the size of V. Finally, the total run time of $\mathcal{A}^{(\mathrm{Laa})}_{\mathsf{tw\text{-}query}}$ is bounded by $2^{\widetilde{O}(\log^{1-2\omega}|\Delta_K|)} + \mathrm{T}_{\mathsf{Su}}(K)$. □

In practice, as shown in §3.4, the special properties of our twisted lattice L_{tw} suggest replacing Laarhoven's CVP solving by Babai's Nearest Plane algorithm for solving Approx-CVP in L_{tw}. In this eventuality, $\mathcal{A}^{(\mathrm{np})}_{\mathsf{tw\text{-}query}}$ would become quantumly polynomial, and classically only subexponential in $\ln|\Delta_K|$.

[16] This is at the heart of the analytic class number formula.

3.4 Experimental Data

This is the first time to our knowledge that this type of algorithm is completely implemented and tested for fields of degrees up to 60. As a point of comparison, the experiments of [PHS19a] constructed the log-$\mathcal{S}$-unit lattice L_{phs} for cyclotomic fields of degrees at most 24 and $h_K \leq 3$, all but the last two being principal [PHS19a, Fig. 4.1].

Hardware and library description. All $\mathcal{S}$-units and class group computations, for the log-$\mathcal{S}$-unit lattice description and the ClDlp resolution, were performed using Magma v2.24-10 [BCP97].[17] The BKZ reductions and Cvp/Svp computations used fplll v5.3.2 [FpL16]. All other parts of the experiments rely on SageMath v9.0 [Sag20]. All the sources and scripts are available as supplementary material on GitHub: ob3rnard/Twisted-PHS[7]. The experiments took less than a week on a server with 36 cores and 768 GB RAM.

Number fields. As announced in §2.1.1, we consider two families of number fields, namely non-principal cyclotomic fields $\mathbb{Q}(\zeta_m)$ of prime conductors $m \in [\![23, 71]\!]$, and NTRU Prime fields $\mathbb{Q}(z_q)$ where z_q is a root of $x^q - x - 1$, for $q \in [\![23, 47]\!]$ prime. These correspond to the range of what is feasible in a reasonable amount of time, as the asymptotics of $\mathrm{T}_{\mathsf{Su}}(K)$ rapidly express in a classical setting.

For cyclotomic fields, we managed to compute $\mathcal{S}$-units up to $\mathbb{Q}(\zeta_{71})$ for all factor bases in less than a day, and all log-$\mathcal{S}$-unit lattice variants up to $\mathbb{Q}(\zeta_{61})$. For NTRU Prime fields, we managed all computations up to $\mathbb{Q}(z_{47})$.

Targeted lattices. We evaluate the lattices computed by three algorithms: the original PHS algorithm, as implemented in [PHS19b]; our optimized version Opt-PHS from §3.2.3, and our new twisted variant Twisted-PHS described in §3.3. This yields three different lattices, denoted by resp. L_{phs}, L_{opt} and L_{tw}. There are a few differences between [PHS19a] and its implementation [PHS19b], but we chose to stick to the provided implementation as much as possible.

In order to separate the improvements due to $\mathcal{A}_{\mathsf{tw\text{-}FB}}$ outputting smaller factor bases from those purely induced by our specific use of the product formula to describe the log-$\mathcal{S}$-unit lattice, we also built lattices $L^{(0)}_{\mathsf{phs}}$ and $L^{(0)}_{\mathsf{opt}}$ corresponding to PHS and Opt-PHS algorithms, but using the *same* factor base as L_{tw}.

BKZ reductions and CVP solving. We applied the same reduction strategy to all of our lattices. Namely, lattices of dimension less than 60 were HKZ reduced, while lattices of greater dimension were reduced using at most 300 loops of BKZ with block size 40. This yields reasonably good bases for a small computational cost [CN11, p.2]. The loop limit was in practice never hit.

For Cvp computations, we applied with these reduced bases Babai's Nearest Plane algorithm, as described in [Gal12, §18.1, Alg. 26].

Precision issues. Choosing the right bit precision for floating point arithmetic in the experiments is particularly tricky. We generically used at most 500 bits of precision in our experiments (corresponding to the lattice volume logarithm in base 2 plus some extra margin). There are two notable exceptions:

[17] Note that SageMath is significantly faster than Magma for computing class groups, but behaves surprisingly poorly when it comes to computing $\mathcal{S}$-units.

[7] https://github.com/ob3rnard/Twisted-PHS

1. The $\mathcal{S}$-units w.r.t. FB can have *huge* coefficients. Computing the absolute values of their embeddings must then be performed at very high precision. All our lattice constructions were conducted using 10000 bits of precision.
2. Computing the target involves the challenge and the ClDLP solution, whose coefficients are potentially *huge* rational numbers, up to 2^{25000} for e.g., $\mathbb{Q}(\zeta_{53})$. As above, we adjust the precision in order to obtain sensible values.

In all cases, once in the log space the resulting high precision data can be rounded back to the generic precision before lattice reduction or CVP computations.

3.4.1 Geometric characteristics

First, we evaluated the geometric characteristics of each produced lattice, using indicators recalled in §2.4.2, namely: the root Hermite factor δ_0, the orthogonality defect δ, and the minimum θ_{min} (resp. average θ_{avg}) vector basis angle. Each of these indicators is declined before and after BKZ reduction to compare their evolution. We also evaluated experimentally the relevance of Heur. 3.28 and 3.29. Example results are given in Tab. 3.3 and 3.4 for cyclotomic and NTRU Prime fields, aside the lattices dimensions $d = \nu + k$ and reduced volumes $V^{1/d}$. Extensive data can be found in Tab. 3.5 and 3.6 for both field families.

		d	$V^{1/d}$	δ_0		δ		θ_{min}		θ_{avg}		μ_2	μ_∞	$\Vert\cdot\Vert_\infty/\Vert\cdot\Vert_2$	
				raw	bkz	raw	bkz	raw	bkz	raw	bkz			real	H. 3.29
$\mathbb{Q}(\zeta_{41})$	L_{tw}	59	4.825	1.001	1.001	3.596	1.802	11	47	69	81	12.91	5.186	0.615	0.489
	$L^{(0)}_{\mathsf{opt}}$	59	1.786	1.020	1.005	4.525	1.986	34	55	76	83	5.112	2.245	0.629	0.530
	$L^{(0)}_{\mathsf{phs}}$	59	2.767	1.037	0.997	8.986	1.809	45	55	79	84	8.535	4.039	0.639	0.530
	L_{opt}	103	1.379	1.013	1.006	6.514	2.592	25	48	66	84	5.301	2.052	0.596	0.456
	L_{phs}	144	1.306	1.012	1.004	7.982	3.651	29	49	71	83	6.536	2.772	0.687	0.414

TABLE 3.3 – Geometric characteristics of log-$\mathcal{S}$-unit lattices for some prime conductor cyclotomic fields.

		d	$V^{1/d}$	δ_0		δ		θ_{min}		θ_{avg}		μ_2	μ_∞	$\Vert\cdot\Vert_\infty/\Vert\cdot\Vert_2$	
				raw	bkz	raw	bkz	raw	bkz	raw	bkz			real	H. 3.29
$\mathbb{Q}(z_{43})$	L_{tw}	38	4.441	0.911	0.911	1.498	1.357	53	59	82	83	10.64	5.177	0.645	0.528
	$L^{(0)}_{\mathsf{opt}}$	38	5.051	0.937	0.937	4.187	1.865	44	50	81	81	12.50	6.573	0.663	0.590
	$L^{(0)}_{\mathsf{phs}}$	38	9.657	0.952	0.952	7.496	1.877	45	56	81	81	23.73	12.18	0.671	0.590
	L_{opt}	114	1.367	0.979	0.979	5.482	3.256	36	57	79	83	6.119	2.803	0.687	0.443
	L_{phs}	161	1.297	0.987	0.987	9.002	4.135	25	55	79	83	7.484	2.837	0.712	0.400
$\mathbb{Q}(z_{47})$	L_{tw}	40	4.576	0.913	0.913	1.650	1.358	49	60	82	84	11.04	5.607	0.632	0.519
	$L^{(0)}_{\mathsf{opt}}$	40	6.231	0.938	0.938	4.628	1.915	37	57	81	81	16.59	8.398	0.658	0.583
	$L^{(0)}_{\mathsf{phs}}$	40	12.06	0.951	0.951	7.908	1.946	38	55	81	81	30.85	15.50	0.662	0.583
	L_{opt}	129	1.376	0.981	0.981	6.189	3.632	21	56	80	83	6.575	2.925	0.696	0.427
	L_{phs}	180	1.309	0.989	0.989	10.15	4.527	31	53	80	83	8.022	2.882	0.704	0.387

TABLE 3.4 – Geometric characteristics of log-$\mathcal{S}$-unit lattices for some NTRU Prime fields.

Orthogonality indicators.

We first remark that minimum and average vector basis angles seem difficult to interpret. They are slightly better for Twisted-PHS on NTRU Prime fields but it is harder to extract a general

tendency for cyclotomic fields.

After a light BKZ reduction, twisted lattices show significantly better root Hermite factor and orthogonality defect than any other log-$\mathcal{S}$-unit lattice representations, *even* when the lattices have the same dimension, i.e., when the same factor base is used. Second, the evolution of the orthogonality defect before and after the reduction is more restricted in the twisted case than in the others. In particular, we observe that the BKZ-reduced versions of $L^{(0)}_{\mathsf{opt}}$ and $L^{(0)}_{\mathsf{phs}}$ can have *bigger* orthogonality defects than the *unreduced* L_{tw}. This last observation is true for all NTRU Prime fields we tested except $\mathbb{Q}(z_{23})$.

These two phenomenons (better values and small variations) are particularly clear for NTRU Prime fields. We remark that in this case, the twisted version of the log-$\mathcal{S}$-unit lattice fully expresses, since for NTRU Prime fields most factor base elements have distinct norms. On the contrary, factor bases for our targeted cyclotomic fields are composed of one (or two, as for $\mathbb{Q}(\zeta_{59})$) Galois orbits whose elements all have the same norm. Finally, we stress that reducing L_{tw} lattices is much faster in practice than reducing $L^{(0)}_{\mathsf{opt}}$ and $L^{(0)}_{\mathsf{phs}}$. This is corroborated by the graphs of the Gram-Schmidt log norms in §3.4.2.

Evaluating heuristic on covering radius (Heur. 3.28).

Computing the covering radius of a given lattice is a very difficult problem in general. To evaluate in practice μ_2 and μ_∞ for our computed lattices, we used a slightly modified version of the strategy of [PHS19a, §4.1]. More precisely, for each lattice L, we picked 500 random target vectors $\mathbf{t}_i$ in the span of L from a continuous Gaussian distribution of deviation $\sigma = 100 \cdot \dim L$, then used Babai's Nearest Plane algorithm with the reduced basis of L to obtain vectors $\mathbf{w}_i \in L$ close to $\mathbf{t}_i$. Finally, we majorate $\mu_\infty(L)$ and $\mu_2(L)$ by respectively $\max_i \|\mathbf{t}_i - \mathbf{w}_i\|_\infty$ and $\max_i \|\mathbf{t}_i - \mathbf{w}_i\|_2$.

Results show that all lattices equally match Heur. 3.28. We noticed, for L_{phs} and for the number fields tested in [PHS19a, Fig. 4.1], a significant gap between our estimations and the published numerical values. We stress that using in our code a standard deviation of only $\sigma = 100$ as in [PHS19b] reproduces their results.

Evaluating heuristic on infinity norm (Heur. 3.29).

In order to support Heur. 3.29, we compared the average $\|f_H^{-1}(\mathbf{t}_i - \mathbf{w}_i)\|_\infty / \|\mathbf{t}_i - \mathbf{w}_i\|_2$ with the expected value $\left(\ln(n+k)/\sqrt{n+k}\right)$ for L_{tw}. The evolution of Heur. 3.29 from [PHS19a, Heur. 5–6] is quantified by relating, for all four PHS log-$\mathcal{S}$-unit variants, the ratio $\|\mathbf{t}_i - \mathbf{w}_i\|_\infty / \|\mathbf{t}_i - \mathbf{w}_i\|_2$ to their expected ratio $\left(\ln(\nu+k)/\sqrt{\nu+k}\right)$.

The data show that all lattices follow exactly the same behaviour w.r.t. [PHS19a, Heur. 5–6] and Heur. 3.29. All these values are tagged with a unique label "$\|\cdot\|_\infty/\|\cdot\|_2$ (real/H. 3.29)" in the tables, but correspond to Heur. 3.29 for Twisted-PHS and to [PHS19a, Heur. 5–6] for PHS.

3.4.2 Plotting Gram-Schmidt log norms

For our second experiment, we evaluate the Gram-Schmidt norms of each produced lattice. We propose two comparisons, the first one is before and after BKZ reduction to see the evolution of the norms for each case at iso factor base in Fig. 3.2; the second one is between the different lattices (after BKZ reduction) in Fig. 3.3. Again, extensive data for other examples can be found in §3.5.2 for both cyclotomic fields and NTRU Prime fields.

We first remark that in Fig. 3.2 the two curves, before and after BKZ reduction, are almost superposed for the Twisted-PHS lattice. This does not seem to be the case for the two other PHS variants we consider here.

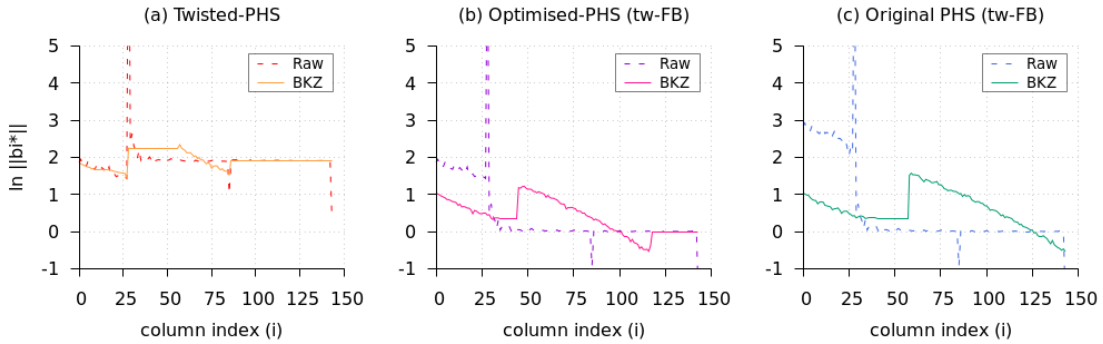

FIGURE 3.2 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{59})$: Gram-Schmidt log norms before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for: (a) L_{tw}; (b) $L^{(0)}_{\text{opt}}$; (c) $L^{(0)}_{\text{phs}}$.

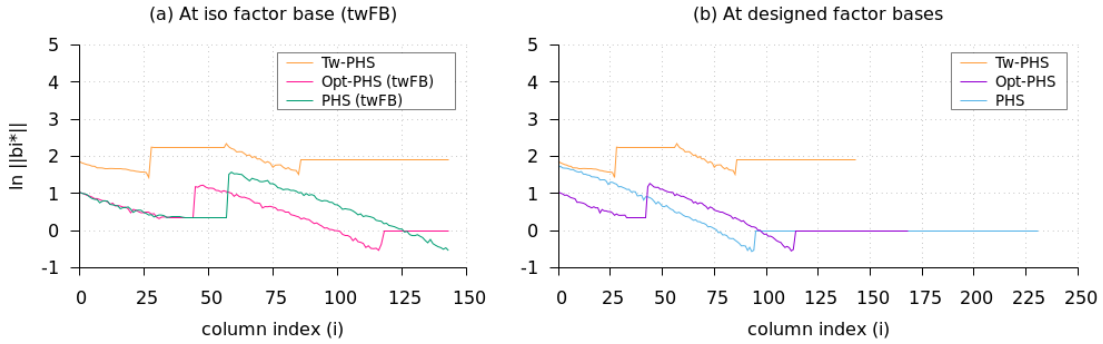

FIGURE 3.3 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{59})$: Gram-Schmidt log norms after BKZ_{40} reduction: (a) at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$; (b) at designed factor bases.

Since the volume of L_{tw} is bigger, by roughly the average log norm of the factor base elements by Lem. 3.20, the Gram-Schmidt log norms of our bases have bigger values. The important phenomenon to consider is how these log norms decrease. Figure 3.3 emphasises that the decrease of the Gram-Schmidt log norms is very limited in the twisted case, compared to other cases (with iso factor bases on the left, and the original algorithms on the right), where the decrease of the log norms seems significant. This observation seems to corroborate the fact that the Twisted-PHS lattice is already quite orthogonal. Finally, we note that both phenomenons do not depend on the lattices having the same dimension.

3.4.3 Approximation factors

We implemented all three algorithms from end to end and used them on numerous challenges to estimate their practically achieved approximation factors. This is to our knowledge the first time that these types of algorithms are completely run on concrete examples.

Ideal SVP challenges and ClDLP computations.

For each targeted field, we chose 50 prime ideals $\mathfrak{b}$ of prime norm q. Indeed, these are the most interesting ideals: in the extreme opposite case, taking $\mathfrak{b}$ inert of norm q^n implies that q reaches the lower bound of Eq. (2.41), as $\|q\|_2 = \sqrt{n} \cdot q$, hence the id-SVP solution is trivial.

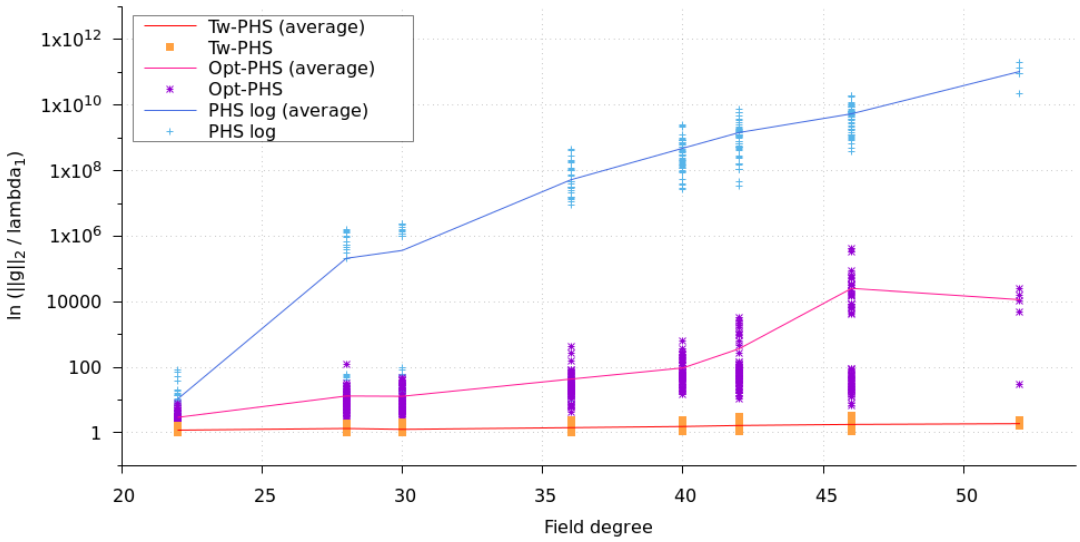

Figure 3.4 – Approximation factors reached by Twisted-PHS, Opt-PHS and PHS for cyclotomic fields of conductors 23, 29, 31, 37, 41, 43, 47 and 53 (in log scale).

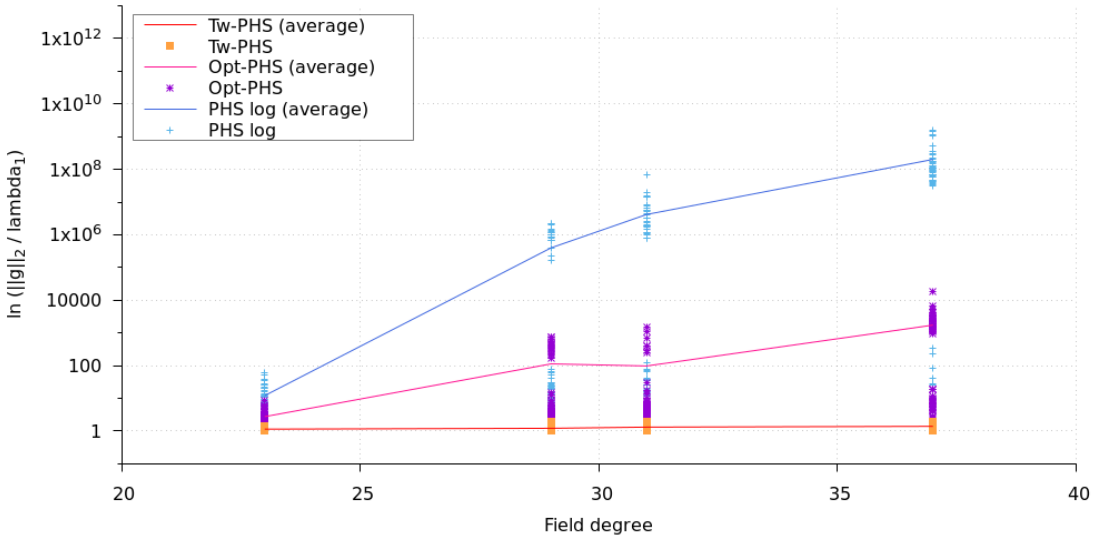

Figure 3.5 – Approximation factors reached by Twisted-PHS, Opt-PHS and PHS for NTRU Prime fields of degrees 23, 29, 31 and 37 (in log scale).

We then tried to solve the ClDlp for these challenges w.r.t. all targeted factor bases. We stress that, using Magma, $\mathcal{S}$-units computations for the ClDlp become harder as the norm of the challenge grows. This is especially true when the factor base inflates, hence providing an additional motivation for taking as small as possible factor bases. Therefore, we restricted ourselves to challenges of norms around 100 bits. Computing the ClDlp solutions for these challenges revealed much harder than computing $\mathcal{S}$-units on all factor bases, which contain only relatively small prime ideals. As a consequence, we were able to compute the ClDlp step only up to $\mathbb{Q}(\zeta_{53})$ (partially) and $\mathbb{Q}(z_{47})$.

Query algorithm.

We exclusively used Babai's Nearest Plane algorithm on the BKZ reduced bases of all log-$\mathcal{S}$-unit lattices to solve the Approx-CVP instances. Actually, the hardest computational task was to compute the output α/s, which necessitates a multi-exponentiation over huge $\mathcal{S}$-units.

As a particular point of interest, we stress that using directly the drift proposed in [PHS19a] would be especially unfair. Hence, for a challenge $\mathfrak{b}$, the target drifts $\mathbf{b}_{\mathsf{phs}}$, $\widetilde{\mathbf{b}}_{\mathsf{phs}}$ and $\mathbf{b}_{\mathsf{tw}}$ were all minimized using an iterative dichotomic approach on β and $\widetilde{\beta}$, taking a bigger value if the output $x \notin \mathfrak{b}$, and a smaller value if $x \in \mathfrak{b}$. After 5 iterations, the shortest x that verified $x \in \mathfrak{b}$ is returned.

Exact approximation factors.

Figures 3.4 and 3.5 report the obtained approximation factors. Note that for these dimensions, it is still possible to *exactly* solve id-SVP in the Minkowski space, so that these graphs show *real* approximation factors. We stress that we used a logarithmic scale to represent on the same graphs the performance of the Twisted-, Opt-PHS and PHS algorithms. The figures suggest that the approximation factor reached by our algorithm increases very slowly with the dimension, in a way that could reveal subexponential or even better. This feature would be particularly interesting to prove.

As a final remark, we point out that increasing the factor base for our Twisted-PHS algorithm has very little impact on the quality of the output. This is expected, since the log norm of the prime ideals constrain the valuation of the output, as in the proof of Pr. 3.24. On the contrary, increasing the factor base for the PHS and Opt-PHS variants clearly sabotages the quality of their output, as their lattice description is blind to these prime norms.

3.5 Supplementary Experimental Data

This section provides extensive additional data for all targeted fields.

3.5.1 Geometric characteristics

First, Tab. 3.5 and 3.6 extend the example given in Tab. 3.4 to respectively all targeted cyclotomic and NTRU Prime fields.

		d	$V^{1/d}$	δ_0		δ		θ_{min}		θ_{avg}		μ_2	μ_∞	$\lVert\cdot\rVert_\infty/\lVert\cdot\rVert_2$	
				raw	bkz	raw	bkz	raw	bkz	raw	bkz			real	H. 3.29
$\mathbb{Q}(\zeta_{23})$	L_{tw}	32	3.796	0.999	0.999	1.667	1.437	31	50	69	77	7.637	4.349	0.636	0.570
	$L^{(0)}_{\mathsf{opt}}$	32	1.515	1.030	1.009	2.477	1.615	40	60	76	81	3.120	1.706	0.676	0.612
	$L^{(0)}_{\mathsf{phs}}$	32	2.083	1.056	0.998	4.689	1.490	34	60	75	81	4.287	2.621	0.690	0.612
	L_{opt}	44	1.334	1.023	1.009	2.711	1.843	37	58	76	82	3.244	1.451	0.640	0.570
	L_{phs}	65	1.246	1.021	1.002	3.141	2.067	21	58	76	82	3.703	1.588	0.640	0.517
$\mathbb{Q}(\zeta_{29})$	L_{tw}	41	4.175	1.001	1.001	1.622	1.579	47	50	77	81	9.594	4.214	0.633	0.537
	$L^{(0)}_{\mathsf{opt}}$	41	1.616	1.025	1.005	2.742	1.870	40	41	78	82	3.772	1.925	0.660	0.580
	$L^{(0)}_{\mathsf{phs}}$	41	2.333	1.047	0.996	5.885	1.664	34	48	77	83	5.850	3.175	0.675	0.580
	L_{opt}	63	1.350	1.018	1.006	3.116	2.143	43	48	78	83	3.910	1.546	0.617	0.522
	L_{phs}	90	1.271	1.017	1.005	4.211	2.560	36	30	77	82	4.547	2.123	0.664	0.474

		d	$V^{1/d}$	δ_0		δ		θ_{min}		θ_{avg}		μ_2	μ_∞	$\|\cdot\|_\infty/\|\cdot\|_2$	
				raw	bkz	raw	bkz	raw	bkz	raw	bkz			real	H. 3.29
$\mathbb{Q}(\zeta_{31})$	L_{tw}	20	4.144	1.004	1.004	1.682	1.330	16	48	77	79	6.877	4.026	0.753	0.597
	$L^{(0)}_{\mathsf{opt}}$	20	10.36	1.051	0.930	3.029	1.269	41	58	76	82	26.93	18.10	0.808	0.669
	$L^{(0)}_{\mathsf{phs}}$	20	21.14	1.071	0.897	4.168	1.186	30	62	76	84	70.21	49.41	0.825	0.669
	L_{opt}	69	1.353	1.017	1.006	4.438	2.120	26	48	69	82	4.049	1.911	0.610	0.509
	L_{phs}	99	1.273	1.016	1.005	5.606	2.650	23	30	70	83	4.699	2.341	0.660	0.461
$\mathbb{Q}(\zeta_{37})$	L_{tw}	53	5.092	0.999	0.999	6.393	1.688	3	48	53	82	13.16	5.894	0.651	0.504
	$L^{(0)}_{\mathsf{opt}}$	53	1.694	1.020	1.007	6.969	1.977	8	55	61	82	4.481	2.079	0.635	0.545
	$L^{(0)}_{\mathsf{phs}}$	53	2.621	1.040	0.998	9.801	1.767	28	55	74	83	7.578	3.901	0.641	0.545
	L_{opt}	89	1.369	1.015	1.004	9.976	2.371	8	41	52	83	4.735	1.870	0.592	0.475
	L_{phs}	126	1.292	1.013	1.005	11.80	3.082	10	37	53	83	5.938	2.567	0.682	0.430
$\mathbb{Q}(\zeta_{41})$	L_{tw}	59	4.825	1.001	1.001	3.596	1.802	11	47	69	81	12.91	5.186	0.615	0.489
	$L^{(0)}_{\mathsf{opt}}$	59	1.786	1.020	1.005	4.525	1.986	34	55	76	83	5.112	2.245	0.629	0.530
	$L^{(0)}_{\mathsf{phs}}$	59	2.767	1.037	0.997	8.986	1.809	45	55	79	84	8.535	4.039	0.639	0.530
	L_{opt}	103	1.379	1.013	1.006	6.514	2.592	25	48	66	84	5.301	2.052	0.596	0.456
	L_{phs}	144	1.306	1.012	1.004	7.982	3.651	29	49	71	83	6.536	2.772	0.687	0.414
$\mathbb{Q}(\zeta_{43})$	L_{tw}	62	5.413	1.000	1.000	19.05	1.800	0	48	50	82	15.12	6.541	0.647	0.483
	$L^{(0)}_{\mathsf{opt}}$	62	1.773	1.018	1.005	19.51	2.019	2	60	53	83	5.165	2.246	0.622	0.524
	$L^{(0)}_{\mathsf{phs}}$	62	2.826	1.035	0.997	21.51	1.806	7	55	62	84	9.056	4.253	0.641	0.524
	L_{opt}	111	1.377	1.012	1.005	38.17	2.678	2	48	36	84	5.320	2.358	0.594	0.447
	L_{phs}	154	1.307	1.012	1.007	48.72	3.997	2	56	32	82	6.968	2.796	0.709	0.405
$\mathbb{Q}(\zeta_{47})$	L_{tw}	68	5.896	0.999	0.999	38.31	1.736	0	47	50	83	17.09	7.888	0.664	0.471
	$L^{(0)}_{\mathsf{opt}}$	68	1.819	1.017	1.007	39.31	2.171	1	60	52	83	5.525	2.597	0.618	0.511
	$L^{(0)}_{\mathsf{phs}}$	68	2.952	1.033	0.999	41.95	1.940	3	60	57	84	10.09	4.343	0.645	0.511
	L_{opt}	125	1.385	1.011	1.005	89.69	2.961	0	32	34	84	5.817	2.006	0.614	0.431
	L_{phs}	173	1.316	1.011	1.004	137.8	4.360	1	55	26	83	7.570	2.862	0.713	0.391
$\mathbb{Q}(\zeta_{53})$	L_{tw}	77	5.385	1.002	1.002	149.6	1.891	0	47	49	83	15.71	6.309	0.617	0.455
	$L^{(0)}_{\mathsf{opt}}$	77	1.928	1.016	1.005	152.0	2.315	0	60	49	83	6.381	2.382	0.611	0.495
	$L^{(0)}_{\mathsf{phs}}$	77	3.145	1.030	0.998	154.6	2.053	0	47	51	84	11.88	5.677	0.635	0.495
	L_{opt}	147	1.397	1.010	1.005	526.7	3.265	0	43	28	84	6.613	2.024	0.609	0.411
	L_{phs}	202	1.330	1.010	1.006	763.2	5.214	0	56	22	83	8.930	3.166	0.708	0.373
$\mathbb{Q}(\zeta_{59})$	L_{tw}	144	6.871	0.999	0.999	813.3	2.045	0	46	33	85	28.12	7.960	0.570	0.391
	$L^{(0)}_{\mathsf{opt}}$	144	1.490	1.010	1.004	821.0	3.301	0	37	34	84	7.091	2.378	0.602	0.414
	$L^{(0)}_{\mathsf{phs}}$	144	1.785	1.016	1.003	831.9	3.064	0	48	34	85	9.122	2.849	0.575	0.414
	L_{opt}	169	1.404	1.009	1.004	1181.	3.637	0	41	28	84	7.213	2.427	0.620	0.394
	L_{phs}	232	1.338	1.009	1.006	1753.	6.011	0	55	21	83	10.30	3.598	0.723	0.357
$\mathbb{Q}(\zeta_{61})$	L_{tw}	89	6.550	1.000	1.000	868.4	1.763	0	46	49	84	21.69	7.293	0.614	0.437
	$L^{(0)}_{\mathsf{opt}}$	89	1.971	1.014	1.005	881.5	2.425	0	57	49	84	7.081	2.807	0.610	0.475
	$L^{(0)}_{\mathsf{phs}}$	89	3.365	1.027	0.999	895.4	2.192	0	57	50	84	14.36	6.576	0.637	0.475
	L_{opt}	177	1.407	1.009	1.004	4765.	3.766	0	36	28	84	7.639	2.451	0.632	0.389
	L_{phs}	242	1.342	1.008	1.006	7948.	6.408	0	49	21	83	10.74	4.002	0.732	0.352

Table 3.5 – Geometric characteristics of log-$\mathcal{S}$-unit lattices for all non-principal cyclotomic fields $\mathbb{Q}(\zeta_m)$ of prime conductor $m \leq 61$.

		d	$V^{1/d}$	δ_0		δ		θ_{min}		θ_{avg}		μ_2	μ_∞	$\lVert\cdot\rVert_\infty/\lVert\cdot\rVert_2$	
				raw	bkz	raw	bkz	raw	bkz	raw	bkz			real	H. 3.29
$\mathbb{Q}(z_{23})$	L_{tw}	15	2.728	0.831	0.831	1.314	1.161	51	58	81	82	4.686	3.245	0.701	0.634
	$L^{(0)}_{\mathsf{opt}}$	15	14.70	0.879	0.879	2.923	1.241	45	49	80	80	25.14	18.09	0.766	0.699
	$L^{(0)}_{\mathsf{phs}}$	15	19.18	0.888	0.888	3.470	1.252	47	66	78	81	32.44	22.05	0.773	0.699
	L_{opt}	48	1.298	0.958	0.958	2.949	1.966	32	57	76	81	3.480	1.702	0.657	0.558
	L_{phs}	72	1.220	0.976	0.976	3.780	2.229	28	54	75	82	3.991	1.709	0.656	0.504
$\mathbb{Q}(z_{29})$	L_{tw}	21	3.333	0.863	0.863	1.357	1.219	55	60	81	83	6.592	3.923	0.679	0.597
	$L^{(0)}_{\mathsf{opt}}$	21	9.010	0.904	0.904	3.315	1.377	47	58	79	80	16.87	9.879	0.730	0.664
	$L^{(0)}_{\mathsf{phs}}$	21	14.68	0.917	0.917	4.538	1.428	47	55	79	80	27.83	18.60	0.735	0.664
	L_{opt}	66	1.329	0.967	0.967	3.733	2.336	38	56	77	82	4.184	1.972	0.665	0.515
	L_{phs}	97	1.252	0.981	0.981	5.385	2.745	27	55	76	82	4.794	2.098	0.678	0.464
$\mathbb{Q}(z_{31})$	L_{tw}	21	3.487	0.860	0.860	1.339	1.193	55	59	82	82	6.811	4.599	0.675	0.593
	$L^{(0)}_{\mathsf{opt}}$	21	13.52	0.898	0.898	3.113	1.456	43	57	81	79	26.65	16.66	0.728	0.664
	$L^{(0)}_{\mathsf{phs}}$	21	22.25	0.909	0.909	4.039	1.463	43	57	80	80	42.47	25.75	0.738	0.664
	L_{opt}	73	1.333	0.970	0.970	3.906	2.423	34	57	78	82	4.526	2.064	0.656	0.502
	L_{phs}	106	1.258	0.982	0.982	5.677	2.920	30	53	77	82	5.179	2.195	0.681	0.452
$\mathbb{Q}(z_{37})$	L_{tw}	30	4.069	0.893	0.893	1.405	1.313	58	60	82	83	9.299	5.202	0.653	0.556
	$L^{(0)}_{\mathsf{opt}}$	30	6.056	0.925	0.925	3.773	1.703	42	56	81	80	14.23	7.068	0.687	0.621
	$L^{(0)}_{\mathsf{phs}}$	30	11.46	0.940	0.940	6.147	1.692	41	53	80	80	24.99	15.60	0.688	0.621
	L_{opt}	93	1.348	0.975	0.975	4.627	2.897	38	55	78	82	5.337	2.247	0.662	0.470
	L_{phs}	133	1.277	0.985	0.985	7.306	3.506	25	52	78	82	6.147	2.677	0.695	0.424
$\mathbb{Q}(z_{41})$	L_{tw}	32	4.406	0.896	0.896	1.474	1.268	53	62	82	84	9.595	5.355	0.645	0.545
	$L^{(0)}_{\mathsf{opt}}$	32	7.279	0.925	0.925	3.895	1.740	41	58	81	81	17.06	8.830	0.690	0.612
	$L^{(0)}_{\mathsf{phs}}$	32	14.89	0.939	0.939	6.352	1.709	41	54	81	80	33.76	19.22	0.684	0.612
	L_{opt}	108	1.355	0.978	0.978	5.385	3.073	35	58	79	83	5.740	2.448	0.686	0.450
	L_{phs}	152	1.288	0.987	0.987	8.442	3.834	29	55	79	83	6.751	2.768	0.712	0.407
$\mathbb{Q}(z_{43})$	L_{tw}	38	4.441	0.911	0.911	1.498	1.357	53	59	82	83	10.64	5.177	0.645	0.528
	$L^{(0)}_{\mathsf{opt}}$	38	5.051	0.937	0.937	4.187	1.865	44	50	81	81	12.50	6.573	0.663	0.590
	$L^{(0)}_{\mathsf{phs}}$	38	9.657	0.952	0.952	7.496	1.877	45	56	81	81	23.73	12.18	0.671	0.590
	L_{opt}	114	1.367	0.979	0.979	5.482	3.256	36	57	79	83	6.119	2.803	0.687	0.443
	L_{phs}	161	1.297	0.987	0.987	9.002	4.135	25	55	79	83	7.484	2.837	0.712	0.400
$\mathbb{Q}(z_{47})$	L_{tw}	40	4.576	0.913	0.913	1.650	1.358	49	60	82	84	11.04	5.607	0.632	0.519
	$L^{(0)}_{\mathsf{opt}}$	40	6.231	0.938	0.938	4.628	1.915	37	57	81	81	16.59	8.398	0.658	0.583
	$L^{(0)}_{\mathsf{phs}}$	40	12.06	0.951	0.951	7.908	1.946	38	55	81	81	30.85	15.50	0.662	0.583
	L_{opt}	129	1.376	0.981	0.981	6.189	3.632	21	56	80	83	6.575	2.925	0.696	0.427
	L_{phs}	180	1.309	0.989	0.989	10.15	4.527	31	53	80	83	8.022	2.882	0.704	0.387

TABLE 3.6 – Geometric characteristics of log-$\mathcal{S}$-unit lattices for all targeted NTRU Prime fields $\mathbb{Q}(z_q)$, for prime $q \in [\![23, 47]\!]$.

These extensive data confirm the discussion made in §3.4.1. Note that these observations are especially valid for NTRU Prime fields. An explanation of this phenomenon might lie in the fact that for NTRU Prime fields, the norms of the factor base prime ideals are almost all distinct, so that the twisted characteristic of our lattices is fully used.

3.5.2 Gram-Schmidt norms of the lattice bases

We also provide the graphs showing the log norms of the Gram-Schmidt vectors for each field and each log-$\mathcal{S}$-unit lattice variant. These graphs confirm the discussion in §3.4.1, namely that in the twisted case, the BKZ reduction has very little impact and that the sequence of norms does not vary much. This corroborates the claim that our twisted lattices are much more orthogonal than expected.

In the case of PHS and its variants, there is always a significant gap between the Gram-Schmidt norms before and after the small block BKZ reduction, and the decrease of the log norms is very pronounced and going down to 0.

Prime conductor cyclotomic fields.

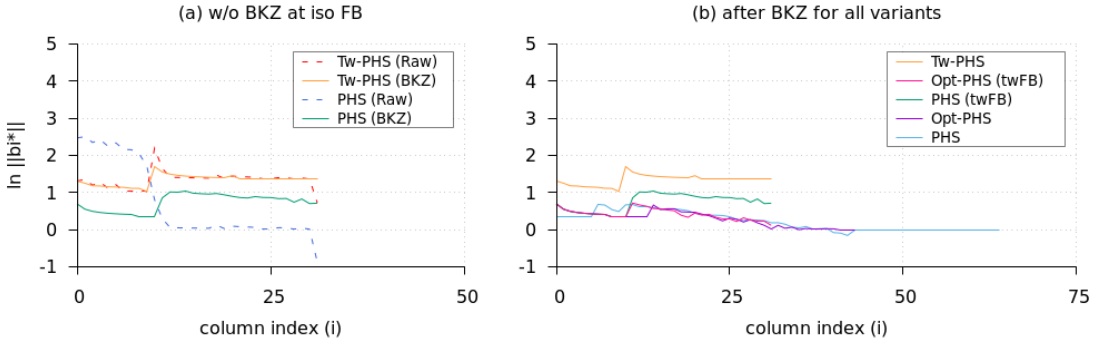

FIGURE 3.6 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{23})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\mathsf{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\mathsf{opt}}$, $L^{(0)}_{\mathsf{phs}}$, L_{opt} and L_{phs}.

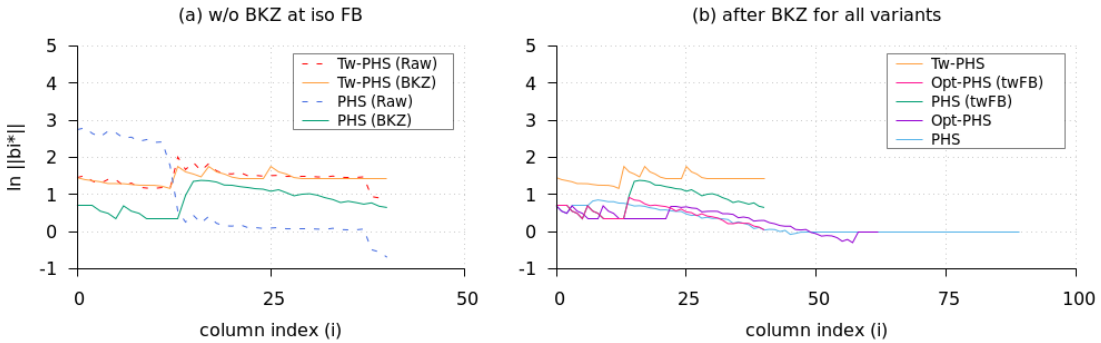

FIGURE 3.7 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{29})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\mathsf{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\mathsf{opt}}$, $L^{(0)}_{\mathsf{phs}}$, L_{opt} and L_{phs}.

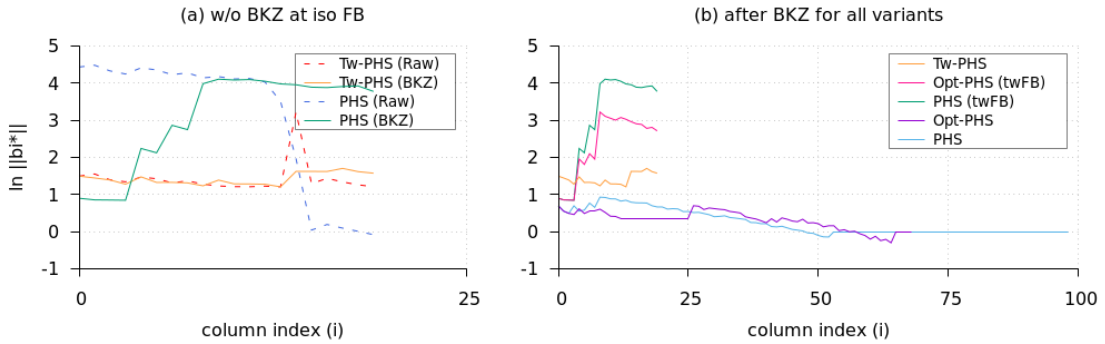

FIGURE 3.8 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{31})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

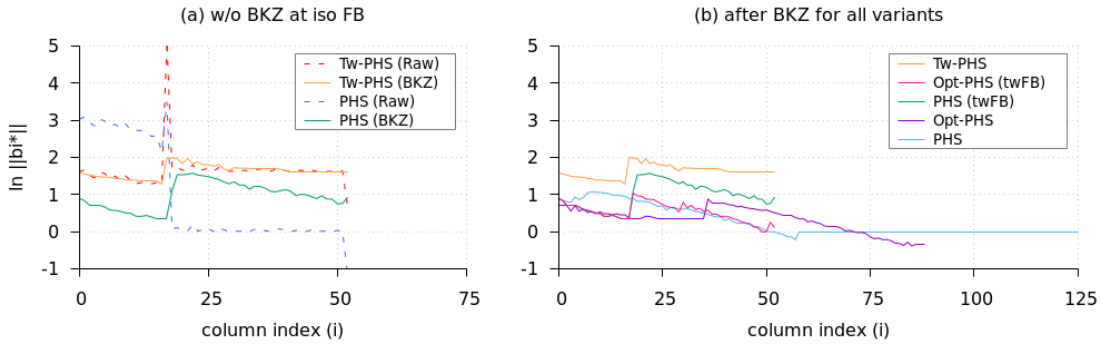

FIGURE 3.9 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{37})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

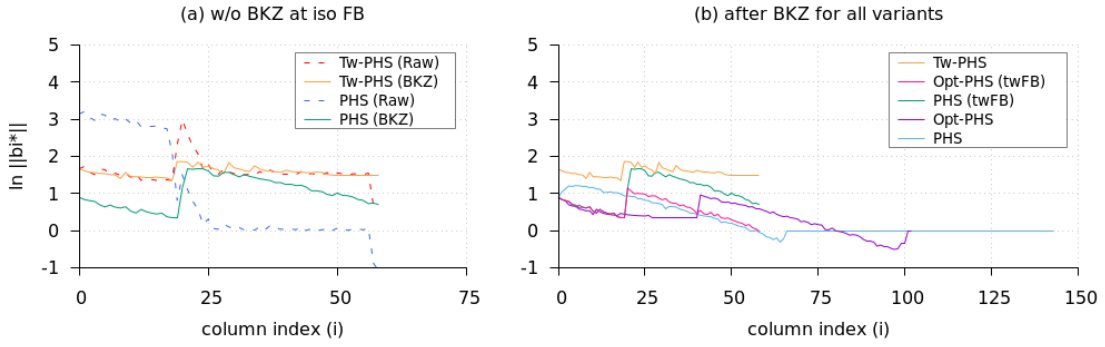

FIGURE 3.10 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{41})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

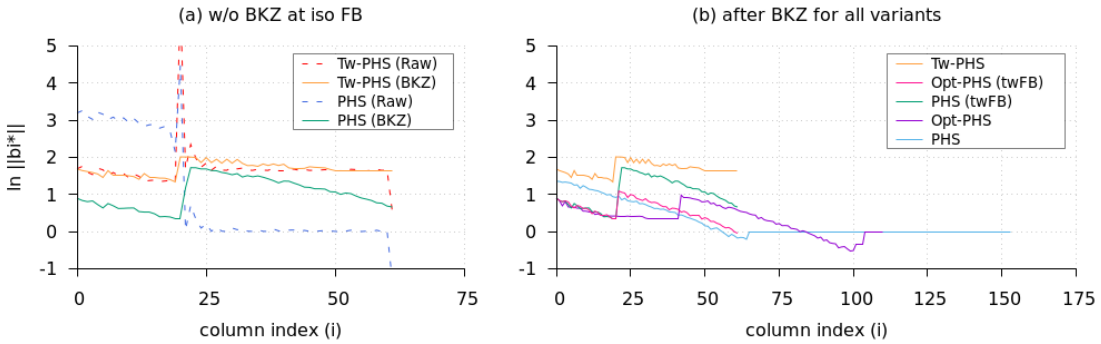

FIGURE 3.11 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{43})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

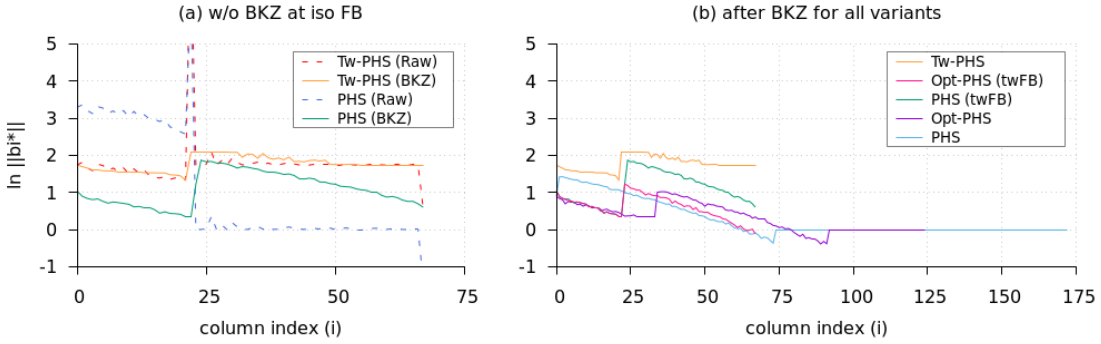

FIGURE 3.12 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{47})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

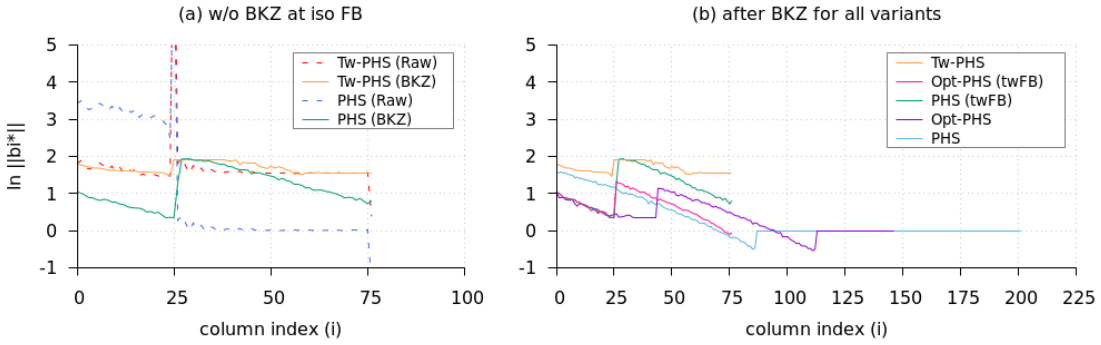

FIGURE 3.13 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{53})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

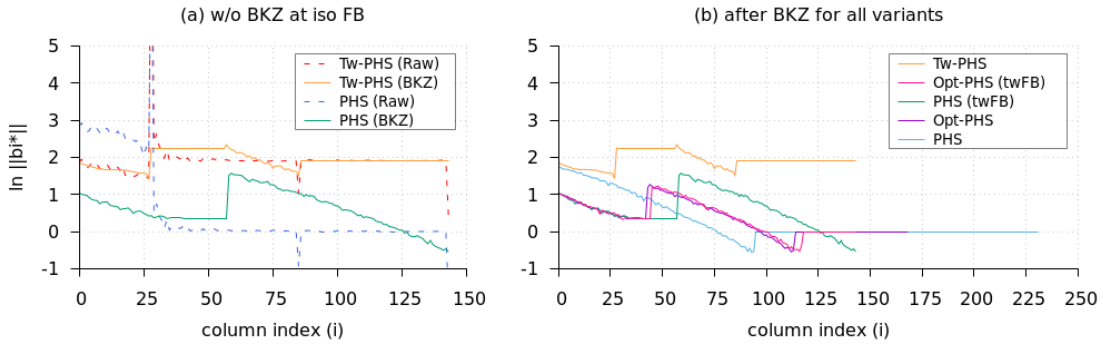

FIGURE 3.14 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{59})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

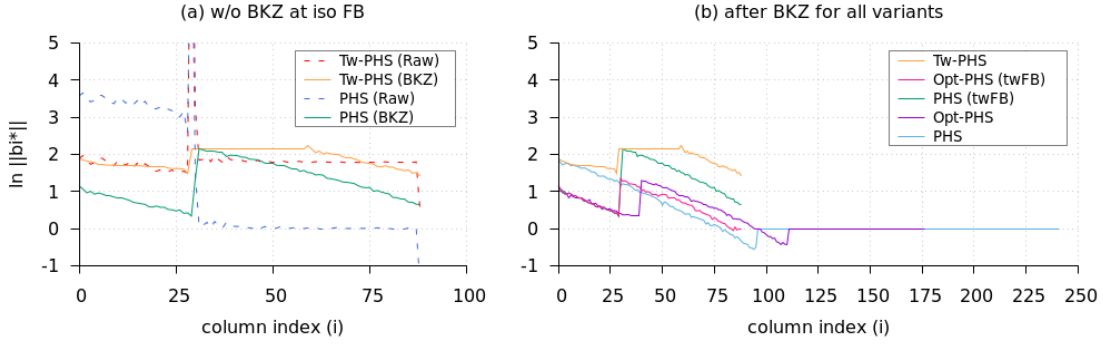

FIGURE 3.15 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(\zeta_{61})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

NTRU Prime fields.

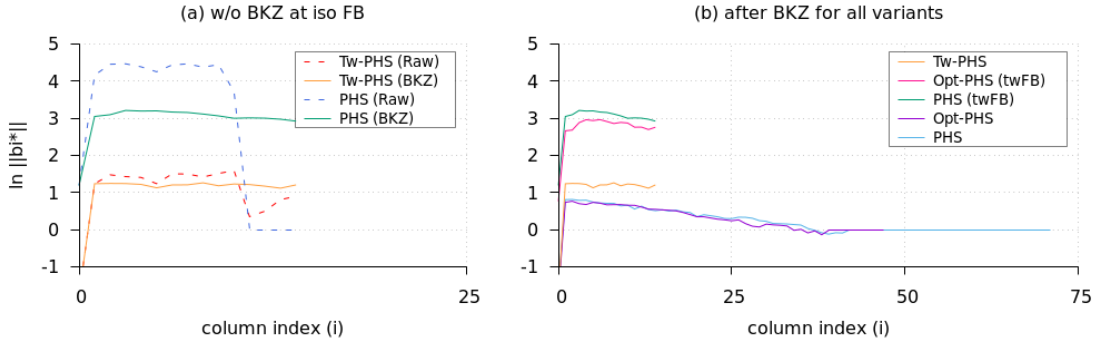

FIGURE 3.16 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(z_{23})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

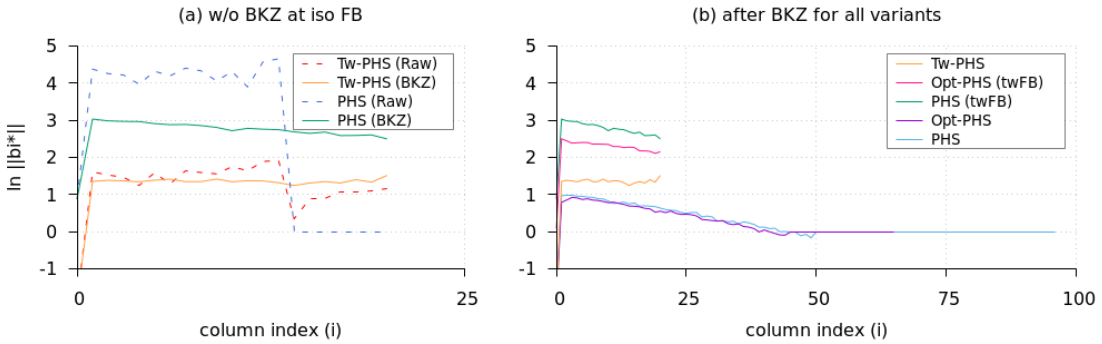

FIGURE 3.17 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(z_{29})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

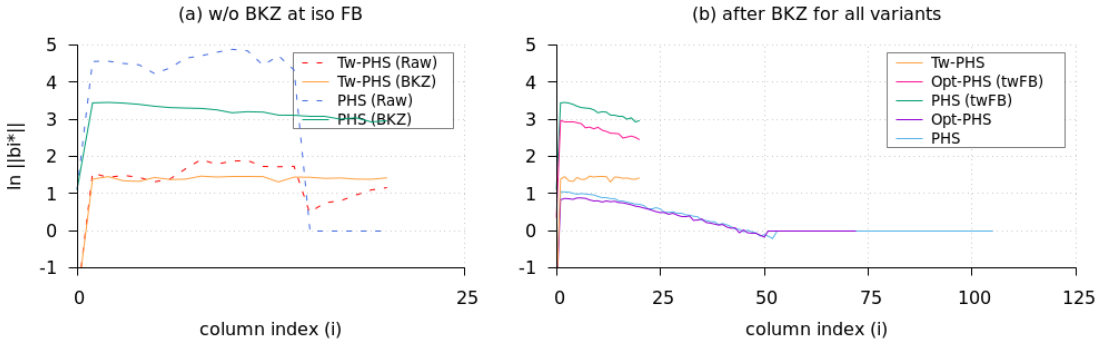

FIGURE 3.18 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(z_{31})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

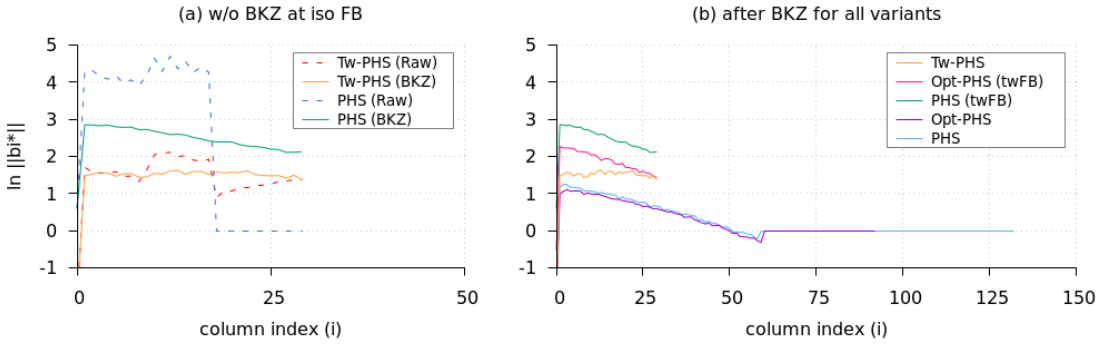

FIGURE 3.19 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(z_{37})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

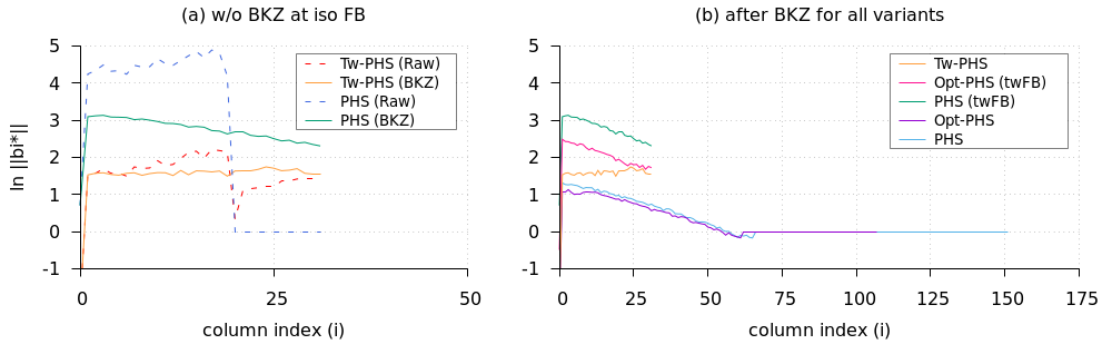

FIGURE 3.20 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(z_{41})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

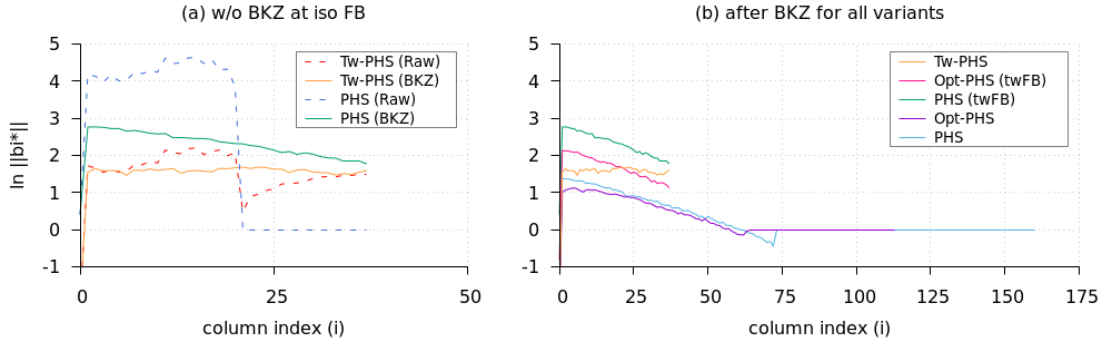

FIGURE 3.21 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(z_{43})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

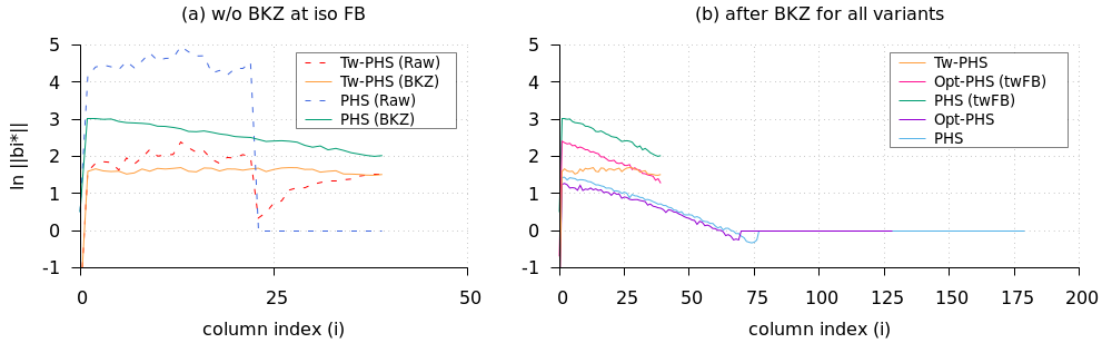

FIGURE 3.22 – Log-$\mathcal{S}$-unit lattices for $\mathbb{Q}(z_{47})$, Gram-Schmidt log norms: (a) before and after BKZ_{40} reduction at iso factor base $\mathcal{A}_{\text{tw-FB}}(K)$ for L_{tw} and $L^{(0)}_{\text{phs}}$; (b) after BKZ_{40} reduction for all variants L_{tw}, $L^{(0)}_{\text{opt}}$, $L^{(0)}_{\text{phs}}$, L_{opt} and L_{phs}.

Chapter 4

A Short Basis of the Stickelberger Ideal of a Cyclotomic Field

In this chapter, the material come from a fruitful collaboration with Pr. Radan Kučera, and constitute the second contribution of this thesis.

> [BK21] **A short basis of the Stickelberger ideal of a cyclotomic field.**
> Olivier Bernard and Radan Kučera.

Submitted to AMS :: Mathematics of Computation (*American Mathematical Society*).

2010 MSC classes: 11R18 (Primary), 11R29, 11Y40 (Secondary).
Keywords: Cyclotomic fields, Stickelberger ideal, short basis, relative class number.
Links: [arXiv: 2109.13329 [math.NT][10]]

Contents

[10] https://arxiv.org/abs/2109.13329

4.1 Introduction

A popular choice for lattice-based cryptography is to consider fractional ideals in some cyclotomic field $K_m = \mathbb{Q}[\zeta_m]$ of conductor $m \not\equiv 2 \pmod 4$, e.g., $m = 2048$. In the last decade, there has been a significant cryptanalytic effort trying to benefit from this additional algebraic structure to solve Approx-Ideal-SVP, giving rise to a long series of works [CGS14, CDPR16, CDW17, DPW19, PHS19a, BR20, CDW21]. All approaches start from a solution to the Class Group Discrete Logarithm Problem (ClDLP), which is, given a fixed set of finite places corresponding to prime ideals $\{\mathfrak{p}_1, \dots, \mathfrak{p}_k\}$ of K_m, and any challenge ideal $\mathfrak{b}$ whose class in the class group of K_m belongs to the subgroup generated by the classes of the $\mathfrak{p}_i$'s, to find $\alpha \in K_m$ and $e_1, \dots, e_k \in \mathbb{Z}$ such that:

$$\langle \alpha \rangle = \mathfrak{b} \cdot \prod_{1 \le i \le k} \mathfrak{p}_i^{e_i}.$$

In a quantum world, this problem is not hard to solve [EHKS14, BS16], so the most difficult part of these cryptanalyses resides in reducing the Euclidean norm of α.

In the case of cyclotomic fields K_m of conductor m, the Stickelberger ideal $\mathcal{S}_m$ of K_m annihilates its class group, by Stickelberger's theorem. Thus, it was proposed in [CDW17, CDW21] to use these free relations to help to reduce the algebraic norm of the ClDLP solution. More precisely, since by [Sin78] $(1-\tau)\mathcal{S}_m$, viewed as a $\mathbb{Z}$-module, has full rank in $(1-\tau)\mathbb{Z}[G_m]$, where $G_m = \mathrm{Gal}(K_m/\mathbb{Q})$ and $\tau \in G_m$ is induced by complex conjugation, it is a lattice of class relations for the relative class group. Therefore, choosing a challenge ideal $\mathfrak{b}$ and prime ideals for the ClDLP in the relative class group, e.g., exactly one Galois orbit $\{\mathfrak{p}^\sigma\}$ for all $\sigma \in G_m$, it is possible to express the reduction of a solution $\langle \alpha \rangle = \mathfrak{b} \cdot \mathfrak{p}^{\sum e_\sigma \sigma}$ as a closest vector problem in $(1-\tau)\mathcal{S}_m$, where the target is the vector $(e_\sigma - e_{\tau\sigma})_\sigma$.

As noticed in [CDW21, Lem. 4.4 and 4.6], this lattice contains many short elements, i.e., elements of $\mathbb{Z}[G_m]$ of the form $\sum a_\sigma \sigma$, where all $a_\sigma \in \{0, 1\}$. *In fine*, this yields a good description for finding sufficiently close vectors. Also, the plus part of the class group seems to be much smaller than the relative part,[18] hence every challenge $\mathfrak{b}$ can be reduced to this case by randomly searching for a small norm ideal $\mathfrak{c}$ such that the class of $\mathfrak{cb}$ belongs to the relative class group [CDW21, Alg. 5].

4.1.1 In praise of short Stickelberger bases

Unfortunately, while in the prime conductor case the exhibited set of short elements from [CDW21, §4.2] form a $\mathbb{Z}$-basis of $\mathcal{S}_m$, in the general case this family is only known to generate $\mathcal{S}_m$ as a $\mathbb{Z}$-module. This comes at the expense of constructing a linearly independent subset of vectors [CDW21, 2.2]. Whereas this is certainly possible without any geometric loss, using e.g., [MG02, Lem. 7.1], it induces a slight growth of the Euclidean norm of the obtained basis vectors. For some applications, this can have dramatic consequences and it is not clear whether it is always possible to find a basis among all subsets of such a short generating set.

A very important point is that the proof of Stickelberger's theorem, i.e., that the Stickelberger ideal annihilates the class group, is completely explicit [Was97, §6.2]. Namely, for any prime ideal $\mathfrak{p}$, and any $\alpha \in \mathcal{S}_m$, it builds an explicit $\gamma \in K_m$ such that $\langle \gamma \rangle = \mathfrak{p}^\alpha$. However, if α has even moderately large coefficients, this has an exponential impact on the height of the coefficients of γ, that renders its computation rapidly intractable. On the contrary, having only short elements

[18] This is backed up by several theoretical and computational observations, see e.g., Weber's conjecture $h_{2^e}^+ = 1$, Buhler, Pomerance and Robertson's conjecture for odd prime powers [BPR04], and Schoof's extensive calculations in [Was97, Tab., §4] and [Sch03].

in the basis keeps the algebraic norm of the generators as low as possible, namely $\mathcal{N}(\mathfrak{p})^{\varphi(m)/2}$. Explicitly computing Stickelberger generators is useful in at least two situations:

1. the first one is when reducing the algebraic norm of the ClDlp solution as in [CDW21], as knowing explicit generators prevents to perform a quantum step – or, a classically costly step – to recover the generator of the reduced ideal (see [CDW21, Th. 5.1] for the complete workflow);
2. the second one occurs when one wants to use the knowledge of the Stickelberger relations to approach some log-$\mathcal{S}$-unit lattice. Indeed, suppose the finite places of $\mathcal{S}$ correspond to one split Galois orbit $\{\mathfrak{p}^\sigma\}$ for all $\sigma \in G_m$. Then, from a maximal set of independent real $\mathcal{S}^+$-units, where the finite places of $\mathcal{S}^+$ correspond to all relative norm ideals $\mathcal{N}_{K_m/K_m^+}(\mathfrak{p}^\sigma)$, adding explicit generators corresponding to a basis of the Stickelberger ideal, besides the absolute norm, yields a maximal set of independent $\mathcal{S}$-units, at the much smaller cost of finding generators in the maximal real subfield.

In the latter case, note that knowing merely a short generating set of $\mathcal{S}_m$ instead of a $\mathbb{Z}$-basis is not sufficient to provide a full-rank family of independent $\mathcal{S}$-units. Building a basis from such a generating set using e.g., the Hermite Normal Form, an LLL reduction or [MG02, Lem. 7.1], increases dramatically the size of the (possibly rational) coefficients of the respective generators. Not to mention the computational burden to manipulate such elements, this significantly hinders their potential use: for example, in the *saturation* process (see e.g., §5.2.4) that allows to approach further log-$\mathcal{S}$-unit lattices, it is vital to constrain both the number of elements and their size. Hence, having in the first place an *explicit short basis* of $\mathcal{S}_m$ as a $\mathbb{Z}$-module is particularly useful.

Historical results.

The first explicitly known basis of $\mathcal{S}_m$, viewed as a $\mathbb{Z}$-module, for *any* conductor m was given in [Kuč92, Th. 6.2], but elements of this basis have rather large coefficients. In the prime conductor case, a short basis can be found in [Sch08, Th. 9.3(i)], the shortness being proven in [Sch08, Ex. 9.3]. This result has been extended to prime-power conductors in [CDW17], at the price of allowing slightly larger coefficients [CDW17, Lem. 4(2)]. Finally, a large set of short *generators* has been given in [CDW21, §4.2] in the general case for any conductor.

4.1.2 Contributions

In this work, our main result (see Th. 4.29) is to provide the first explicit *basis* of the Stickelberger ideal $\mathcal{S}_m$ for any conductor m, viewed as a $\mathbb{Z}$-module, that is constituted *only* of short elements, i.e., elements of the form

$$\sum_{\sigma \in G_m} a_\sigma \sigma \in \mathcal{S}_m \subset \mathbb{Z}[G_m], \qquad \text{where } a_\sigma \in \{0,1\} \text{ for all } \sigma \in G_m.$$

Actually, besides the absolute norm element, all other members of this short basis have exactly $\varphi(m)/2$ non-zero coordinates. In the prime conductor case, our short basis coincides with the basis given in [Sch08, Th. 9.3(i)]. One ingredient of independent interest in the proof is Pr. 4.15, which describes a large family of short elements of $\mathcal{S}_m$ that encompasses the set from [CDW21, §4.2], using a very simple arithmetic criterion in the spirit of [Was97, Lem. 16.3] when m is an odd prime power. Picking wisely some elements $\alpha_m(b)$ in this large family yields our proposed short basis.

We also show how to explicitly compute algebraic integers generating $\mathfrak{L}^{\alpha_m(b)}$, for any unramified prime ideal $\mathfrak{L}$ and any element $\alpha_m(b)$ of our short basis. These generators can be expressed

as Jacobi sums that turn out to be drastically more efficient to compute than the generators given e.g., in [Was97, §6.2].

Finally, a nice theoretical consequence of our result is to derive an explicit upper bound on the relative part h_m^- of the class number of K_m. More precisely, for any conductor $m \not\equiv 2 \pmod 4$, Cor. 4.32 gives that

$$h_m^- \leq 2^{1-a} \cdot \left(\frac{\varphi(m)}{8}\right)^{\varphi(m)/4},$$

where

$$a = \begin{cases} 0 & \text{if } m \text{ is a prime-power,} \\ 2^{t-2} - 1 & \text{if } m \text{ has } t > 1 \text{ prime divisors.} \end{cases}$$

To our knowledge, the best explicit upper bound on the relative class number which is valid for any conductor is given by [Lou14, 6]. However, whereas our bound is given by a simple formula and easy to manipulate, Louboutin's bound is difficult to instantiate for comparison in the general case. As an example, the special case $m = 4p$, where $p \geq 3$ is an odd prime, is concretely treated in [Lou14, Th. 2], which results in the following upper bound:

$$h_{4p}^- \leq 8\sqrt{p} \cdot \left(\frac{p}{16}\right)^{(p-1)/2}.$$

We stress that in this example, this upper bound is sharper than ours.

We should also mention that the proof of our bound indirectly gives an algorithm to compute the relative class number by computing the determinant of some scaled Hadamard matrix: incidentally, this method seems to be significantly more efficient than when using the traditional analytic formula [Was97, Th. 4.17], when the number t of prime factors of m is small.

4.2 On Bases of $\mathcal{S}'_m$

Recall that $m > 1$ is a positive integer such that $m = q_1 q_2 \dots q_t \not\equiv 2 \pmod 4$, where $q_1, \dots, q_t$ are pairwise coprime prime powers greater than 2.

4.2.1 A first basis of $\mathcal{S}'_m$

We first give a basis of $\mathcal{S}'_m$ constructed in the spirit of [Kuč92, Th. 4.2]. Let X_m and M_m^- be the subsets of $[\![1, m]\!]$ defined in §2.2.1, and recall that M_m^- is exactly the set M_- defined in [Kuč92, p.293]. This set M_m^- has the following stability property:

Lemma 4.1. *Let $r \mid m$, $0 < r < m$, such that $\left(r, \frac{m}{r}\right) = 1$. Let the set $M_{\frac{m}{r}}^-$ be defined using the ordering of prime power divisors of $\frac{m}{r}$ induced by the chosen ordering of prime power divisors of m. Then*

$$\left\{a \in M_m^-;\ r \mid a\right\} = \left\{rb;\ b \in M_{\frac{m}{r}}^-\right\} = r \cdot M_{\frac{m}{r}}^-.$$

Proof. For any integer b, $0 < b < \frac{m}{r}$, we have $b \in X_{\frac{m}{r}}$ if and only if for each $i \in [\![1, t]\!]$ such that $q_i \mid \frac{m}{r}$, either $(q_i, b) = 1$ or $q_i \mid b$. This is the case if and only if for each $i \in [\![1, t]\!]$, either $(q_i, rb) = 1$ or $q_i \mid rb$, thus if and only if $rb \in X_m$.

If $q_i \mid \frac{m}{r}$ for some $i \in [\![1, t]\!]$ then $(q_i, r) = 1$, and so $q_i \nmid rb$ if and only if $q_i \nmid b$, moreover $b \not\equiv -(b, \frac{m}{r}) \pmod{q_i}$ if and only if $br \not\equiv -(br, m) \pmod{q_i}$.

If $b \nmid \frac{m}{r}$ then for any $i \in [\![1, t]\!]$ such that $q_i \mid \frac{m}{r}$ we have $b \not\equiv (b, \frac{m}{r}) \pmod{q_i}$ if and only if $br \not\equiv (br, m) \pmod{q_i}$. Therefore we get the same k for $b \in X_{\frac{m}{r}}$ and for $br \in X_m$. Moreover $\frac{b}{(b, \frac{m}{r})q_k} = \frac{br}{(br, m)q_k}$.

If $b \mid \frac{m}{r}$ then $\left\{i \in [\![1, t]\!];\ q_i \mid \frac{m}{r}, q_i \nmid b\right\} = \left\{i \in [\![1, t]\!];\ q_i \nmid rb\right\}$. □

Theorem 4.2. *For any integer $m > 1$, $m \not\equiv 2 \pmod 4$, the set*

$$\left\{\omega_m(a);\ a \in M_m^-\right\} \cup \left\{\tfrac{1}{2}N_m\right\} \tag{4.3}$$

is a $\mathbb{Z}$-basis of $\mathcal{S}'_m$.

Proof. This can be proved similarly to the part of [Kuč92, Th. 4.2] about the Stickelberger ideal, using Lem. 2.19, 2.31 and 2.32 instead of [Kuč92, Lem. 3.1, 3.2 and 3.4]. Indeed, the proof of [Kuč92, Th. 4.2] about the Stickelberger ideal and its preparatory statements [Kuč92, Lem. 3.3 and 4.1 (for Ψ)] need the validity of only the following facts (using notations $\omega(a)$ and ω^* from [Kuč92]):

- the Stickelberger ideal is generated by $\left\{\omega(a);\ 0 < a < m\right\} \cup \left\{\omega^*\right\}$ as a group ([Kuč92, Lem. 3.1]);
- these generators satisfy the relations of Lem. 2.31, where we write $\omega(a)$ instead of $\omega_m(a)$ ([Kuč92, Lem. 3.2]);
- these generators satisfy the relations of Lem. 2.32, where we write $\omega(a)$ instead of $\omega_m(a)$ ([Kuč92, Lem. 3.4]).

Therefore, this proof can be used *mutatis mutandis* to get a basis for any group generated by generators satisfying these relations. Hence, plugging $\omega^* = \frac{1}{2}N_m$ and $\omega(a) = \omega_m(a)$, we deduce the theorem from Lem. 2.19, 2.31 and 2.32. □

The above basis inherits the stability property given in Lem. 4.1.

Proposition 4.4. *For any given $b \in \mathbb{Z}$, $0 < b < m$, let r_b be the maximal divisor of (b, m) satisfying $\left(r_b, \frac{m}{r_b}\right) = 1$, i.e., r_b is the product of all $q_i, i \in [\![1, t]\!]$ which divide b, and write $\omega_m(b) \in \mathcal{S}'_m$ as a unique $\mathbb{Z}$-linear combination of basis elements (4.3). Then for each $a \in M_m^-$ such that $r_b \nmid a$, the coefficient of $\omega_m(a)$ in this $\mathbb{Z}$-linear combination is equal to zero.*

Proof. For brevity's sake, let $r = r_b$. By Eq. (2.29), $\omega_m(b) = \mathrm{Cor}_{K_m/K_{\frac{m}{r}}}\left(\omega_{\frac{m}{r}}(\frac{b}{r})\right)$. Using Th. 4.2 for $\frac{m}{r}$ implies $\omega_{\frac{m}{r}}(\frac{b}{r}) \in \mathcal{S}'_{\frac{m}{r}}$ is a unique $\mathbb{Z}$-linear combination of

$$\left\{\omega_{\frac{m}{r}}(a);\ a \in M_{\frac{m}{r}}^-\right\} \cup \left\{\tfrac{1}{2}N_{\frac{m}{r}}\right\}.$$

Since by Eq. (2.29), $\mathrm{Cor}_{K_m/K_{\frac{m}{r}}}\left(\omega_{\frac{m}{r}}(a)\right) = \omega_m(ra)$ and $\mathrm{Cor}_{K_m/K_{\frac{m}{r}}}\left(N_{\frac{m}{r}}\right) = N_m$, and since the corestriction map $\mathrm{Cor}_{K_m/K_{\frac{m}{r}}}$ is a linear map, the proposition follows from Lem. 4.1. □

In particular, for any positive $r \mid m$, $1 < r < m$, such that $\left(r, \frac{m}{r}\right) = 1$, the corestriction subgroup $\mathrm{Cor}_{K_m/K_{m/r}}\left(\mathcal{S}'_{m/r}\right)$ of $\mathcal{S}'_m$ has the following $\mathbb{Z}$-basis

$$\left\{\omega_m(a);\ a \in r \cdot M_{\frac{m}{r}}^-\right\} \cup \left\{\tfrac{1}{2}N_m\right\}.$$

4.2.2 An alternative basis of $\mathcal{S}'_m$: the prime-power case

In this section we shall suppose that m is a prime power $q = p^e$, where p is a prime and e is a positive integer. Let us mention explicitly that the case $p = 2$ is allowed whenever $e \geq 2$ to ensure $q \not\equiv 2 \pmod 4$. We set

$$M'_q = M'_{p^e} = \left\{1, \ldots, \tfrac{\varphi(p^e)}{2}\right\}. \tag{4.5}$$

Theorem 4.6. *For any prime power $q = p^e > 2$, the set*

$$\left\{\omega_q(a);\ a \in M'_q\right\} \cup \left\{\tfrac{1}{2}N_q\right\} \tag{4.7}$$

is a $\mathbb{Z}$-basis of $\mathcal{S}'_q$.

Proof. We shall prove the theorem by induction with respect to e. If q is an odd prime or $q = 4$, we have $M'_q = M^-_q$ so this is just a special case of Th. 4.2.

Let us suppose that the theorem has been proved for $p^e > 2$ and let us prove it for $q = p^{e+1}$. Let H be the subgroup of $\mathcal{S}'_q$ generated by the set (4.7). We shall show that H contains all $\omega_q(a)$, $0 < a < q$, so that it generates $\mathcal{S}'_q$ by Eq. (2.30).

Since $\omega_q(q-a) = -\omega_q(a)$ by Eq. (2.28), the subgroup H contains also $\omega_q(a)$ for each $a \in \mathbb{Z}$ satisfying $q - \frac{\varphi(q)}{2} \leq a < q$. Suppose $a = bp$, using Eq. (2.29) we get

$$\omega_q(a) = \mathrm{Cor}_{K_q/K_{q/p}}\big(\omega_{q/p}(b)\big).$$

Since $\mathrm{Cor}_{K_q/K_{q/p}}$ is an injective linear map, the induction hypothesis implies $\omega_q(a)$ is a linear combination of $\frac{1}{2}N_q$ and of $\mathrm{Cor}_{K_q/K_{q/p}}\big(\omega_{q/p}(t)\big) = \omega_q(tp)$ for $t \in M'_{q/p}$, which implies $tp \in M'_q$. Thus, H contains $\omega_q(a)$ whenever $p \mid a$. As for the remaining cases, let $a \in \mathbb{Z}$ be such that $\frac{\varphi(q)}{2} < a < q - \frac{\varphi(q)}{2}$ and $p \nmid a$. Lemma 2.31 states that

$$\sum_{\substack{t=0,\dots,q-1\\ t\equiv a \pmod{q/p}}} \omega_q(t) = \omega_q(ap) \in H.$$

Since $\left(q - \frac{\varphi(q)}{2}\right) - \frac{\varphi(q)}{2} = \frac{q}{p}$, there is only one t in the sum on the left hand side satisfying $\frac{\varphi(q)}{2} < t < q - \frac{\varphi(q)}{2}$, namely $t = a$. All other summands are known to belong to H, and since we just proved that $\omega_q(ap) \in H$, we deduce $\omega_q(a) \in H$.

We have shown that H generates $\mathcal{S}'_q$. Since $\left|M'_q\right| = \left|M^-_q\right|$, the theorem follows. □

4.2.3 An alternative basis of $\mathcal{S}'_m$: the general case

Now, we return to the general case where $m = q_1 q_2 \dots q_t \not\equiv 2 \pmod 4$. Let us fix $i \in [\![1, t]\!]$. Lemma 4.1 gives that

$$\left\{a \in M^-_m;\ \tfrac{m}{q_i} \mid a\right\} = \tfrac{m}{q_i} \cdot M^-_{q_i} = \left\{\tfrac{mb}{q_i};\ p_i \nmid b, 0 < b < \tfrac{q_i}{2}\right\}.$$

Since $\mathrm{Cor}_{K_m/K_{q_i}}$ is an injective linear map, Pr. 4.4 and respectively Th. 4.6 combined with Eq. (2.29) imply that the sets

$$\left\{\omega_m(a);\ a \in \tfrac{m}{q_i} \cdot M^-_{q_i}\right\} \cup \left\{\tfrac{1}{2}N_m\right\}$$

and

$$\left\{\omega_m(b);\ b \in \tfrac{m}{q_i} \cdot M'_{q_i}\right\} \cup \left\{\tfrac{1}{2}N_m\right\}$$

are $\mathbb{Z}$-bases of the same subgroup $\mathrm{Cor}_{K_m/K_{q_i}}\big(\mathcal{S}'_{q_i}\big)$ of $\mathcal{S}'_m$, so that there is an integral transition matrix between these bases of determinant ± 1. We stress that the sets $\frac{m}{q_i} \cdot M^-_{q_i}$ (resp. $\frac{m}{q_i} \cdot M'_{q_i}$) for $i \in [\![1, t]\!]$ are pairwise disjoint. Hence, it is natural to define

$$\begin{aligned} M'_m &= \Big(M^-_m \ \setminus \ \bigcup_{i=1}^{t} \tfrac{m}{q_i} \cdot M^-_{q_i}\Big) \cup \Big(\bigcup_{i=1}^{t} \tfrac{m}{q_i} \cdot M'_{q_i}\Big) \\ &= \Big\{a \in M^-_m;\ \forall i \in [\![1, t]\!], \tfrac{m}{q_i} \nmid a\Big\} \cup \Big(\bigcup_{i=1}^{t} \Big\{\tfrac{mb}{q_i};\ 1 \leq b \leq \tfrac{\varphi(q_i)}{2}\Big\}\Big). \end{aligned} \tag{4.8}$$

which agrees with the previous definition of M'_{p^e}. Easily adapting the proof of Lem. 4.1 gives that for any $r \mid m$, $0 < r < m$, such that $\left(r, \frac{m}{r}\right) = 1$, we have

$$\left\{a \in M'_m;\ r \mid a\right\} = \left\{rb;\ b \in M'_{\frac{m}{r}}\right\} = r \cdot M'_{\frac{m}{r}}.$$

Thus, we have proved that Th. 4.2 and Pr. 4.4 implies the following:

Theorem 4.9. *For any integer $m > 1$, $m \not\equiv 2 \pmod 4$, the set*

$$\left\{\omega_m(a);\ a \in M'_m\right\} \cup \left\{\tfrac{1}{2}N_m\right\} \tag{4.10}$$

is a $\mathbb{Z}$-basis of $\mathcal{S}'_m$.

Proposition 4.11. *For any given $b \in \mathbb{Z}$, $0 < b < m$, let r_b be the maximal divisor of (b, m) satisfying $\left(r_b, \frac{m}{r_b}\right) = 1$, i.e., r_b is the product of all $q_i, i \in [\![1, t]\!]$ which divide b, and write $\omega_m(b) \in \mathcal{S}'_m$ as a unique $\mathbb{Z}$-linear combination of basis elements (4.10). Then, for each $a \in M'_m$ such that $r_b \nmid a$, the coefficient of $\omega_m(a)$ in this $\mathbb{Z}$-linear combination is equal to zero.*

Finally, keeping in mind that $\omega_m(a) = \theta_m(a) - \frac{1}{2}N_m$ if $m \nmid a$, we stress that all results of this whole section are equally valid when replacing $\omega_m(\cdot)$ by $\theta_m(\cdot)$, for example:

Corollary 4.12. *For any integer $m > 1$, $m \not\equiv 2 \pmod 4$, the set*

$$\left\{\theta_m(a);\ a \in M'_m\right\} \cup \left\{\tfrac{1}{2}N_m\right\} \tag{4.13}$$

is a $\mathbb{Z}$-basis of $\mathcal{S}'_m$.

Corollary 4.14. *For any given $b \in \mathbb{Z}$, $0 < b < m$, let r_b be the maximal divisor of (b, m) satisfying $\left(r_b, \frac{m}{r_b}\right) = 1$, i.e., r_b is the product of all $q_i, i \in [\![1, t]\!]$ which divide b, and write $\theta_m(b) \in \mathcal{S}'_m$ as a unique $\mathbb{Z}$-linear combination of basis elements (4.13). Then, for each $a \in M'_m$ such that $r_b \nmid a$, the coefficient of $\theta_m(a)$ in this $\mathbb{Z}$-linear combination is equal to zero.*

4.3 Short Basis of the Stickelberger Ideal

Elements of $\mathbb{Z}[G_m]$ are called *short* if they are of the form

$$\sum_{\sigma \in G_m} a_\sigma \sigma \in \mathbb{Z}\big[G_m\big], \qquad \text{where } a_\sigma \in \{0, 1\} \text{ for all } \sigma \in G_m.$$

We first exhibit a large family of short elements of $\mathcal{S}_m$. Choosing carefully elements from this family yields a basis (4.22) of $\mathcal{S}'_m$ with almost only short elements and also our short basis (4.30) of the Stickelberger ideal $\mathcal{S}_m = \mathcal{S}'_m \cap \mathbb{Z}\big[G_m\big]$.

4.3.1 A family of short elements of $\mathcal{S}_m$

In this section, we construct numerous short elements of $\mathcal{S}_m \subset \mathcal{S}'_m$ which we shall use later on.

Proposition 4.15. *Let $a, b, c \in \mathbb{Z}$ satisfy $m \nmid a$, $m \nmid b$, $m \nmid c$, $m \mid a + b + c$. Then*

$$\alpha = \theta_m(a) + \theta_m(b) + \theta_m(c) - N_m$$

is a short element of $\mathcal{S}_m$. Moreover $(1 + \sigma_{m,-1})\alpha = N_m$, so exactly one half of the coefficients of α are zeros.

Proof. Using $\theta_m(c) + \theta_m(-c) = N_m$ when $m \nmid c$ (see Eq. (2.17)), we obtain

$$\alpha = \theta_m(a) + \theta_m(b) - \theta_m(-c) = \sum_{\substack{0<s\le m\\(s,m)=1}} \Big(\big\{-\tfrac{as}{m}\big\} + \big\{-\tfrac{bs}{m}\big\} - \big\{\tfrac{cs}{m}\big\}\Big)\sigma_{m,s}^{-1}.$$

Since $0 \le \{x\} < 1$, every coefficient in the above sum is trivially bounded by

$$-1 < \big\{-\tfrac{as}{m}\big\} + \big\{-\tfrac{bs}{m}\big\} - \big\{\tfrac{cs}{m}\big\} < 2.$$

Moreover, let $[x] = x - \{x\} \in \mathbb{Z}$ be the integral part of x for any $x \in \mathbb{Q}$. Then,

$$\big\{-\tfrac{as}{m}\big\} + \big\{-\tfrac{bs}{m}\big\} - \big\{\tfrac{cs}{m}\big\} = -\tfrac{(a+b+c)s}{m} - \big[-\tfrac{as}{m}\big] - \big[-\tfrac{bs}{m}\big] + \big[\tfrac{cs}{m}\big] \in \mathbb{Z},$$

which proves that α is short. The last equality of the proposition follows again from Eq. (2.17) and an easy observation that $\sigma_{m,-1}\theta_m(a) = \theta_m(-a)$. □

4.3.2 Bases of $\mathcal{S}'_m$ with many short elements

We first describe the map α_m, which associates to any $b \in \mathbb{Z}$, $0 < b < m$, one short element from the family of Pr. 4.15. For any given $b \in \mathbb{Z}$, let r_b be the maximal divisor r of (b, m) satisfying the condition $(r, \frac{m}{r}) = 1$. In other words,

$$r_b = \prod_{i \in J_b} q_i, \qquad \text{where } J_b = \big\{i \in [\![1,t]\!];\ q_i \mid b\big\}.$$

Let $J'_b = [\![1,t]\!] \setminus J_b = \big\{i \in [\![1,t]\!];\ q_i \nmid b\big\}$, and let us suppose that $0 < b < m$ so that $J'_b \neq \emptyset$. We define $\alpha_m(b)$ as follows:

- If $\big|J'_b\big| > 1$, let $u = q_{\min J'_b}$, and $v = \frac{m}{ur_b}$. Since $(u, v) = 1$, the equation

  $$ux + vy = -1$$

  has a solution $x, y \in \mathbb{Z}$, where x is well-defined modulo v and y modulo u, so bux and bvy are well-defined modulo m. Let

  $$\alpha_m(b) = \theta_m(b) + \theta_m(bux) + \theta_m(bvy) - N_m. \tag{4.16}$$

- If $J'_b = \{j\}$ then $b = \frac{mc}{q_j}$ for a unique $c \in \mathbb{Z}$, $0 < c < q_j$. If $c > 1$ we define

  $$\alpha_m(b) = \theta_m(-b) + \theta_m\big(b - \tfrac{m}{q_j}\big) + \theta_m\big(\tfrac{m}{q_j}\big) - N_m, \tag{4.17}$$

  whereas if $c = 1$, so that $b = \frac{m}{q_j}$, we put

  $$\alpha_m(b) = 2\theta_m\big(\tfrac{m\cdot\varphi(q_j)}{2q_j}\big) + \theta_m\big(\tfrac{m}{p_j}\big) - N_m. \tag{4.18}$$

Intuitively, $\alpha_m(\cdot)$ is constructed by means of layers on $\big|J'_b\big|$, similarly to what happens for M_m^- as shown by Lem. 4.1. For $\big|J'_b\big| = 1$, we follow the prime power case of Th. 4.6, which is very similar to [Sch08, Th. 9.3(i)] when $m = p$. For $\big|J'_b\big| > 1$ we use Bezout's equality to write $-b$ as the sum of two summands bux and buy in such a way that both $\big|J'_{bux}\big|$ and $\big|J'_{bvy}\big|$ are strictly smaller than $\big|J'_b\big|$, so that both $\theta_m(bux)$ and $\theta_m(bvy)$ are generated by basis elements that were already chosen in the previous layers. Any way of achieving this property works. In particular, note that in the case $\big|J'_b\big| > 1$ we could use any other decomposition of $\frac{m}{r_b}$ into the product of relatively prime integers $u > 1$, $v > 1$.

Lemma 4.19. *For any integer $m > 1$, $m \not\equiv 2 \pmod 4$, the element $\alpha_m(b)$ is short and satisfies $(1 + \sigma_{m,-1})\alpha_m(b) = N_m$ for each positive integer $b < m$.*

Proof. In the former case $\left|J'_b\right| > 1$, we have $b + bux + bvy = 0$. Since $u \nmid b$, we have $u \nmid bvy$; similarly $v \nmid b$ implies $v \nmid bux$. Hence $\alpha_m(b)$ is short by Pr. 4.15. In the latter case $J'_b = \{j\}$ for some $j \in [\![1, t]\!]$, we have that b writes as $\frac{mc}{q_j}$ with $c \in \mathbb{Z}$ and $0 < c < q_j$, then $\alpha_m(b)$ is short by Pr. 4.15 again, because $-b + (b - \frac{m}{q_j}) + \frac{m}{q_j} = 0$ and $2 \cdot \frac{m \cdot \varphi(q_j)}{2q_j} + \frac{m}{p_j} = m$. □

Theorem 4.20. *For any integer $m > 1$, $m \not\equiv 2 \pmod 4$, the sets*

$$\left\{\alpha_m(b);\ b \in M'_m, \left|J'_b\right| > 1\right\} \cup \left\{\theta_m(b);\ b \in M'_m, \left|J'_b\right| = 1\right\} \cup \left\{\tfrac{1}{2}N_m\right\}, \tag{4.21}$$

$$\left\{\alpha_m(b);\ b \in M'_m \setminus \left\{\tfrac{m}{q_1}, \dots, \tfrac{m}{q_t}\right\}\right\} \cup \left\{\theta_m\left(\tfrac{m}{q_1}\right), \dots, \theta_m\left(\tfrac{m}{q_t}\right)\right\} \cup \left\{\tfrac{1}{2}N_m\right\} \tag{4.22}$$

are $\mathbb{Z}$-bases of $\mathcal{S}'_m$.

Proof. By definition of $\alpha_m(b)$ in Eqs. (4.16) and (4.17), we know that all elements of these sets belong to $\mathcal{S}'_m$. We shall show that the transition matrices from the set (4.13) to the set (4.21) and from the set (4.21) to the set (4.22) are, after a suitable reordering of elements of M'_m, triangular with ± 1 on the diagonal, which will prove the theorem.

At first, we deal with the transition matrix from the set (4.13) to the set (4.21) and we shall use induction with respect to $\left|J'_b\right|$. If $\left|J'_b\right| = 1$ then $\theta_m(b)$ belongs to both sets (4.13) and (4.21). So suppose that $\left|J'_b\right| > 1$. Then the transition from $\theta_m(b)$ to $\alpha_m(b)$ given in Eq. (4.16) uses $\theta_m(bux)$ and $\theta_m(bvy)$ and the coefficient of $\theta_m(b)$ is 1. By Cor. 4.14, $\theta_m(bux)$ is a $\mathbb{Z}$-linear combination of $\theta_m(a)$ for a running over M'_m such that $r_{bux} \mid a$. For these a's, we have that

$$J'_a \subseteq J'_{r_{bux}} = J'_{bux} \subsetneq J'_b,$$

since $\min J'_b \notin J'_{bux}$ by definition of u. Hence, all these $\theta_m(a)$ are covered by induction, and so is $\theta_m(bux)$. The case of $\theta_m(bvy)$ can be treated similarly.

Now, let us consider the transition matrix from the set (4.21) to the set (4.22). Suppose that $J'_b = \{j\}$ and $b = \frac{mc}{q_j}$ for some $c \in \mathbb{Z}$, $1 \le c \le \frac{\varphi(q_j)}{2}$. If $c = 1$ then $\theta_m(b)$ belongs to both sets (4.21) and (4.22). If $c > 1$ then the transition from $\theta_m(b)$ to $\alpha_m(b)$, by Eqs. (4.17) and (2.17), writes as

$$\alpha_m(b) = -\theta_m(b) + \theta_m\left(b - \tfrac{m}{q_j}\right) + \theta_m\left(\tfrac{m}{q_j}\right).$$

Since $J'_{b-m/q_j} = J'_{m/q_j} = J'_b$, both $\theta_m\left(b - \frac{m}{q_j}\right) = \theta_m\left(\frac{m}{q_j}(c-1)\right)$ and $\theta_m\left(\frac{m}{q_j}\right)$ were already covered by induction. The coefficient of $\theta_m(b)$ is -1. □

4.3.3 A basis of $\mathcal{S}_m$ with only short elements

Recall that the Stickelberger ideal of K_m is the intersection $\mathcal{S}_m = \mathcal{S}'_m \cap \mathbb{Z}[G_m]$. Let $\mathcal{S}''_m$ be the subgroup of $\mathcal{S}'_m$ generated by the set

$$\left\{\alpha_m(a);\ a \in M'_m\right\} \cup \left\{\tfrac{1}{2}N_m\right\}. \tag{4.23}$$

We shall prove that $\mathcal{S}''_m = \mathcal{S}_m + \frac{1}{2}N_m \cdot \mathbb{Z}$ and that Eq. (4.23) is its basis. We shall start by computing its finite index in $\mathcal{S}'_m$. First, we treat the prime power case.

Lemma 4.24. *Let $q = p^e > 2$, where p is a prime and e is a positive integer. Then the index of $\mathcal{S}''_q$ in $\mathcal{S}'_q$ is finite and*

$$[\mathcal{S}'_q : \mathcal{S}''_q] = \begin{cases} \frac{q}{2} & \text{if } p = 2, \\ q & \text{if } p > 2. \end{cases}$$

Proof. To obtain the index $[\mathcal{S}'_q : \mathcal{S}''_q]$, let us compute the transition matrix from

$$\left\{\theta_q(a);\ a \in \mathbb{Z}, 1 \le a \le \tfrac{\varphi(q)}{2}\right\} \cup \left\{\tfrac{1}{2}N_q\right\}, \tag{4.25}$$

which is a $\mathbb{Z}$-basis of $\mathcal{S}'_q$ by Cor. 4.12, to the system of generators of $\mathcal{S}''_q$, i.e.,

$$\left\{\alpha_q(a);\ a \in \mathbb{Z}, 1 \le a \le \tfrac{\varphi(q)}{2}\right\} \cup \left\{\tfrac{1}{2}N_q\right\}. \tag{4.26}$$

This transition matrix is given by Eqs. (4.18) and (4.17). More precisely, using also Eq. (2.17), we obtain in the studied special case that

$$\alpha_q(a) = \begin{cases} \theta_q\left(p^{e-1}\right) + 2\theta_q\left(\frac{\varphi(q)}{2}\right) - N_q & \text{if } a = 1, \\ \theta_q(1) + \theta_q(a-1) - \theta_q(a) & \text{if } 2 \le a \le \frac{\varphi(q)}{2}. \end{cases}$$

Since $\frac{1}{2}N_q$ belongs to both sets (4.25) and (4.26), we can ignore this element in the computation of the determinant of the transition matrix.

At first, let us assume that $p > 3$. Then $p^{e-1} < p^{e-1} \cdot \frac{p-1}{2} = \frac{\varphi(q)}{2}$. We shall compute the determinant of the following square matrix of dimension $\frac{\varphi(q)}{2}$

$$\begin{pmatrix} 0 & 0 & 0 & 0 & \cdots & 1 & \cdots & 0 & 0 & 2 \\ 2 & -1 & 0 & 0 & \cdots & 0 & \cdots & 0 & 0 & 0 \\ 1 & 1 & -1 & 0 & \cdots & 0 & \cdots & 0 & 0 & 0 \\ 1 & 0 & 1 & -1 & \cdots & 0 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & & \vdots & & \vdots & \vdots & \vdots \\ 1 & 0 & 0 & 0 & \cdots & 0 & \cdots & 1 & -1 & 0 \\ 1 & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & 1 & -1 \end{pmatrix}, \tag{4.27}$$

where the 1 in the first row belongs to the p^{e-1}th column (which is the first column if $e = 1$). The sum of all rows but the first one, multiplied by 2, equals

$$\left(\varphi(q)\ 0\ 0\ 0 \cdots 0 \cdots 0\ 0\ -2\right).$$

We add this row to the first row of our matrix. If $e > 1$, we also add to the first row the sum of all rows from the second one to the p^{e-1}th one, i.e.,

$$\left(p^{e-1}\ 0\ 0\ 0 \cdots -1 \cdots 0\ 0\ 0\right).$$

After this computation we get a lower triangular matrix of determinant $\pm q$. As this determinant is nonzero, the set (4.26) is a $\mathbb{Z}$-basis of $\mathcal{S}''_q$ and the index $[\mathcal{S}'_q : \mathcal{S}''_q]$ equals the absolute value of the determinant. The lemma follows for $p > 3$.

Now, suppose $p = 3$. Then $p^{e-1} = \frac{\varphi(q)}{2}$ and the square transition matrix of dimension 3^{e-1} writes as

$$\begin{pmatrix} 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 3 \\ 2 & -1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 1 & 1 & -1 & 0 & \cdots & 0 & 0 & 0 \\ 1 & 0 & 1 & -1 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots \\ 1 & 0 & 0 & 0 & \cdots & 1 & -1 & 0 \\ 1 & 0 & 0 & 0 & \cdots & 0 & 1 & -1 \end{pmatrix}.$$

If $e = 1$ then the only entry of our matrix of dimension 1 is 3. If $e > 1$, the sum of all rows but the first one, multiplied by 3, is equal to

$$\left(3^e \; 0 \; 0 \; 0 \cdots 0 \cdots 0 \; 0 \; -3\right).$$

Adding this row to the first row, we again get a lower triangular matrix of determinant $\pm q$, which gives the lemma in the case $p = 3$.

Finally, we treat the case $p = 2$. Then, by Eqs. (2.16) and (2.17), we have

$$\theta_q(2^{e-1}) = \theta_q(2^{e-1} - q) = \theta_q(-2^{e-1}) = N_q - \theta_q(2^{e-1}),$$

so $\theta_q(2^{e-1}) = \frac{1}{2}N_q$. Therefore we have got almost the same matrix as written in Eq. (4.27), except that in the first row the only non-zero element is the 2 at the very end. By the same approach as above, we obtain that the determinant of this matrix is equal to $\pm\varphi(q) = \pm\frac{q}{2}$ and the lemma in the case $p = 2$ follows. □

Proposition 4.28. *For any integer $m > 1$, $m \not\equiv 2 \pmod 4$, the set (4.23) is a basis of $\mathcal{S}''_m$, whose finite index in $\mathcal{S}'_m$ is given by*

$$[\mathcal{S}'_m : \mathcal{S}''_m] = \begin{cases} \frac{m}{2} & \text{if } m \text{ is even,} \\ m & \text{if } m \text{ is odd.} \end{cases}$$

Proof. This is similar to the proof of Th. 4.9. The following sets are pairwise disjoint for $i \in [\![1, t]\!]$

$$\left\{\tfrac{mb}{q_i};\ b \in M'_{q_i}\right\} = \left\{a;\ a \in M'_m,\ \tfrac{m}{q_i} \mid a\right\}.$$

Since $\mathrm{Cor}_{K_m/K_{q_i}}$ is an injective linear map, the transition matrix from the $\mathbb{Z}$-basis (4.21) of $\mathcal{S}'_m$, given by Th. 4.20, to the system of generators (4.23) of $\mathcal{S}''_m$ is a block diagonal matrix, having (besides plenty of trivial blocks of dimension 1 containing 1) one nontrivial block for each $i \in [\![1, t]\!]$. For a given i, the nontrivial block is equal to the matrix considered in Lem. 4.24 for $q = q_i$. Since the determinant of this transition matrix is equal to the product of determinants of these nontrivial blocks, it is nonzero and the proposition follows. □

We are now ready to state our main theorem, which in particular implies the afore-mentioned relation $\mathcal{S}''_m = \mathcal{S}_m + \frac{1}{2}N_m \cdot \mathbb{Z}$.

Theorem 4.29. *For any integer $m > 1$, $m \not\equiv 2 \pmod 4$, the set*

$$\left\{\alpha_m(a);\ a \in M'_m\right\} \cup \left\{N_m\right\} \tag{4.30}$$

is a $\mathbb{Z}$-basis of the Stickelberger ideal $\mathcal{S}_m$ of K_m having only short elements.

Proof. Let $\widetilde{\mathcal{S}_m}$ denote the subgroup of $\mathcal{S}'_m$ generated by the set (4.30). Each element of (4.30) is short by Lem. 4.19, in particular it belongs to $\mathbb{Z}[G_m]$, so that

$$\widetilde{\mathcal{S}_m} \subseteq \mathbb{Z}[G_m] \cap \mathcal{S}'_m = \mathcal{S}_m. \tag{4.31}$$

The indices $[\mathcal{S}'_m : \mathcal{S}_m] = w$ and $[\mathcal{S}'_m : \mathcal{S}''_m] = \frac{w}{2}$ are given by Lem. 2.26 and Pr. 4.28, respectively. In particular, by Pr. 4.28, the set (4.23) is linearly independent; comparing with the set (4.30), we see that the set (4.30) is also linearly independent and that $\widetilde{\mathcal{S}_m}$ is a subgroup of $\mathcal{S}''_m$ of index $[\mathcal{S}''_m : \widetilde{\mathcal{S}_m}] = 2$. Hence,

$$[\mathcal{S}'_m : \widetilde{\mathcal{S}_m}] = [\mathcal{S}'_m : \mathcal{S}''_m] \cdot [\mathcal{S}''_m : \widetilde{\mathcal{S}_m}] = w = [\mathcal{S}'_m : \mathcal{S}_m],$$

and the inclusion (4.31) gives $\widetilde{\mathcal{S}_m} = \mathcal{S}_m$. The theorem follows. □

4.4 An Upper Bound on the Relative Class Number

Our short basis of the Stickelberger ideal $\mathcal{S}_m$, given in Th. 4.29, allows to derive a simple upper bound on the relative class number of *any* cyclotomic field.

Corollary 4.32. *Let $m > 1$ be an integer satisfying $m \not\equiv 2 \pmod 4$, let t be the number of primes dividing m. The relative class number h_m^- of the m-th cyclotomic field satisfies*

$$h_m^- \leq 2^{1-a} \cdot \left(\tfrac{\varphi(m)}{8}\right)^{\varphi(m)/4},$$

where $\varphi()$ is Euler's totient function and

$$a = \begin{cases} 0 & \text{if } t = 1, \\ 2^{t-2} - 1 & \text{if } t \geq 2. \end{cases} \tag{4.33}$$

Proof. Recall that, for any integer s relatively prime to m, $\sigma_{m,s} \in G_m$ denotes the automorphism of the m-th cyclotomic field K_m sending any m-th root of unity to its s-th power. In particular, $\sigma_{m,-1}$ is the restriction of the complex conjugation. Following Sinnott, let $\mathcal{R}_m = \mathbb{Z}[G_m]$ and

$$\begin{aligned} \mathcal{R}_m^- &= \{\alpha \in \mathcal{R}_m;\ (1 + \sigma_{m,-1})\alpha = 0\}, \\ \mathcal{A}_m &= \{\alpha \in \mathcal{R}_m;\ (1 + \sigma_{m,-1})\alpha \in N_m\mathbb{Z}\}. \end{aligned}$$

Moreover, for any submodule $M \subseteq \mathcal{R}_m$ we define $M^- = M \cap \mathcal{R}_m^-$. Using [Sin80, Lem. 1.2(a)], multiplication by $1 + \sigma_{m,-1}$ gives

$$[\mathcal{A}_m : \mathcal{S}_m] = [(1 + \sigma_{m,-1})\mathcal{A}_m : (1 + \sigma_{m,-1})\mathcal{S}_m] \cdot [\mathcal{A}_m^- : \mathcal{S}_m^-].$$

It is clear that $(1 + \sigma_{m,-1})\mathcal{A}_m = (1 + \sigma_{m,-1})\mathcal{S}_m = N_m\mathbb{Z}$ and that $\mathcal{A}_m^- = \mathcal{R}_m^-$. Therefore, using [Sin78, Th., p.107] and the remark following Lem. 2.19, we have

$$[\mathcal{A}_m : \mathcal{S}_m] = [\mathcal{R}_m^- : \mathcal{S}_m^-] = 2^a \cdot h_m^-, \tag{4.34}$$

where a is defined by Eq. (4.33).

We use our short basis Eq. (4.30) of $\mathcal{S}_m$ given in Th. 4.29 to get a bound on $[\mathcal{A}_m : \mathcal{S}_m]$. First, a $\mathbb{Z}$-basis of $\mathcal{A}_m$ is given by

$$\left\{\beta_m(s);\ 1 \leq s < \tfrac{m}{2}, (s,m) = 1\right\} \cup \left\{\gamma_m\right\}, \tag{4.35}$$

where $\beta_m(s) = \sigma_{m,s} - \sigma_{m,-s}$ and

$$\gamma_m = \sum_{\substack{1 \leq s < \frac{m}{2} \\ (s,m)=1}} \sigma_{m,s}.$$

An easy calculation gives

$$N_m = 2\gamma_m - \sum_{\substack{1 \leq s < \frac{m}{2} \\ (s,m)=1}} \beta_m(s).$$

For each $b \in M_m'$, let us define integers $a_{b,s}$, where $1 \leq s < m$, $(s,m) = 1$, by

$$\alpha_m(b) = \sum_{\substack{1 \leq s < m \\ (s,m)=1}} a_{b,s}\sigma_{m,s}.$$

By Lem. 4.19, we have $a_{b,s} + a_{b,m-s} = 1$, so that

$$\alpha_m(b) = \gamma_m + \sum_{\substack{1 \le s < \frac{m}{2} \\ (s,m)=1}} (a_{b,s} - 1)\beta_m(s).$$

The index $[\mathcal{A}_m : \mathcal{S}_m]$ is given by the absolute value of the determinant of the transition matrix from the basis (4.30) of $\mathcal{S}_m$ to the basis (4.35) of $\mathcal{A}_m$, i.e.,

$$\left[\mathcal{A}_m : \mathcal{S}_m\right] = \left| \det \left(\begin{array}{c|ccc} 2 & -1 & \cdots & -1 \\ \hline 1 & & & \\ \vdots & & \left(a_{b,s} - 1\right)_{\substack{b \in M'_m \\ 1 \le s < \frac{m}{2},\, (s,m)=1}} & \\ 1 & & & \end{array} \right) \right|.$$

We subtract one half of the first row from each of the other rows to get

$$\left[\mathcal{A}_m : \mathcal{S}_m\right] = \left| \det \left(\begin{array}{c|ccc} 2 & -1 & \cdots & -1 \\ \hline 0 & & & \\ \vdots & & \left(a_{b,s} - \frac{1}{2}\right)_{\substack{b \in M'_m \\ 1 \le s < \frac{m}{2},\, (s,m)=1}} & \\ 0 & & & \end{array} \right) \right| = 2 \cdot \left| \det \left(\left(a_{b,s} - \frac{1}{2}\right)_{\substack{b \in M'_m \\ 1 \le s < \frac{m}{2},\, (s,m)=1}} \right) \right|. \tag{4.36}$$

By Lem. 4.19 we know that $a_{b,s} \in \{0, 1\}$, and so $a_{b,s} - \frac{1}{2} \in \left\{-\frac{1}{2}, \frac{1}{2}\right\}$. So the length of each row of this matrix, as a vector in the Euclidean space of dimension $\frac{\varphi(m)}{2}$, is equal to $\frac{1}{2}\sqrt{\frac{\varphi(m)}{2}}$. Therefore, by Hadamard's inequality,

$$\left[\mathcal{A}_m : \mathcal{S}_m\right] \le 2 \cdot \left(\frac{1}{2}\sqrt{\frac{\varphi(m)}{2}}\right)^{\varphi(m)/2}.$$

A comparison with Eq. (4.34) gives the corollary. □

Remark 4.37. For the marginal cases where $4 \nmid \frac{\varphi(m)}{2}$, better bounds exist for these scaled Hadamard matrices (see [BEHC21]) that directly translate into slightly better bounds for h_m^-. We do not dive into the details here.

4.5 Effective Short Stickelberger Generators

Let $m > 1$ satisfy $m \not\equiv 2 \pmod 4$. Let ℓ be any prime such that $(\ell, m) = 1$ and let $\mathfrak{L}$ be a fixed (unramified) prime ideal above ℓ of inertia degree f in the m-th cyclotomic field K_m. The aim of this section is to describe an algebraic integer of K_m generating the principal ideal $\mathfrak{L}^{\alpha_m(b)}$ for each $b \in M'_m$.

Of course, we shall use Gauss sums. Recall that for any positive integer r, we let $\zeta_r = e^{2\pi i/r}$. Let $\mathbb{F} = \mathbb{Z}[\zeta_m]/\mathfrak{L}$ be the finite field of cardinality $\mathcal{N}(\mathfrak{L}) = \ell^f$, and let $\chi_{\mathfrak{L}}$ be the m-th power Legendre symbol with respect to $\mathfrak{L}$, i.e., for any $a \in \mathbb{F}^\times$, the m-th root of unity $\chi_{\mathfrak{L}}(a) \in \langle \zeta_m \rangle$ is

determined by the condition that $\chi_{\mathfrak{L}}(a)$ belongs to the class $a^{(\mathcal{N}(\mathfrak{L})-1)/m}$. We extend as usual characters to $\mathbb{F}$ by setting $\chi_{\mathfrak{L}}(0) = 0$. For any integer b, we have the following Gauss sum, where $\mathrm{Tr} : \mathbb{F} \to \mathbb{F}_\ell$ is the trace map in the field extension $\mathbb{F}/\mathbb{F}_\ell$,

$$g_{\mathfrak{L}}(b) = -\sum_{y\in\mathbb{F}} \chi_{\mathfrak{L}}(y)^b \zeta_\ell^{\mathrm{Tr}(y)} \in \mathbb{Z}[\zeta_{m\ell}].$$

For any integers $u \equiv 1 \pmod m$, $\ell \nmid u$, and $v \equiv 1 \pmod \ell$, $(v, m) = 1$, an easy computation gives (see e.g., [Sin80, (3.3) and (3.5)])

$$\sigma_{m\ell,u}\big(g_{\mathfrak{L}}(b)\big) = \chi_{\mathfrak{L}}(u)^{-b} \cdot g_{\mathfrak{L}}(b), \tag{4.38}$$

$$\sigma_{m\ell,v}\big(g_{\mathfrak{L}}(b)\big) = g_{\mathfrak{L}}(vb). \tag{4.39}$$

Hence, $g_{\mathfrak{L}}(b)^m \in \mathbb{Z}[\zeta_m]$ by Eq. (4.38). Moreover, we have the well-known Stickelberger factorization (see e.g., [Sin80, (3.4)])

$$g_{\mathfrak{L}}(b)^m \cdot \mathbb{Z}[\zeta_m] = \mathfrak{L}^{m\theta_m(b)}. \tag{4.40}$$

We want to describe an explicit generator of the principal ideal $\mathfrak{L}^{\alpha_m(b)}$ for each $b \in M'_m$. Since each $\alpha_m(b)$ is given by the general construction from Pr. 4.15 (see the proof of Lem. 4.19), we shall start more generally.

Proposition 4.41. *For any $a, b \in \mathbb{Z}$ such that $m \nmid a$, $m \nmid b$, $m \nmid a + b$, let*

$$\alpha = \theta_m(a) + \theta_m(b) - \theta_m(a + b) \in \mathbb{Z}[G_m]$$

be one of the short elements given by Pr. 4.15. Then the Jacobi sum

$$\mathcal{J}_{\mathfrak{L}}(a, b) = -\sum_{y\in\mathbb{F}} \chi_{\mathfrak{L}}(y)^a \chi_{\mathfrak{L}}(1 - y)^b \in \mathbb{Z}[\zeta_m]$$

satisfies $\mathcal{J}_{\mathfrak{L}}(a, b) \cdot \mathbb{Z}[\zeta_m] = \mathfrak{L}^\alpha$.

Proof. By [Was97, Lem. 6.2(d)], we have

$$\mathcal{J}_{\mathfrak{L}}(a, b) = \frac{g_{\mathfrak{L}}(a)g_{\mathfrak{L}}(b)}{g_{\mathfrak{L}}(a + b)}.$$

Thus, the result directly follows from Eq. (4.40) and the fact $\mathcal{J}_{\mathfrak{L}}(a, b) \in \mathbb{Z}[\zeta_m]$. □

As an example of application of Pr. 4.41, let us consider any $b \in M'_m$ such that $\big|J'_b\big| > 1$. Then the short element $\alpha_m(b)$ is given by Eq. (4.16), so that

$$\mathfrak{L}^{\alpha_m(b)} = \mathcal{J}_{\mathfrak{L}}(bux, bvy) \cdot \mathbb{Z}[\zeta_m],$$

where $u = q_{\min J'_b}$, $v = \frac{m}{ur_b}$, and $x, y \in \mathbb{Z}$ satisfy $ux + vy = -1$.

Furthermore, it is clear that u, v, x, y do not depend on b but only on J'_b. Therefore, having another $c \in M'_m$ such that $J'_c = J'_b$, there is an integer s relatively prime to m satisfying $c \equiv sb \pmod m$, so that Eq. (4.39) gives

$$\mathcal{J}_{\mathfrak{L}}(cux, cvy) = \mathcal{J}_{\mathfrak{L}}(sbux, sbvy) = \sigma_{m,s}\big(\mathcal{J}_{\mathfrak{L}}(bux, bvy)\big).$$

Hence, computing generators of $\mathfrak{L}^{\alpha_m(b)}$, for all $b \in M'_m$ with $\big|J'_b\big| > 1$, comes down to the computation of exactly one representative Jacobi sum per set J'_b, then applying a suitable automorphism to obtain the generator for $\mathfrak{L}^{\alpha_m(c)}$ whenever $J'_c = J'_b$.

4.6 Practical Results

We implemented in practice the computation of our short Stickelberger bases from Th. 4.29 using SageMath [Sag20] on an Intel® Core™ i7-8650U @3.2GHz.

All involved algebraic criteria are very easy to compute, so that obtaining the short bases is actually a matter of seconds for any reasonable conductor. We verified, for all conductors $m < 10000$, $m \not\equiv 2 \pmod 4$, such that $\varphi(m) \leq 2000$, that the Hermite Normal Form (HNF) of the short basis from Th. 4.29 coincides with the HNF of the large basis from [Kuč92, Th. 6.2].

We stress that using a naive trial-and-error strategy to extract a short basis from a large set of short vectors, e.g., from the set W of [CDW21, §4.2], may converge only after a huge number of iterations, each involving the computation of a costly HNF. This is especially hazardous when the number of prime divisors of m grows, e.g., our brute force experiment for $m = 780 = 2^2 \cdot 3 \cdot 5 \cdot 13$ never finished despite the relatively small dimension $\varphi(m) = 192$.

More interestingly, we used the determinant formula for $\big[\mathcal{A}_m : \mathcal{S}_m\big]$ given in Eq. (4.36) to derive the relative class number h_m^- from Eq. (4.34). We checked, for the same range of conductors as above, that the obtained values coincide with the values given by the analytic class number formula given in Eq. (2.10). Surprisingly, we observed that the determinant computation is very competitive, especially when the number of distinct prime factors of m is small. Some comparative timings are provided in Tab. 4.1.

m	$q_1 \dots q_t$	$\varphi(m)$	Time h_m^- (s)	
			Analytic	$[\mathcal{A}_m : \mathcal{S}_m]$
1139	$17 \cdot 67$	1056	12.6	8.1
1495	$5 \cdot 13 \cdot 23$	1056	7.6	7.9
4140	$2^2 \cdot 3^2 \cdot 5 \cdot 23$	1056	4.8	8.5
2283	$3 \cdot 761$	1520	25.1	21.8
2865	$3 \cdot 5 \cdot 191$	1520	16.3	21.0
1951	1951	1950	78.8	60.3
2171	$13 \cdot 167$	1992	57.6	35.6
2495	$5 \cdot 499$	1992	53.8	41.7
6012	$2^2 \cdot 3^2 \cdot 167$	1992	28.3	40.2

Table 4.1 – Comparative timings for computing the relative class number h_m^- using resp. the analytic formula Eq. (2.10) and the index formula for $[\mathcal{A}_m : \mathcal{S}_m]$ in Eq. (4.36), for a few representative examples.

Finally, we verified that relations $\mathfrak{L}^{\alpha_m(b)} = \mathcal{J}_{\mathfrak{L}}(a_1, a_2) \cdot \mathbb{Z}[\zeta_m]$ hold true in small dimensions (up to $\varphi(m) = 80$). We note that computing explicitly such generators using the Jacobi sum formalism is very easy for any m. For instance, taking $m = 2003$ and $\ell = 48073 \equiv 1 \pmod m$, the computation of all $\varphi(m)/2$ generators corresponding to $\mathfrak{L}^{\alpha_m(b)}$, for all $b \in M'_m$ and some $\mathfrak{L}$ above ℓ takes under 15 minutes, i.e., less than 1 second per generator.

By contrast, using suitable combinations of Gauss sums to obtain e.g., generators for the relations $\mathfrak{L}^{(a-\sigma_{m,a})\cdot\theta_m(-1)}$ of [Was97, 6.9] imposes to work in $\mathbb{Q}\big[\zeta_{m\ell}\big]$. Even using all available algorithmic tricks, such as using sparse polynomials modulo $x^{m\ell} - 1$, replacing divisions by the use of the identity $g_{\mathfrak{L}}(b) \cdot g_{\mathfrak{L}}(-b) = \pm\mathcal{N}(\mathfrak{L})$ [Was97, 6.1(b)] and profitting from Eq. (4.39), this is arguably intractable in the above case when $m\ell = 96\,290\,219$, and still takes over 39 seconds per generator when restricting to the first split prime $\ell = 4007$.

Chapter 5

Log-$\mathcal{S}$-unit Lattices Using Explicit Stickelberger Generators to Solve Approx Ideal-SVP

Now, we present our last contribution, which is a joint effort with Andrea Lesavourey, Tuong-Huy Nguyen and Adeline Roux-Langlois.

> **[BLNR21] Log-$\mathcal{S}$-unit lattices using Explicit Stickelberger Generators to solve Approx Ideal-SVP.**
> Olivier Bernard, Andrea Lesavourey, Tuong-Huy Nguyen and Adeline Roux-Langlois.

Accepted for publication in the proceedings of Asiacrypt 2022, in *Lecture Notes in Computer Science* (LNCS) Series, Springer..

Keywords: Ideal lattices, Approx-SVP, Stickelberger, S-unit attacks, Twisted-PHS algorithm.
Links: [ePrint: 2021/1384[11] | GitHub: ob3rnard/Tw-Sti[8] | Blog: H2020 Prometheus[12]]

Contents

[11] https://eprint.iacr.org/2021/1384
[8] https://github.com/ob3rnard/Tw-Sti
[12] https://www.h2020prometheus.eu/dissemination/blog

5.1 Introduction

Even though the theoretically proven trade-off between runtime and approximation factor is the same for the Twisted-PHS algorithm as for the PHS algorithm (see Th. 3.1 and 3.14), experimentally, very significant improvements compared to the original PHS algorithm are illustrated in Fig. 3.4 and 3.5. In particular, the implementation provided in GitHub: ob3rnard/Twisted-PHS[7] allowed us to test the Twisted-PHS algorithm in number fields of degree up to 60, while achieving much better approximation factors than the original [PHS19a] implementation. However, reaching larger degrees was limited by the classical complexity of the algorithm.

5.1.1 Our contributions

Our first contribution is to succeed in performing new experiments on the Twisted-PHS algorithm, in almost all cyclotomic fields up to degree 210, thanks to a novel approach which allows us to significantly improve the running time of the preprocessing phase. The approximation factors obtained in our experiments, as detailed in Fig. 5.1, show that the Twisted-PHS algorithm performs much better (over the considered experimental range) than the CDW algorithm, which was the previously best-known algorithm. More interestingly, the obtained approximation factors are comparable to the volumetric lower bound for the CDW algorithm experimentally obtained in [DPW19] in the prime conductor case, and sometimes even smaller.

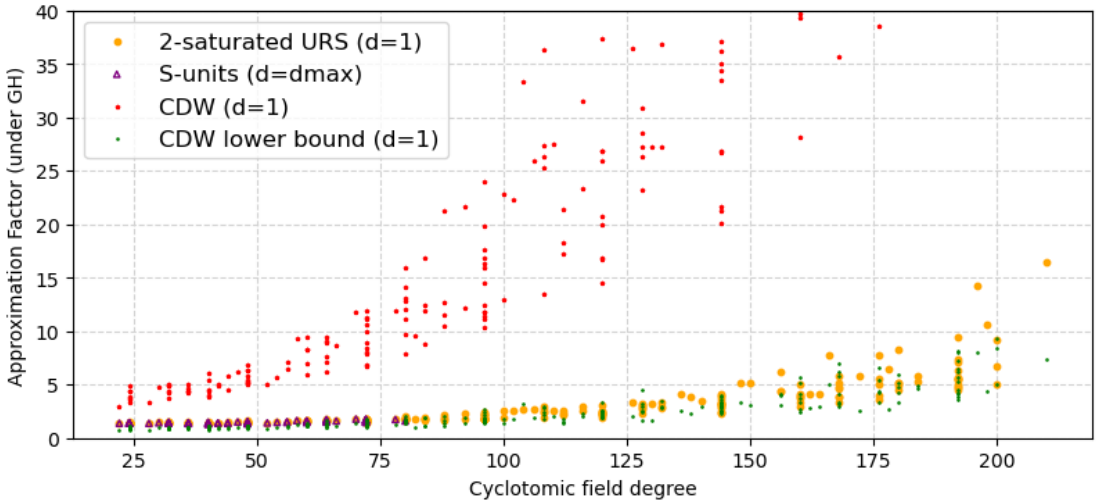

FIGURE 5.1 – Approximation factors comparison for cyclotomic fields K_m of degree $\varphi(m) \leq 210$ with $h_m^+ = 1$, under the Gaussian Heuristic. Our results, labelled as "2-saturated URS", bound Twisted-PHS from above.

We stress that one main goal of our experiments is to break the small dimension barrier and reach ranges of parameters where asymptotic phenomena — e.g., the exponential growth of the

[7] https://github.com/ob3rnard/Twisted-PHS

class number — start to express. Though experimental data as the one we obtain are not enough to provide an asymptotical result concerning the approximation factor, we stress that pushing experiments up to degree 210 is a significant breakthrough, and to our knowledge, no other experiments of $\mathcal{S}$-unit attacks beyond degree 70 have been publicly reported two years after those presented in Ch. 3. We run the (practical version of the) Twisted-PHS algorithm using full-rank log-$\mathcal{S}$-unit sublattices on simulated random targets to see how the final approximation factor evolves with the dimension in our regime. Additionally, we compute several geometrical parameters on the basis obtained by our implementation to study their quality, as done in §3.4: the root-Hermite factor δ_0, the orthogonality defect δ and the logarithm of the Gram-Schmidt norms. We are able to confirm the peculiar geometric nature of the log-$\mathcal{S}$-unit lattice already observed in §§3.4.1 and 3.4.2, across all considered cyclotomic fields, sublattices and factor bases. These recurrent observations in very different regimes suggest that this phenomenon has a possibly deep explanation, an observation that has been recently developed by Bernstein and Lange [BL21]. For example, to give an idea of the striking ease of reduction of these log-$\mathcal{S}$-unit sublattices, we report that in our biggest field example, BKZ_{40} terminates in around 7 minutes (resp. 30) for our lattices in dimension 1154 (resp. 1574), which is unusually fast at these dimensions. Moreover, we provide a full implementation, which is publicly available at GitHub: ob3rnard/Tw-Sti[8].

Due to the classical complexity of computing $\mathcal{S}$-units, reaching degrees beyond 100 is not directly possible using the algorithms given in Ch. 3, and is the result of both theoretical and implementational improvements. We compute full-rank sublattices of the log-$\mathcal{S}$-unit lattice for cyclotomic fields K_m of *any* conductor m from degree 20 up to degree 210. To obtain these results, our main theoretical contribution is to exhibit in §5.2 a full-rank family of *independent* $\mathcal{S}$-units lifted from the maximal real subfield K_m^+ of K_m. One step of this construction is to use explicit Stickelberger generators that are easy to compute using Jacobi sums, as shown in §4.5 and specialized in §5.2.1 with additional insights. Hence, we obtain a full-rank sublattice of the log-$\mathcal{S}$-unit lattice, at the much lower cost of computing class group relations in the maximal real subfield of half degree. We also provide in Th. 5.14 a closed formula for the multiplicative index of this full-rank family inside the whole $\mathcal{S}$-unit group. This index allows to quantify the comparison between our new approach and the previous one from Ch. 3. Though we first obtain sublattices of large index in the full log-$\mathcal{S}$-unit lattice, we are able to mitigate it by using classical saturation techniques recalled in §5.2.4.

As a minor contribution, we apply these results to show in §5.3 how to benefit from these explicit Stickelberger generators to remove most quantum steps of the CDW algorithm [CDW21]. Namely, we remove the last Pip resolution, and also, under a relatively harmless restriction that the plus part of the class number verifies $h_m^+ \leq O(\sqrt{m})$ (Hyp. 5.18), the random walk to the relative class group, replaced with a single call to a quantum class group computation in dimension $\varphi(m)/2$. The latter should also yield in practice better approximation factors, by allowing to choose the finite places of $\mathcal{S}$ of smallest possible norms.

5.1.2 Technical overview

Let $\mathcal{S}$ be a set of places where the finite places correspond to a collection of full Galois orbits of split prime ideals. Our full-rank family $\mathfrak{F}$ of independent $\mathcal{S}$-units is composed of three parts:

1. circular units, defined e.g., in [Was97, §8] and for which an explicit basis can be found in [Kuč92, Th. 6.1];
2. Stickelberger generators, as explicitly given by the proof of Stickelberger's theorem, see for example [Sin80, Eq. (3.4)];

[8] https://github.com/ob3rnard/Tw-Sti

3. real $\mathcal{S}^+$-units (apart from real units), where $\mathcal{S}^+$ is the set $\mathcal{S} \cap K_m^+$ of places of $\mathcal{S}$ restricted to the maximal real subfield K_m^+ of K_m.

In the context of the cryptanalysis of id-SVP, the set of circular units has already been used to reduce the size of principal ideal generators in [CDPR16, CDW17] for m being a prime power, in [HWB17] when m has two distinct prime factors and finally in [CDW21] in the general case. Using free relations in the class group Cl_m coming from Stickelberger's theorem was suggested in [CDW17, CDW21], where many *short* relations were identified [CDW21, Lem. 4.4]. For the first time, we use for cryptanalysis the main results from Ch. 4:

- the knowledge of an explicit *short* $\mathbb{Z}$*-basis* of the Stickelberger ideal for *any* conductor from Th. 4.29,
- the effective computation of generators corresponding to these short relations, using Jacobi sums as in §4.5.

Compared to [CDW17, CDW21], we stress that only knowing a short generating set of the Stickelberger ideal is not necessarily sufficient for our purpose. Indeed, though it would be possible to build a basis from such a generating set to solve the CVP like in [CDW21, Cor. 2.2] without any geometric loss, using e.g., [MG02, Lem. 7.1], the slight Euclidean norm growth of the obtained basis vectors however translates into a dramatic increase of the size of the (possibly rational) coefficients of the corresponding generators, in a way that significantly hinders subsequent computations. In particular, in order to climb dimensions as far as possible and best approach log-$\mathcal{S}$-unit lattices using the saturation process described in §5.2.4, it is crucial to constrain both the number of elements we use and their size, i.e., to use a *short basis* of the Stickelberger lattice. As for the last part, obtaining a full-rank lattice of class relations was done in [CDW17] using relative norm relations $\mathcal{N}_{K_m/K_m^+}(\mathfrak{L}) = \mathfrak{L}^{1+\tau}$, where the $\mathfrak{L}$'s are chosen in the relative class group, to obtain the so-called "extended Stickelberger lattice". We extend this result by considering the lattice of all real class relations between the relative norms of ideals of any class.

The multiplicative index of this family in the full $\mathcal{S}$-unit group is explicitly given by our Th. 5.14. This index contains a large power of 2 that can be removed using classical 2-saturation techniques of §5.2.4, leading to a family $\mathfrak{F}_{\mathsf{sat}}$.

Removing quantum steps from the CDW algorithm.

In the context of the CDW algorithm, we first propose in §5.3 an equivalent rewriting of [CDW21, Alg. 7] that enlightens some hidden steps that reveal useful for subsequent modifications. Then, we plug the explicit Stickelberger generators and real generators described above to remove the last call to the quantum PIP solver. Finally, by considering the module of *all* real class group relations, we remove the need of a random walk mapping any ideal of K_m into Cl_m^-, at the small price of restricting to cyclotomic fields such that $h_m^+ \leq O(\sqrt{m})$ (Hyp. 5.18), whereas [CDW21, Ass. 2] uses $h_m^+ \leq \mathrm{poly}(m)$. Then, only two quantum steps remain: the first is performed only once in dimension $\frac{\varphi(m)}{2}$ to compute real class group relations and generators, the second is for solving the ClDLP for each query.

Simulating the Twisted-PHS algorithm.

Finally, we apply the practical version of the Twisted-PHS algorithm from §3.3 on our full-rank sublattices of the log-$\mathcal{S}$-unit lattice. This is actually an *approximated* mode of the Twisted-PHS algorithm, as Twisted-PHS normally uses the full log-$\mathcal{S}$-unit lattice for an optimal number of orbits $d = d_{\mathsf{max}}$ maximizing the density of the full log-$\mathcal{S}$-unit lattice, as predicted by Alg. 3.3, which we estimated using the analytic class number formula. However, in our case, the family $\mathfrak{F}_{\mathsf{sat}}$

has index roughly $(h_m^-)^{d-1}$, which is sufficiently large so that this optimal factor base phenomenon does not hold. More precisely, the density of the log-$\mathcal{S}$-unit sublattice associated to $\mathfrak{F}_{\mathsf{sat}}$ decreases as soon as $d > 1$.

We fully implement the construction of the lattices associated to $\mathfrak{F}$, $\mathfrak{F}_{\mathsf{sat}}$ and to fundamental elements of the full $\mathcal{S}$-unit group $\mathfrak{F}_{\mathsf{su}}$ when available (up to degree 80) for the first d split prime orbits with $d \in [\![1, d_{\mathsf{max}}]\!]$, including the computation of Stickelberger generators and real generators. We evaluate the geometry of all these lattices with standard indicators described in §2.4.3, and observe consistently the same phenomenons already observed in §§3.4.1 and 3.4.2, that indicate close to orthogonal lattices. Moreover, as computing ClDlp solutions for random ideals is not possible, we simulate the query phase *via* random targets. The approximation factors obtained in this mode give an upper bound on what can be expected when using Twisted-PHS. Notably, they are already much smaller than the ones obtained using the CDW algorithm, and sometimes beat the volumetric lower bound experimentally derived in [DPW19]. We stress that, up to degree 80 when the full $\mathcal{S}$-unit group is computable, our results match, under the Gaussian Heuristic, the *exact* approximation factors obtained by Fig. 3.4.

Remark 5.1. Similar techniques for the construction of $\mathcal{S}$-units may be used in a concurrent work by Bernstein, Eisenträger, Rubin, Silverberg and van Vredendaal, as announced in a talk by Bernstein on 20th August 2021 at SIAM Conference in the power of 2 conductor case up to degree 64 assuming $h_{2^e}^+ = 1$.

5.2 An Explicit Full-Rank Family of Independent $\mathcal{S}$-units

In this section, we exhibit a full-rank family of *independent* $\mathcal{S}$-units, where the finite places of $\mathcal{S}$ correspond to a collection of full Galois orbits of split prime ideals. As mentioned in introduction, this family is composed of three parts:

1. Circular units are given in §2.2.4 using the material from [Kuč92, Th. 6.1];
2. Stickelberger generators are given in §4.5 in the general case, and specialized in §5.2.1 in the split case together with additional remarks on their complex embeddings and on how our short basis relates to the results of [CDW17, CDW21];
3. Real $\mathcal{S}^+$-units (apart from real units), where $\mathcal{S}^+ = \mathcal{S} \cap K_m^+$, are in §5.2.2.

Considering real $\mathcal{S}^+$-units and proving in §5.2.3 the multiplicative index of our family in the full $\mathcal{S}$-unit group constitute our main theoretical contributions. Finally, the saturation process used to mitigate this index is described in §5.2.4.

Remark 5.2. Recall that the prime factorization of $m \not\equiv 2 \mod 4$ is written as $m = q_1 q_2 \cdots q_t$, where $q_i = p_i^{e_i} > 2$ for $i \in [\![1, t]\!]$. The rest of the section uses the subsets M_m^+ and M_m' of $[\![1, m]\!]$ from resp. §2.2.1 and Eq. (4.8) to describe resp. a fundamental family of circular units and a short $\mathbb{Z}$-basis of the Stickelberger ideal of K_m.

5.2.1 Stickelberger generators

Recall from §2.2.5 that the Stickelberger ideal provides free relations in the class group of K_m, by Stickelberger's fundamental theorem Th. 2.22. In this section, we essentially rephrase Pr. 4.15 and Th. 4.29, giving additional insight on how these results relate to [CDW17, CDW21].

A short basis of the Stickelberger lattice.

An element of the integral group ring $\mathbb{Z}[G_m]$ is called *short* if it is of the form $\sum_{\sigma \in G_m} a_\sigma \sigma$ in $\mathbb{Z}[G_m]$, where $a_\sigma \in \{0, 1\}$ for all $\sigma \in G_m$. Short elements of $\mathcal{S}_m$ have been identified in

[Sch08, Th. 9.3(i) and Ex. 9.3] in the prime conductor case, and the proof has been adapted to any conductor in [CDW21, Lem. 4.4] to prove the shortness of the following generating set:

$$W = \{w_a;\ a \in [\![2, m]\!]\}, \quad \text{with } w_a = \theta_m(1) + \theta_m(a-1) - \theta_m(a). \tag{5.3}$$

Note that using $\theta_m(a)+\theta_m(-a) = N_m$ when $m \nmid a$, we obtain $w_a = w_{m-a+1}$ whenever $1 < a < m$, and that $w_m = N_m$ using also $\theta_m(m) = 0$. Hence, W is the set $\{w_a;\ 2 \le a \le \lceil \frac{m}{2} \rceil\} \cup \{N_m\}$.

We emphasize that only knowing a generating set of short elements as in [CDW21] is not necessarily sufficient. Indeed, though it would be possible to build a basis from this generating set to solve the CVP like in [CDW21, Cor. 2.2] without any geometric loss using e.g., [MG02, Lem. 7.1], the slight Euclidean norm growth of the obtained basis vectors however translates into a dramatic increase of the size of the (possibly rational) coefficients of the corresponding generators, in a way that significantly hinders subsequent computations. In particular, in order to climb dimensions as far as possible and best approach log-$\mathcal{S}$-unit lattices using the saturation process described in §5.2.4, it is crucial to constrain both the number of elements we use and their size, i.e., to use a *basis* of the Stickelberger lattice containing only *short* elements. Such a basis has been explicitly given in Th. 4.29, and can be computationally easily extracted from a very large family of short elements Pr. 4.15 encompassing $W \setminus \{N_m\}$ by Eq. (5.3):

Proposition 5.4 (Adapted from Pr. 4.15)**.** *Let $a, b \in \mathbb{Z}$ satisfying $m \nmid a$, $m \nmid b$ and $m \nmid (a+b)$. Then $\alpha = \theta_m(a) + \theta_m(b) - \theta_m(a+b)$ is a short element of $\mathcal{S}_m$. Moreover, $(1+\tau)\cdot\alpha = N_m$, so exactly one half of the coefficients of α are zeros.*

Note that the second part of the proposition actually specifies [CDW21, Lem. 4.4(3)]: it implies that the ℓ_2-norm of any $w \in W \setminus \{N_m\}$, viewed as a vector in $\mathbb{Z}^{\varphi(m)} \simeq_{\mathbb{Z}} \mathbb{Z}[G_m]$, is *exactly* $\sqrt{\varphi(m)/2}$.

Theorem 5.5 (Adapted from Th. 4.29)**.** *There exists efficiently computable elements $\alpha_m(b)$, for $b \in [\![1, m]\!]$, such that $\alpha_m(b)$ is a short element from Pr. 5.4 and $\{\alpha_m(b);\ b \in M'_m\} \cup \{N_m\}$ is a $\mathbb{Z}$-basis of the Stickelberger lattice $\mathcal{S}_m$ of K_m.*

We stress that when m is a prime, this basis coincides with the one given by [Sch08, Th. 9.3(i)] and with the set W described in Eq. (5.3).

Effective Stickelberger generators using Jacobi sums.

As previously mentioned, the proof of Stickelberger's theorem (see Th. 2.22) is explicit, i.e., for any $\alpha \in \mathcal{S}_m$ and any fractional ideal $\mathfrak{b}$ of K_m, it builds an explicit $\gamma \in K_m$ such that $\langle\gamma\rangle = \mathfrak{b}^\alpha$ [Was97, §6.2], [Sin80, §3.1]. Moreover, when α is a short basis element from Th. 5.5, it turns out that γ has a suprisingly simple expression using Jacobi sums as in §4.5.

We briefly specialize §4.5 to the split case here. Let $\ell \in \mathbb{Z}$ be a prime such that $\ell \equiv 1 \mod m$, and let $\mathfrak{L}$ be any fixed (split) prime ideal of K_m above ℓ. Let a, b be such as in Pr. 5.4, then for $\alpha = \theta_m(a)+\theta_m(b)-\theta_m(a+b)$, we have that $\mathfrak{L}^\alpha$ is a principal ideal generated by the following Jacobi sum (see Pr. 4.41:

$$\mathcal{J}_{\mathfrak{L}}(a, b) = -\sum_{u \in \mathcal{O}_{K_m}/\mathfrak{L}} \chi_{\mathfrak{L}}^a(u)\chi_{\mathfrak{L}}^b(1-u) \quad \in K_m, \tag{5.6}$$

where $\chi_{\mathfrak{L}}(u) \in \langle\zeta_m\rangle$ verifies $\chi_{\mathfrak{L}}(u) \equiv u^{(\ell-1)/m} \mod \mathfrak{L}$, for any $u \in (\mathcal{O}_{K_m}/\mathfrak{L})^\times$, and $\chi_{\mathfrak{L}}(0) = 0$. When $\alpha = \alpha_m(c)$ for $c \in M'_m$, we shall write $\gamma_{\mathfrak{L},c}^-$ for the generator of $\mathfrak{L}^{\alpha_m(c)}$. Using a discrete logarithm table for elements of $\mathcal{O}_{K_m}/\mathfrak{L}^\times$, the computation, for a fixed prime $\mathfrak{L}$, of all Jacobi

sums corresponding to the short basis $\{\alpha_m(c);\ c \in M'_m\}$ is very fast. As noted in §4.5, the Galois group also acts on the involved Jacobi sums in a way that allows to replace some of the Jacobi sum computations by the application of a suitable automorphism.

Finally, as a direct consequence of [Was97, Lem. 6.1], all these Jacobi sums are ℓ-Weil numbers, i.e., they verify the Weil relation $\mathcal{J}_{\mathfrak{L}}(a,b)\overline{\mathcal{J}_{\mathfrak{L}}(a,b)} = \ell$, for a and b as above. This implies that for all $\sigma \in G_m$, we actually have $\left|\sigma\left(\mathcal{J}_{\mathfrak{L}}(a,b)\right)\right| = \sqrt{\ell}$, meaning that any of these elements is *the shortest* generator of its corresponding ideal $\mathfrak{L}^{\alpha}$, which has algebraic norm $\ell^{\varphi(m)/2}$.

Remark 5.7. By Eq. (5.3), generators corresponding to the short relations of W write as:

$$\mathcal{J}_{\mathfrak{L}}(1,a) \cdot \mathcal{O}_{K_m} = \mathfrak{L}^{w_a}, \qquad \text{for any } a \in [\![2, m-1]\!].$$

5.2.2 Real $\mathcal{S}^+$-units

Since $\sharp M'_m = \frac{\varphi(m)}{2}$, a consequence of Th. 5.5 is that the Stickelberger lattice has rank $\frac{\varphi(m)}{2}+1$ in $\mathbb{Z}[G_m]$; in particular, it is not full rank, hence cannot be directly used as a lattice of class relations. In previous works, obtaining a full-rank lattice in $\mathbb{Z}[G_m]$ from $\mathcal{S}_m$ was done by projecting into $(1-\tau)\mathcal{S}_m$ [CDW21, §4.3], or by the adjonction of $(1+\tau)\mathbb{Z}[G_m]$ [CDW17, Def. 2]. Both can be used as a lattice of class relations for the *relative* class group Cl_m^-. In particular, the so-called *augmented* Stickelberger lattice $\mathcal{S}_m + (1+\tau)\mathbb{Z}[G_m]$ annihilates the relative class group and has full rank in $\mathbb{Z}[G_m]$, as shown in [CDW17, Lem. 2].

We generalize this result by considering the module of all real class group relations between relative norm ideals of ideals from the entire class group Cl_m. In §5.2.3, we shall prove that the Stickelberger lattice augmented with these real class group relations yields a lattice of class relations for the *whole* class group. Note that, as opposed to other modules like $(1-\tau)\mathcal{S}_m$ or $\mathcal{S}_m + (1+\tau)\mathbb{Z}[G_m]$, real class group relations actually depend on the underlying prime ideals.

On one hand, this affects negatively the shortness of the obtained relation vectors: putting those in Hermite Normal Form, we shall see later that each relation, viewed as a vector of integer valuations, has ℓ_2-norm at most h_m^+. On the other hand, removing the constraint to belong to the relative class group brings a significant practical and theoretical gap: first, it allows to choose prime ideals of smallest possible norms, which as shown in §3.2.3 or [CDW21, Th. 4.8] lowers in practice the obtained approximation factor; second, whereas prime ideals of norm at most Bach's bound are sufficient to generate the entire class group, prime generators for the *relative* class group are only proven to be of norm bounded by the *larger* bound $(2.71 \cdot h_m^+ \cdot \ln \Delta_{K_m} + 4.13)^2$ from [Wes18].

Lifting real class group relations.

Let $\ell_1, \ldots, \ell_d$ be distinct prime integers satisfying $\ell_i \equiv 1 \mod m$, so that ℓ_i splits in K_m. For each $i \in [\![1,d]\!]$, fix a prime ideal $\mathfrak{L}_i \mid \ell_i$ in K_m of norm ℓ_i, and let $\mathfrak{l}_i = \mathcal{N}_{K_m/K_m^+}(\mathfrak{L}_i) = \mathfrak{L}_i^{1+\tau} \cap K_m^+$ be the relative norm ideal of $\mathfrak{L}_i$. Since $\mathfrak{L}_i$ is a split prime ideal of K_m dividing ℓ_i, the ideal $\mathfrak{l}_i$ is a split prime ideal of K_m^+ of norm ℓ_i, and by Kummer-Dedekind's theorem we have $\mathfrak{l}_i \cdot \mathcal{O}_{K_m} = \mathfrak{L}_i^{1+\tau}$. This justifies the slight abuse of notation of writing $\mathfrak{l}_i^{\sigma} = \mathfrak{L}_i^{(1+\tau)\sigma} \cap K_m^+$, for any $\sigma \in G_m$.

We are interested in the real class group relations between all prime ideals in the G_m^+-orbits of the $\mathfrak{l}_i$, i.e., between the following prime ideals of K_m^+:

$$\left\{\mathfrak{l}_i^{\sigma_s};\ i \in [\![1,d]\!], 0 < s < \tfrac{m}{2}, (s,m) = 1\right\}. \tag{5.8}$$

The important point is that any class group relation in K_m^+ between ideals from Eq. (5.8) translates to a class group relation in K_m using repeatedly that $\mathfrak{l}_i^{\sigma} \cdot \mathcal{O}_{K_m} = \mathfrak{L}_i^{(1+\tau)\sigma}$. More precisely,

let $(r_1, \dots, r_d) \in \mathbb{Z}[G_m^+]^d$ represent a real class group relation in K_m^+ between ideals $\{\mathfrak{l}_i^{\sigma_s}\}$ of Eq. (5.8), i.e., there exists $\gamma_r^+ \in K_m^+$ such that $\gamma_r^+ \cdot \mathcal{O}_{K_m^+} = \prod_{i=1}^d \mathfrak{l}_i^{r_i}$. Then, this real class group relation naturally lifts to a class group relation $((1+\tau)\cdot r_1, \dots, (1+\tau)\cdot r_d)$ in K_m between prime ideals in the G_m-orbits $\{\mathfrak{L}_i^\sigma;\ i \in [\![1,d]\!], \sigma \in G_m\}$ as:

$$\gamma_r^+ \cdot \mathcal{O}_{K_m} = \prod_{i=1}^{d} \mathfrak{L}_i^{(1+\tau)r_i}. \tag{5.9}$$

Let $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ denote the lattice of all real class group relations between elements of the G_m^+-orbits of $\{\mathfrak{l}_i;\ i \in [\![1,d]\!]\}$. Concretely, it is the kernel of the following map:

$$\mathfrak{f}_{\mathfrak{l}_1,\dots,\mathfrak{l}_d} : (r_{i,s})_{\substack{1\le i\le d,\\ 0<s<m/2,(s,m)=1}} \in \mathbb{Z}^{d\cdot\frac{\varphi(m)}{2}} \longmapsto \prod_{i,s} \left[\mathfrak{l}_i^{\sigma_s}\right]^{r_{i,s}} \in \mathrm{Cl}_m^+ . \tag{5.10}$$

Using the canonical isomorphism of $\mathbb{Z}$-modules $\mathbb{Z}^{d\cdot\frac{\varphi(m)}{2}} \simeq_{\mathbb{Z}} \mathbb{Z}[G_m^+]^d$, the lattice of real class group relations $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ may be viewed as a $\mathbb{Z}$-submodule of $\mathbb{Z}[G_m^+]^d$. Lifting all these relations back to K_m as in Eq. (5.9), we therefore obtain the submodule $(1+\tau)\cdot C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d} \subseteq (1+\tau)\mathbb{Z}[G_m]^d$, that we shall call the lattice of *real class group relations* between the G_m-orbits of $\{\mathfrak{L}_i;\ i \in [\![1,d]\!]\}$.

Remark 5.11. When $h_m^+ = 1$, $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ is isomorphic to d copies of the integral group ring $\mathbb{Z}[G_m^+]$ and the lattice of real class relations is simply $(1+\tau)\mathbb{Z}[G_m]^d$.

Euclidean norm of real class relations.

We now identify a real class group relation from $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ to a vector in $\mathbb{Z}^{d\cdot\frac{\varphi(m)}{2}}$. In other words, we consider only the valuations of these relations on the G_m^+-orbits of the prime ideals $\mathfrak{l}_1, \dots, \mathfrak{l}_d$. Furthermore, $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ is put in Hermite Normal Form, conveniently for the proof, but better bounds might easily be obtained using e.g., the LLL algorithm.

Proposition 5.12. *Suppose the lattice $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ of real class relations is in HNF. Then, for all $\mathbf{w} \in C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d} \subseteq \mathbb{Z}[G_m^+]^d$, we have $\|\mathbf{w}\|_2 \le \|\mathbf{w}\|_1 \le h_m^+$.*

This means that $(1+\tau)\cdot C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ can be used in the CDW algorithm instead of $(1+\tau)\mathbb{Z}[G_m]$, as we will see in §5.3, while still reaching the same asymptotic approximation factor as long as $h_m^+ \le O(\sqrt{\varphi(m)})$ (Hyp. 5.18). This slightly more restrictive (see the discussion in §2.2.3) hypothesis will be more than compensated by the fact that it removes the need for the $\mathfrak{l}_i$'s to be principal, which has a significant impact in practice on the algebraic norm of the chosen ideals, and thus on the final approximation factor reached in [CDW21, Alg. 6].

Proof. The image of the map $\mathfrak{f}_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ given in Eq. (5.10) is a subgroup of Cl_m^+, so the volume of its kernel $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ is at most h_m^+. By definition of the Hermite Normal Form,[19] $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ has diagonal elements $h_1, \dots, h_{\varphi(m)/2} > 0$, and the j-th column contains integers c_{ij} such that $0 \le c_{ij} < h_j$ for $i < j$ and $c_{ij} = 0$ for $i > j$. We shall prove $h_i + \sum_{i<j} c_{ij} \le h_i \cdot \prod_{i<j} h_j$ for any row of fixed index $i \in [\![1, \frac{\varphi(m)}{2}]\!]$, which yields the result. This is done by induction on the dimension, using repeatedly the fact that for any integers $x, y \ge 1$, $x + (y-1) \le (xy)$. □

[19] In this proof, we consider an upper-triangular HNF with row vectors.

Explicit real generators.

For each relation $r = (r_1, \dots, r_d) \in C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$, we compute an explicit $\gamma_r^+ \in K_m^+ \subsetneq K_m$ that verifies Eq. (5.9). Together with the unit group $\mathcal{O}^\times_{K_m^+}$ of K_m^+, they form a fundamental system of $\mathcal{S}^+$-units, where the finite places of $\mathcal{S}^+$ are the G_m^+-orbits of the relative norm ideals $\mathfrak{l}_i$.

In the next section, we shall see that adding the explicit Stickelberger generators of §5.2.1 to these real generators yields a maximal set of independent $\mathcal{S}$-units in the degree $\varphi(m)$ cyclotomic field K_m, at the much smaller cost of computing a fundamental system of real $\mathcal{S}^+$-units in K_m^+ of degree only $\frac{\varphi(m)}{2}$.

In practice, though this remains the main bottleneck of our experimental setting, it allows us to push effectively our experiments up to degree $\varphi(m) = 210$, whereas the (full) $\mathcal{S}$-units computations of §3.4 were bound to $\varphi(m) = 70$.

5.2.3 An $\mathcal{S}$-unit subgroup of finite index

As in §5.2.2, let $\ell_1, \dots, \ell_d$ be prime integers satisfying $\ell_i \equiv 1 \mod m$; for each i, fix a (split) prime ideal $\mathfrak{L}_i \mid \ell_i$ in K_m and let $\mathfrak{l}_i = \mathfrak{L}_i \cap K_m^+$. Let $\mathcal{S}$ be a set of places containing, apart the infinite places of K_m, all G_m-orbits of the $\mathfrak{L}_i$'s. Combining the results of §§2.2.4, 5.2.1 and 5.2.2, we get the following family of $\mathcal{S}$-units:

$$\mathfrak{F} = \left\{v_a;\ a \in M_m^+\right\} \cup \left\{\gamma^-_{\mathfrak{L}_i,b};\ i \in [\![1,d]\!], b \in M'_m\right\} \cup \left\{\gamma_r^+;\ r \in C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}\right\} \tag{5.13}$$

where the first set is the set of *circular units* given by Th. 2.14, the second is the set of explicit *Stickelberger generators* stated at the end of §5.2.1 and the last one is the set of *real generators* as in Eq. (5.9).

This family has $\left(\varphi(m)/2 - 1\right) + d \cdot \varphi(m)$ elements, which matches precisely the multiplicative rank of the full $\mathcal{S}$-unit group modulo torsion $\mathcal{O}^\times_{K_m,\mathcal{S}}/\mu\left(\mathcal{O}^\times_{K_m}\right)$.[20] In this section, we prove that these $\mathcal{S}$-units are indeed independent and we compute the index of the subgroup of $\mathcal{O}^\times_{K_m,\mathcal{S}}$ generated by those elements.

Theorem 5.14. *Let $h_{m,(\mathfrak{L}_1,\dots,\mathfrak{L}_d)}$ (resp. $h^+_{m,(\mathfrak{l}_1,\dots,\mathfrak{l}_d)}$) be the cardinal of the subgroup of Cl_m (resp. Cl_m^+) generated by the G_m-orbits of $\mathfrak{L}_1, \dots, \mathfrak{L}_d$ (resp. the G_m^+-orbits of $\mathfrak{l}_1, \dots, \mathfrak{l}_d$). The family $\mathfrak{F}$ given in Eq. (5.13) is a maximal set of independent $\mathcal{S}$-units. The subgroup generated by $\mathfrak{F}$ in $\mathcal{O}^\times_{K_m,\mathcal{S}}/\mu\left(\mathcal{O}^\times_{K_m}\right)$ has index:*

$$\left(\frac{h_m \cdot h^+_{m,(\mathfrak{l}_1,\dots,\mathfrak{l}_d)}}{h_{m,(\mathfrak{L}_1,\dots,\mathfrak{L}_d)}}\right) \cdot 2^b \cdot \left(h_m^-\right)^{d-1} \cdot \left(2^{\frac{\varphi(m)}{2}-1} \cdot 2^a\right)^d,$$

where $a = b = 0$ if m is a prime power, and $a = 2^{t-2} - 1$, $b = 2^{t-2} + 1 - t$ whenever m has t distinct prime divisors.

Note that when the G_m-orbits of the $\mathfrak{L}_i$'s generate Cl_m, the first term in this index equals h_m^+. As we shall see in §5.2.4, the powers of 2 can be killed by standard saturation techniques, so the real problem comes from the $(h_m^-)^{d-1}$ part, which has generically *huge* prime factors. Intuitively, this comes from the fact that the Stickelberger relations miss all class group relations that exist between two (or more) distinct G_m-orbits.

First, we show that the lattice obtained by adding one copy of the Stickelberger ideal for each G_m-orbit, to the lattice $(1+\tau) \cdot C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ of real class relations, yields a full-rank submodule

[20] Note that for our purpose, the torsion units play no role and can thus be put aside.

of $\mathbb{Z}[G_m]^d$. Hence, we have obtained a full-rank lattice of class relations for the union of all the G_m-orbits above $\ell_1, \dots, \ell_d$.

We begin by restricting our attention to the case $d = 1$. We need the following lemma, which extends and proves an observation already made in [DPW19, Rem. 3] in the prime conductor case:

Lemma 5.15. *The index of $\mathcal{S}_m + (1+\tau)\cdot\mathbb{Z}[G_m^+]$ in $\mathbb{Z}[G_m]$ is finite:*

$$\left[\mathbb{Z}[G_m] : \mathcal{S}_m + (1+\tau)\cdot\mathbb{Z}[G_m^+]\right] = 2^{\varphi(m)/2-1}\cdot 2^a \cdot h_m^-,$$

where $a = 0$ if $t = 1$ and $a = 2^{t-2}-1$ otherwise, where m has t prime divisors.

Proof. The proof is due to R. KUČERA. First, note that $(1+\tau)\cdot\mathbb{Z}[G_m^+]$ contains N_m, hence by Th. 5.5, $\mathcal{S}_m + (1+\tau)\cdot\mathbb{Z}[G_m^+]$ is generated by the following $\varphi(m)$ elements:

$$\left\{\alpha_m(b);\ b \in M_m'\right\} \cup \left\{(1+\tau)\sigma_s;\ 0 < s < \tfrac{m}{2}, (s,m) = 1\right\}.$$

Therefore, its index is given by the absolute value of the determinant of the transition matrix from the canonical basis of $\mathbb{Z}[G_m]$ to the above generating set:

$$\left[\mathbb{Z}[G_m] : \mathcal{S}_m + (1+\tau)\cdot\mathbb{Z}[G_m^+]\right] = \left|\det\begin{pmatrix} \left\{a_{b,s}\right\}_{\substack{b\in M_m' \\ 0<s<m,\,(s,m)=1}} \\ \hline 1 \quad 1 \\ \ddots \qquad \ddots \\ 1 \qquad\qquad 1 \end{pmatrix}\right|,$$

where for any $b \in M_m'$, we write $\alpha_m(b) = \sum_{\sigma_s \in G_m} a_{b,s}\sigma_s$. Subtracting suitable combinations of rows of the lower half of this matrix to rows of the upper half to cancel the upper right block, this is the absolute value of the determinant of the square matrix of dimension $\frac{\varphi(m)}{2}$ with coefficients $\left\{a_{b,s} - a_{b,-s}\right\}$, for $b \in M_m'$ and s prime with m such that $0 < s < \frac{m}{2}$. By Pr. 5.4, $a_{b,s} + a_{b,-s} = 1$, which implies that $a_{b,s} - a_{b,-s} = 2a_{b,s} - 1$, therefore we recognize the matrix appearing in Eq. (4.36) at the very end of the proof of Cor. 4.32 with each coefficient being multiplied by 2. Combining this with Eq. (4.34), we obtain:

$$\left[\mathbb{Z}[G_m] : \mathcal{S}_m + (1+\tau)\cdot\mathbb{Z}[G_m^+]\right] = 2^{\frac{\varphi(m)}{2}}\cdot\tfrac{1}{2}\left[\mathcal{R}_m^- : \mathcal{S}_m^-\right],$$

and the result follows from Sinnott's theorem (see Th. 2.24). □

When $h_m^+ = 1$, the lattice of real class relations is always $(1+\tau)\cdot\mathbb{Z}[G_m^+]$, and Lem. 5.15 gives the whole story. In the general case $h_m^+ \neq 1$, we deduce:

Lemma 5.16. *Let ℓ be a prime integer that splits in K_m, let $\mathfrak{L} \mid \ell$ in K_m and let $\mathfrak{l} = \mathfrak{L}^{1+\tau} \cap K_m^+$. Let $h_{m,(\mathfrak{l})}^+$ be the cardinal of the subgroup of Cl_m^+ generated by the G_m^+-orbit of $\mathfrak{l}$ in K_m^+. The $\mathbb{Z}$-module generated by $\mathcal{S}_m$ and the lattice $(1+\tau)\cdot C_{\mathfrak{l}}^+$ of real class group relations of the G_m-orbit of $\mathfrak{L}$, has finite index in $\mathbb{Z}[G_m]$:*

$$\left[\mathbb{Z}[G_m] : \mathcal{S}_m + (1+\tau)\cdot C_{\mathfrak{l}}^+\right] = 2^{\varphi(m)/2-1}\cdot 2^a \cdot h_m^- \cdot h_{m,(\mathfrak{l})}^+,$$

where $a = 0$ if $t = 1$ and $a = 2^{t-2}-1$ otherwise, where m has t prime divisors.

Proof. By definition of $C_{\mathfrak{l}}^+$ as the kernel of the map $\mathfrak{f}_{\mathfrak{l}}$ of Eq. (5.10), we have:

$$\left[\mathbb{Z}[G_m^+] : C_{\mathfrak{l}}^+\right] = h_{m,(\mathfrak{l})}^+ = \left[(1+\tau)\cdot\mathbb{Z}[G_m^+] : (1+\tau)\cdot C_{\mathfrak{l}}^+\right].$$

Note also that N_m belongs to $(1+\tau)\cdot C_{\mathfrak{l}}^+ \subseteq (1+\tau)\cdot\mathbb{Z}[G_m^+]$, hence, again by means of transition matrix:

$$\left[\mathcal{S}_m + (1+\tau)\cdot\mathbb{Z}[G_m^+] : \mathcal{S}_m + (1+\tau)\cdot C_{\mathfrak{l}}^+\right] = \left[(1+\tau)\cdot\mathbb{Z}[G_m^+] : (1+\tau)\cdot C_{\mathfrak{l}}^+\right].$$

Finally, putting things together with Lem. 5.15, the result comes from:

$$\begin{aligned}\left[\mathbb{Z}[G_m] : \mathcal{S}_m + (1+\tau)\cdot C_{\mathfrak{l}}^+\right] &= \left[\mathbb{Z}[G_m] : \mathcal{S}_m + (1+\tau)\cdot\mathbb{Z}[G_m^+]\right] \\ &\quad\cdot\left[\mathcal{S}_m + (1+\tau)\cdot\mathbb{Z}[G_m^+] : \mathcal{S}_m + (1+\tau)\cdot C_{\mathfrak{l}}^+\right] \\ &= \left(2^{\varphi(m)/2-1}\cdot 2^a\cdot h_m^-\right)\cdot\left[\mathbb{Z}[G_m^+] : C_{\mathfrak{l}}^+\right].\end{aligned}$$

□

Finally, for the case where there are $d \geq 1$ orbits, a reasoning very similar to the proofs of Lem. 5.15 and 5.16 leads to:

Proposition 5.17. *Let $h_{m,(\mathfrak{l}_1,\ldots,\mathfrak{l}_d)}^+$ be the cardinal of the subgroup of Cl_m^+ generated by all G_m^+-orbits of $\mathfrak{l}_1,\ldots,\mathfrak{l}_d$. Then, the $\mathbb{Z}$-module generated by the lattice $(1+\tau)\cdot C_{\mathfrak{l}_1,\ldots,\mathfrak{l}_d}^+ \subseteq (1+\tau)\cdot\mathbb{Z}[G_m^+]^d$ of real class relations between the G_m-orbits of the $\mathfrak{L}_i$'s, and the diagonal block matrix of d copies of $\left(\mathcal{S}_m \setminus N_m\mathbb{Z}\right)$, verifies:*

$$\left[\mathbb{Z}[G_m]^d : \mathcal{S}_m^d + (1+\tau)\cdot C_{\mathfrak{l}_1,\ldots,\mathfrak{l}_d}^+\right] = \left(2^{\varphi(m)/2-1}\cdot 2^a\cdot h_m^-\right)^d\cdot h_{m,(\mathfrak{l}_1,\ldots,\mathfrak{l}_d)}^+.$$

Proof of Th. 5.14. The independence comes from Pr. 5.17 and the trivial fact that circular units are independent from Stickelberger and real generators. The index of the subgroup generated by $\mathfrak{F}$ in $\mathcal{O}_{K_m,\mathcal{S}}^\times/\mu\left(\mathcal{O}_{K_m}^\times\right)$ is given by:

$$\left[\mathcal{O}_{K_m}^\times : C_m\right]\cdot\frac{\left[\mathbb{Z}[G_m]^d : \mathcal{S}_m^d + (1+\tau)\cdot C_{\mathfrak{l}_1,\ldots,\mathfrak{l}_d}^+\right]}{\left|\det\left(\ker\mathfrak{f}_{\mathcal{S}}\right)\right|},$$

where $\ker\mathfrak{f}_{\mathcal{S}}$ is the lattice of all class group relations between finite places of $\mathcal{S}$. The first term is given by Pr. 2.12, the numerator of the second term is given by Pr. 5.17, and by definition of $\mathcal{O}_{K_m,\mathcal{S}}^\times$, the denominator is precisely $h_{m,(\mathfrak{L}_1,\ldots,\mathfrak{L}_d)}$. Rearranging terms adequately yields the result. □

5.2.4 Saturation

Saturation is a standard tool of computational algebraic number theory that has been used in various contexts like unit and class group computations, and can be traced back at least to [PZ89, §5.7].

Intuitively, the e-saturation procedure applied to $\mathfrak{F}$ consists in detecting e-th powers in the subgroup generated by $\mathfrak{F}$, including their e-th roots in the set, using e.g., the generalized Montgomery's e-th-root algorithm from [Tho12, §3], and rebuilding a basis of multiplicatively independent elements. At the end, the index of the new basis is no longer divisible by e. Remark that the output size does not depend on e, but only on the number and size of the elements of $\mathfrak{F}$. However, as the relative class number h_m^- in the index of Th. 5.14 hides *huge* prime factors, this strategy is at first glance hopeless in general to obtain the full $\mathcal{S}$-unit group from $\mathfrak{F}$.

As the index given by Th. 5.14 is divisible by a large power of 2, it is nonetheles natural to 2-saturate $\mathfrak{F}$ in order to mitigate its exponential growth, obtaining the 2-saturated family $\mathfrak{F}_{\mathsf{sat}}$. In the following, we briefly describe the 2-saturation procedure we use, and refer to e.g., [BFHP21, §4.3] for a formal exposition.

Recognizing squares.

Let $U = \langle g_1, \dots, g_k \rangle$ be a finitely generated multiplicative subgroup of $\mathcal{O}_{K_m,\mathcal{S}}^{\times}$. The first step of the 2-saturation process is to recognize squares in $U \cap (\mathcal{O}_{K_m,\mathcal{S}}^{\times})^2$. This is done by using local information provided by quadratic characters.

Fix a prime $\mathfrak{p} \notin \mathcal{S}$ such that $\mathcal{N}(\mathfrak{p}) \equiv 1 \mod \operatorname{lcm}(m, 2)$. Define $\chi_{\mathfrak{p}}$ as the Legendre symbol such that $\chi_{\mathfrak{p}}(a) \equiv a^{(\mathcal{N}(\mathfrak{p})-1)/2} \mod \mathfrak{p}$ for any $a \in U$. As $\mathfrak{p} \notin \mathcal{S}$ and $a \in \mathcal{O}_{K_m,\mathcal{S}}^{\times}$, we have that $\chi_{\mathfrak{p}}(a) \in \{-1, 1\}$. If a is a square, $\chi_{\mathfrak{p}}(a) = 1$ as a is still a square modulo $\mathfrak{p}$. The converse is not true, but by considering many characters $\chi_{\mathfrak{p}_1}, \dots, \chi_{\mathfrak{p}_N}$ as above, it is expected that at least one of them evaluates to -1. Hence, recognizing squares boils down to compute the kernel of:

$$\begin{aligned} \log_{-1,\chi} : U &\longrightarrow \mathbb{F}_2^N \\ a &\longmapsto \left\{\log_{-1} \chi_{\mathfrak{p}_i}(a);\ i \in [\![1, N]\!]\right\}. \end{aligned}$$

An element of this kernel is still not guaranteed to be a square. Nevertheless, a standard heuristic, first stated in the context of integer factorization [BLP93, §8] and also used in multiquadratic fields [BBV+17, §4.2], [BV18, Heur. 4.3], is to assume that if the $\mathfrak{p}_i$ are all distinct (split) prime ideals, then the $\log_{-1} \chi_{\mathfrak{p}_i}$ behave as independent uniform random elements of $\operatorname{Hom}\big(U/\big(U \cap (K_m^{\times})^2\big), \mathbb{F}_2\big)$. Concretely, this means that these should span this dual with probability at least $\big(1 - 1/2^{N-k}\big)$ [BLP93, Lem. 8.2]; in that case, any element of the kernel of $\log_{-1,\chi}$ is indeed a square. In other words, if $\sum_{1 \le i \le k} v_i \log_{1,\chi} g_i = 0$, then with high probability the product $g = \prod_{1 \le i \le k} g_i^{v_i}$ indeed belongs to $U \cap (\mathcal{O}_{K_m,\mathcal{S}}^{\times})^2$.

Square roots algorithm.

Once we have identified combinations of elements of U that are $\mathcal{S}$-unit squares, it remains to compute their square roots explicitly. First, we note that it is useful to systematically reduce those products modulo all squared circular units C_m^2 to contain the coefficients size. This is done as usual by projecting the logarithmic embedding $\operatorname{Log} g$ of the obtained $g \in (\mathcal{O}_{K_m,\mathcal{S}}^{\times})^2$ into $2 \cdot \operatorname{Log} C_m$, finding a closest vector $y = \operatorname{Log} u^2$ and replacing g by g/u^2.

The traditional method to compute the square root of an element $g \in (K_m^{\times})^2$ is to factor the polynomial $x^2 - g$ in $K_m[x]$, using e.g., Trager's method [Coh93, Alg. 3.6.4] or Belabas' p-adic method [Bel04]. As, according to Th. 5.14, we have many square roots to compute, we choose instead to use a batch strategy in the spirit of [LPS20, Alg. 5] using complex embeddings approximations.

Since LLL seminal paper [LLL82], it is known that one can retrieve an algebraic number from approximations of one of its complex embeddings. Indeed, fix an embedding $\sigma \in G_m$ and a $\mathbb{Q}$-basis $(\omega_1, \dots, \omega_n)$ of $\mathcal{O}_{K_m}$, and LLL-reduce:

$$B_\kappa := \begin{pmatrix} -\sigma(\omega_1) & C & 0 & \dots & 0 \\ -\sigma(\omega_2) & 0 & C & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ -\sigma(\omega_n) & 0 & \dots & 0 & C \end{pmatrix}.$$

where $C > 0$ is a constant and approximations are computed at precision $\kappa \in \mathbb{N}$. Then, for any $g \in \mathcal{O}_{K_m}$, applying e.g., Babai's Nearest Plane algorithm on the LLL basis of B_κ and target $\big(\sigma(g), 0, \dots, 0\big)$ gives a combination $(g_1, \dots, g_n)$ such that $g = \sum_{i=1}^{n} g_i \omega_i$. As explained in [LPS20], it is possible to mutualize the computation of B_κ and reuse the unitary transformation to hasten computations when increasing κ is required.

We use an improvement that benefits from the existence of the maximal real subfield K_m^+. Each $g \in K_m = K_m^+[\zeta_m]$ can be uniquely written as $g = g_0 + g_1 \cdot \zeta_m$, with $g_0, g_1 \in K_m^+$. For each $\sigma \in G_m^+$, the *relative Minkowski embedding* of σ w.r.t. to the extension K_m/K_m^+ is defined by $\sigma_{K_m/K_m^+}(g_0^\sigma, g_1^\sigma) = \left(g^\sigma, \overline{g^\sigma}\right) \in \mathbb{C}^2$. This is a linear homomorphism of $\mathbb{C}^2$. When $g = h^2$, its square root $h_0 + h_1\zeta_m$ can be retrieved from approximations of h_0^σ and h_1^σ instead of h^σ, as follows:

1. Compute $\sigma_{K_m/K_m^+}(g_0^\sigma, g_1^\sigma) = \left(g^\sigma, \overline{g^\sigma}\right) \in \mathbb{C}^2$;
2. Choose one complex square root z of g^σ and apply $\sigma^{-1}_{K_m/K_m^+}$ to $(z, \overline{z})$ to get potential approximations $\left(\tilde{h}_0^\sigma, \tilde{h}_1^\sigma\right)$ of h_0^σ and h_1^σ respectively;
3. Using LLL as above in K_m^+ on $\tilde{h}_0^\sigma$ and $\tilde{h}_1^\sigma$, obtain $\left(\tilde{h}_0, \tilde{h}_1\right)$ in K_m^+, which are candidates for resp. h_0 and h_1.
4. If $(\tilde{h}_0 + \tilde{h}_1 \cdot \zeta_m)^2 \neq g$, then increase κ using the fast method of [LPS20].

Hence, this method amounts to LLL reducing a matrix of size $\frac{n}{2} \times (\frac{n}{2} + 1)$ and decoding using e.g., Babai's Nearest Plane algorithm. This offers a great speed-up compared to reducing a matrix of size $n \times (n+1)$. For further details and generalizations to higher order polynomial roots, we refer the interested reader to [Les21].

Rebuilding a basis.

After the square root step, we obtain new elements $h_1, \ldots, h_r$, where $r = \dim\left(\ker \log_{-1,\chi}\right)$. In order to extract a set of k independent elements from the extended set $\left\{h_1, \ldots, h_r, g_1, \ldots, g_k\right\}$, we compute an LLL-basis of the matrix constituted of their valuations at the places of $\mathcal{S}$. Note that this matrix can be computed entirely from the valuations of the initial set $\{g_i\}$ and the basis of $\ker \log_{-1,\chi}$. Using the same trick as for matrix A in [BBV+17, Alg. 5.2], this contains the height of the transformation matrix, sufficiently for our needs.

At the end of this process we obtain a maximal set of independent $\mathcal{S}$-units of index given by Th. 5.14 where no factor 2 remains.

5.3 Removing Quantum Steps from the CDW Algorithm

The complete material for this section is given in [BLNR21, §B], and the main points are briefly summarized here. The CDW algorithm for solving Approx-Svp was introduced in [CDW17] for cyclotomic fields of prime power conductors, and extended to all conductors in [CDW21]. Its main feature is the use of some short relations of the Stickelberger ideal.

In this section, we show how to benefit from the results of §5.2.1 and §5.2.2 to remove most quantum steps of [CDW21]. More precisely, we first propose in [BLNR21, §B.2] an equivalent rewriting of [CDW21, Alg. 7] that enlightens some hidden steps that reveal useful for subsequent modifications. Then, in [BLNR21, §B.3], we plug the explicit generators of §5.2.1 and Eq. (5.9), for relative class group orbits, to remove the last call to the quantum Pip solver. Finally, by considering the module of *all* real class group relations, using Pr. 5.17, we remove in [BLNR21, §B.4] the need of a random walk mapping any ideal of K_m into Cl_m^-, at the (small) price of restricting to cyclotomic fields such that $h_m^+ \leq O(\sqrt{m})$ (Hyp. 5.18).

An equivalent rewriting of CDW ([BLNR21, §B.2]).

Omitting details, the CDW algorithm works as follows, for any challenge ideal $\mathfrak{a}$ of K_m [CDW21, Alg. 7]:

1. Random walk to Cl_m^-: find $\mathfrak{b}$ such that $\left[\mathfrak{a}\mathfrak{b}\right] \in \mathrm{Cl}_m^-$.

2. Solve the ClDLP of $\mathfrak{ab}$ on G_m-orbits of the prime ideals $\mathfrak{L}_1, \ldots, \mathfrak{L}_d$ of Cl_m^-. This gives a vector $\alpha = (\alpha_1, \ldots, \alpha_d) \in \mathbb{Z}[G_m]^d$ such that $\mathfrak{ab} \cdot \prod_i \mathfrak{L}_i^{\alpha_i}$ is principal.
3. Solve the CPMP by projecting each α_i in $\pi(\mathcal{S}_m) = (1-\tau)\mathcal{S}_m$, find a vector $v_i = y_i \cdot \pi(\mathcal{S}_m)$ close to $\pi(\alpha_i)$, and then lift v_i to get some β_i s.t. $\pi(\beta_i) = v_i$, $\|\alpha - \beta\|_1$ is small *with positive coordinates* and $\mathfrak{ab} \cdot \prod_i \mathfrak{L}_i^{\alpha_i - \beta_i}$ is principal.
4. Apply the PIP algorithm of [BS16] to get a generator of this principal ideal.
5. Reduce the obtained generator by circular units like in [CDPR16].

This eventually outputs $h \in \mathfrak{a}$ of length $\|h\|_2 \leq \exp\big(\widetilde{O}(\sqrt{m})\big) \cdot \mathcal{N}(\mathfrak{a})^{1/\varphi(m)}$. [CDW21, Th. 5.1].

We focus on the lift procedure of Step 3. In [CDW21], a vector $v \in \pi(\mathcal{S}_m)$ is lifted to β by keeping positive coordinates for β_σ and sending opposite of negative coordinates to $\beta_{\tau\sigma}$. This works because for any $\mathfrak{c} \in \mathrm{Cl}_m^-$, $[\mathfrak{c}]^{-1} = [\mathfrak{c}^\tau]$, but hides which exact product of relative norm ideals is involved.

We propose a totally equivalent lift procedure: from $v = y \cdot \pi(\mathcal{S}_m)$, consider the preimage vector $\tilde{\beta} = y \cdot \mathcal{S}_m$, from which we remove $\min\{\tilde{\beta}_\sigma, \tilde{\beta}_{\tau\sigma}\}$ to each $\tilde{\beta}_\sigma$ coordinate to obtain β. Now, it is obvious that β is a combination y of relations in $\mathcal{S}_m$, and of relative norm relations given by the min part. Details are given in [BLNR21, Alg. B.6].

Using explicit Stickelberger generators ([BLNR21, §B.3]).

Each element w_a of the generating set W of $\mathcal{S}_m$ corresponds to a generator $\mathcal{J}_{\mathfrak{L}}(1, a-1)$ (see §5.2.1). Similarly, each relative norm ideal writes $\langle \gamma_s^+ \rangle = \mathfrak{L}^{(1+\tau)\sigma_s}$ (see §5.2.2). Hence, from an (explicit) ClDLP solution $\langle g \rangle = \mathfrak{ab} \cdot \mathfrak{L}^\alpha$, and given, as rewritten above, a CPMP solution as $\beta = y \cdot W + u \cdot (1+\tau) \cdot \mathbb{Z}[G_m^+]$, we have that a generator of $\mathfrak{ab} \cdot \mathfrak{L}^{\alpha-\beta}$ is directly given by $g / \big(\prod_a \mathcal{J}_{\mathfrak{L}}(1, a-1)^{y_a} \prod_s (\gamma_s^+)^{u_s}\big)$. Knowing this allows us to remove the quantum PIP in dimension n in step 4 (for each query). In exchange, we need to compute (only once) all real generators for relative norm relations, which can be done in dimension $\varphi(m)/2$ by [BS16, Alg. 2].

Avoiding the random walk ([BLNR21, §B.4]).

Finally, note that several quantum steps are performed (for each query) in the random walk that maps ideals to Cl_m^-. Using the results of §5.2.2, we replace the module $(1+\tau) \cdot \mathbb{Z}[G_m]^d$ by the module of all real class group relations.

Asymptotically, we prove in [BLNR21, Pr. B.7] as a direct consequence of Pr. 5.12 that this does not change the bound on the approximation factor, as long as:

Hypothesis 5.18. *We restrict to cyclotomic fields K_m verifying $h_m^+ \leq O(\sqrt{m})$.*

Remark 5.19. This assumption is certainly not true in general. Nevertheless, by the discussion in §2.2.3, it is expected to be valid for a very large proportion of cyclotomic fields and is likely to hold when m is a power of 2.

On the other hand, this slightly more restrictive hypothesis is largely compensated by the fact that only two quantum steps remain: one is performed only once in dimension $\varphi(m)/2$ to compute real class group relations and generators, and the second is solving the ClDLP for each query (see [BLNR21, Tab. B.1]). Moreover, this removes the need for the factor base prime ideals $\mathfrak{L}_i$ to be in the relative class group Cl_m^-, which happens with probability only roughly $1/h_m^+$. Therefore, we can choose a factor base of prime ideals having the smallest possible norms, which has in practice a significant impact on the algebraic norm of these $\mathfrak{L}_i$'s (see also Wesolowski's bound on $\mathcal{N}(\mathfrak{L}_{\mathsf{max}}^-)$ in §2.3.1), and thus on the final approximation factor reached by the CDW algorithm in [CDW21, Alg. 6].

5.4 Computing Log-$\mathcal{S}$-unit Sublattices in Higher Dimensions

Our main goal is to simulate the Twisted-PHS algorithm for high degree cyclotomic fields. To this end, we compute full-rank sublattices of the full log-$\mathcal{S}$-unit lattice using the knowledge of the maximal set $\mathfrak{F}$ of independent $\mathcal{S}$-units defined by Eq. (5.13) and its 2-saturated counterpart $\mathfrak{F}_{\mathsf{sat}}$ from §5.2.4. These sets are lifted from a complete set of real $\mathcal{S}^+$-units (see §5.2.2), hence are obtained at the classically subexponential cost of working in the half degree maximal real subfield. We note that by Th. 5.14, the index of these families grows rapidly as the number of orbits increases, hence these approximated modes give an upper bound on the approximation factors that can be expected when using Twisted-PHS.

Our experimental setting is detailed in §5.4.1. Then, we analyse in §5.4.2 the geometric characteristics of our log-$\mathcal{S}$-unit sublattices and the obtained approximation factors in §5.4.3.

5.4.1 Experimental settings

Computing the full group of $\mathcal{S}$-units in a classical way is rapidly intractable, even in the case of cyclotomic fields; therefore, the experiments performed on Twisted-PHS in §3.4 were bound to $\varphi(m) \leq 70$. We apply the Twisted-PHS algorithm using our full-rank sublattices of the whole log-$\mathcal{S}$-unit lattice induced by the independent family $\mathfrak{F}$ of Eq. (5.13), its 2-saturated counterpart $\mathfrak{F}_{\mathsf{sat}}$ (§5.2.4) and, when possible, a fundamental system $\mathfrak{F}_{\mathsf{su}}$ for the full $\mathcal{S}$-unit group. Approximated modes with $\mathfrak{F}$ or $\mathfrak{F}_{\mathsf{sat}}$ give a glimpse on how Twisted-PHS scales in higher dimensions, where asymptotic phenomena like the growth of h_m start to express.

Source code and hardware description. All experiments have been implemented using SageMath v9.0 [Sag20], except for the full $\mathcal{S}$-unit groups computations for which we used Magma [BCP97], which appears much faster for this particular task and also offers an indispensable product ("Raw") representation. Moreover, fplll [FpL16] was used to perform all lattice reduction algorithms. The entire source code is provided on GitHub: ob3rnard/Tw-Sti[8].

Most of the computations were performed in less than two weeks on a server with 72 Intel® Xeon® E5-2695v4 @2.1GHz with 768GB of RAM, using 2TB of storage for the precomputations. Real class group computations were performed on a single Intel® Core™ i7-8650U @3.2GHz CPU using 10GB of RAM.

Targetted cyclotomic fields. We consider cyclotomic fields of *any* conductor m s.t. $20 < \varphi(m) \leq 210$ with known real class number $h_m^+ = 1$, including those from Tab. 2.1. The restriction to $h_m^+ = 1$ is only due to technical interface obstructions, i.e., we are not aware of how to access the non-trivial real class group relations internally computed by SageMath. Additionally, for some of the conductors, we were not able to obtain the real class group within a day. Thus, we are left with 210 distinct cyclotomics fields, and Tab. 5.1 lists all ignored conductors.

Finite places choice. The optimal set of places computed by Alg. 3.3 yields a number d_{max} of split G_m-orbits of smallest norms maximizing the density of the corresponding full log-$\mathcal{S}$-unit lattice. However, the index of our log-$\mathcal{S}$-unit *sublattices*, given by Th. 5.14, grows too quickly, roughly in $(h_m^-)^{d-1}$, so that their density always decreases as soon as $d > 1$. This remark motivates us to compute all log-$\mathcal{S}$-unit sublattices for $d = 1$ to d_{max} first split G_m-orbits.

[8] https://github.com/ob3rnard/Tw-Sti

m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+
136	64	2	408	128	2	205	160	2	356	176	†	520	192	4	265	208	†
212	104	5	268	132	†	328	160	†	376	184	†	840	192	†	424	208	†
145	112	2	284	140	†	440	160	5	191	190	11	303	200	†	636	208	†
183	120	4	292	144	†	163	162	4	221	192	†	404	200	†			
248	120	4	504	144	4	332	164	†	388	192	†	309	204	†			
272	128	2	316	156	†	344	168	†	476	192	†	412	204	†			

TABLE 5.1 – List of ignored conductors (†: failure to compute Cl_m^+ within a day).

Full rank log-$\mathcal{S}$-unit sublattices. The first maximal set of independent $\mathcal{S}$-units that we consider is $\mathfrak{F}$ from Eq. (5.13). The 2-saturation process of §5.2.4 mitigates the huge index of $\mathfrak{F}$, yielding family $\mathfrak{F}_{\mathsf{sat}}$. A fundamental system $\mathfrak{F}_{\mathsf{su}}$ of the full $\mathcal{S}$-unit group $\mathcal{O}_{K_m,\mathcal{S}}^{\times}$ (modulo torsion) is also used whenever it is computable in reasonable time, i.e., up to $\varphi(m) < 80$. As noted in §2.1.3, their images under any log-$\mathcal{S}$-embedding φ form full-rank sublattices resp. L_{urs}, L_{sat}, L_{su}, generated by resp. $\varphi(\mathfrak{F})$, $\varphi(\mathfrak{F}_{\mathsf{sat}})$, $\varphi(\mathfrak{F}_{\mathsf{su}})$, of the corresponding full log-$\mathcal{S}$-unit lattice $\varphi(\mathcal{O}_{K_m,\mathcal{S}}^{\times})$.

We consider several choices of the log-$\mathcal{S}$-embedding φ. Namely, we tried to evaluate the advantage of using the expanded $\overline{\mathrm{Log}}_{\mathcal{S}}$ (exp) over $\mathrm{Log}_{\mathcal{S}}$, labelled tw (as twisted by $[\mathbb{C}:\mathbb{R}]=2$). We also considered versions with (iso) or without (noiso) the isometry f_H of Eq. (3.16). This yields four choices for φ, e.g., tag noiso/tw is $\varphi = \mathrm{Log}_{\mathcal{S}}$ and iso/exp gives the original $\varphi_{\mathsf{tw}} = f_H \circ \overline{\mathrm{Log}}_{\mathcal{S}}$.

Compact product representation. In order to avoid the exponential growth of algebraic integers viewed in $\mathbb{Z}[x]/\langle\Phi_m(x)\rangle$, we use a compact product representation, so that any element α in $\mathfrak{F}$ (resp. $\mathfrak{F}_{\mathsf{sat}}$ or $\mathfrak{F}_{\mathsf{su}}$) is written on a set $g_1,\ldots,g_N$ of N small elements as $\alpha = \prod_{i=1}^{N} g_i^{e_i}$. Hence, besides the g_i's, each α is stored as a vector $e \in \mathbb{Z}^N$, and for any choice of φ, we have that $\varphi(\alpha) = \sum_{i=1}^{N} e_i \cdot \varphi(g_i)$. This allows us to compute φ without the coefficient explosion encountered in §3.4, which unlocks the full log-$\mathcal{S}$-unit lattices computations beyond degree 60.

Lattice reductions. For each of the constructed log-$\mathcal{S}$-unit sublattices, i.e., for each number of orbits $d \in [\![1, d_{\mathsf{max}}]\!]$, for each family of independent $\mathcal{S}$-units $\mathfrak{F}$, $\mathfrak{F}_{\mathsf{sat}}$ and (when available) $\mathfrak{F}_{\mathsf{su}}$, and for each choice of log-$\mathcal{S}$-embedding, we compare several levels of reduction: no reduction ("raw"), LLL-reduction and BKZ_{40}-reduction.

5.4.2 Geometry of the lattices

For all described choices of log-$\mathcal{S}$-unit sublattices, we first evaluate several geometrical parameters (see §2.4.3): reduced volume $V^{1/k}$, root-Hermite factor δ_0, orthogonality defect δ. For clarity's sake, we only give here a few examples giving a glimpse of what happens in general, and additional data can be found in §5.5.1.

Table 5.2 contains data for cyclotomic fields $\mathbb{Q}(\zeta_{152})$ and $\mathbb{Q}(\zeta_{211})$ of degrees resp. 72 and 210. All values correspond to the iso/exp log-$\mathcal{S}$-embedding, i.e., $\varphi = \varphi_{\mathsf{tw}}$. Indeed, as illustrated by Tab. 5.4, we experimentally note that using (no)iso/exp seems geometrically slightly better than using (no)iso/tw. Notice how small is the normalized orthogonality defect after only LLL reduction, unambiguously below the tight Minkowski bound $\sqrt{1+\frac{k}{4}}$ given in §2.4.3.

We then look at the logarithm of the Gram-Schmidt norms, for every described choice of log-$\mathcal{S}$-unit sublattices. Figure 5.2 plots the Gram-Schmidt log norms before and after BKZ reduction of the lattices L_{sat}, using the original iso/exp log-$\mathcal{S}$-embedding φ_{tw}. As in Fig. 3.6–3.15, for each field the two curves are almost superposed. We also checked the impact of the

m	d	set	k	$\mathrm{Vol}^{1/k}$	δ			$\max_{1\le i\le k}\|\mathbf{b}_i\|_2$		
					raw	LLL	bkz_{40}	raw	LLL	bkz_{40}
152	1	urs	107	8.691	2.016	1.570	1.551	45.007	38.466	38.202
		sat	107	6.928	4.398	1.787	1.822	752.306	23.280	21.720
		su	107	6.928	28.396	1.805	1.828	3163.723	21.953	21.446
	2	urs	179	9.683	2.157	1.623	1.590	48.754	41.313	41.404
		sat	179	7.384	7.670	1.885	1.896	6273.562	23.280	22.772
		su	179	6.816	65.355	2.226	2.322	3427.134	23.221	24.741
211	1	urs	314	14.325	2.672	2.291	2.257	96.068	97.930	96.569
		sat	314	11.386	9.998	2.581	2.562	9742.552	59.387	59.578
	5	urs	1154	18.232	3.118	2.542	2.497	118.124	119.160	115.888
		sat	1154	13.341	19.443	2.918	2.901	32067.612	71.428	72.752
	7	urs	1574	18.976	3.161	2.557	2.512	120.838	121.129	119.020
		sat	1574	13.771	26.841	2.927	2.910	530646.708	71.428	72.752

Table 5.2 – Geometric characteristics of L_{urs}, L_{sat} and L_{su} for $\mathbb{Q}(\zeta_{152})$ and $\mathbb{Q}(\zeta_{211})$ using log-$\mathcal{S}$-embedding φ_{tw} (of type $\mathsf{iso}/\mathsf{exp}$). For *all* bases, the root-Hermite factor verifies $|\delta_0 - 1| < 10^{-3}$.

log-$\mathcal{S}$-embedding choice among all four options on the Gram-Schmidt logarithm norms of the *unreduced* basis $\varphi(\mathfrak{F}_{\mathsf{sat}})$. As expected, the isometry f_H has absolutely no influence on the Gram-Schmidt norms. On the other hand, using $\mathrm{Log}_{\mathcal{S}}$ or $\overline{\mathrm{Log}}_{\mathcal{S}}$ seems to alter only the first norms, very slightly, as can be seen in Fig. 5.8. Again, increasing the number of orbits does not influence these behaviours.

We stress that these very peculiar geometric characteristics — shape of the logarithm of the norms of the Gram-Schmidt basis, ease of reduction, very small orthogonality defect (after LLL) — already observed in §§3.4.1 and 3.4.2, are consistently viewed across all conductors, degrees, log-$\mathcal{S}$-unit sublattices and number of orbits. To give a concrete idea of e.g., the striking ease of reduction of these log-$\mathcal{S}$-unit sublattices, we report that for $m = 211$, BKZ_{40} terminates in around 7 minutes (resp. 30 minutes) on the log-$\mathcal{S}$-unit sublattice of dimension $k = 1154$ (resp. 1574) corresponding to $d = 5$ (resp. $d_{\mathsf{max}} = 7$), which is unusually fast.

This very broad phenomenon suggests that the explanation is possibly deep, an observation that has been recently developed by Bernstein and Lange [BL21].

5.4.3 Evaluation of the approximation factor

In §3.4.3, evaluating in practice the approximation factors reached by Twisted-PHS is done by choosing random split ideals of prime norm, solving the ClDlp for these challenges and comparing the length of the obtained algebraic integer with the length of the exact shortest element. As the degrees of the fields grow, solving the ClDlp and exact id-Svp becomes rapidly intractable. Hence, we resort to simulating random outputs of the ClDlp, similarly to [DPW19, Hyp. 8], and estimate the obtained approximation factors with inequalities from Eq. (2.42).

Simulation of ClDLP solutions.

To simulate targets that heuristically correspond to the output α of the ClDlp, we assume that for each ideal $\mathfrak{L}_{\mathfrak{i}} \in \mathcal{S}$, the vector $\left(v_{\mathfrak{L}_{\mathfrak{i}}^{\sigma}}(\alpha)\right)_{\sigma\in G_m}$ of $\mathbb{Z}[G_m]$ is uniform modulo the lattice of class relations, and that after projection along the $\mathbf{1}$-axis, $\left(\ln|\sigma(\alpha)|\right)_{\sigma}$ is uniform modulo the log-unit

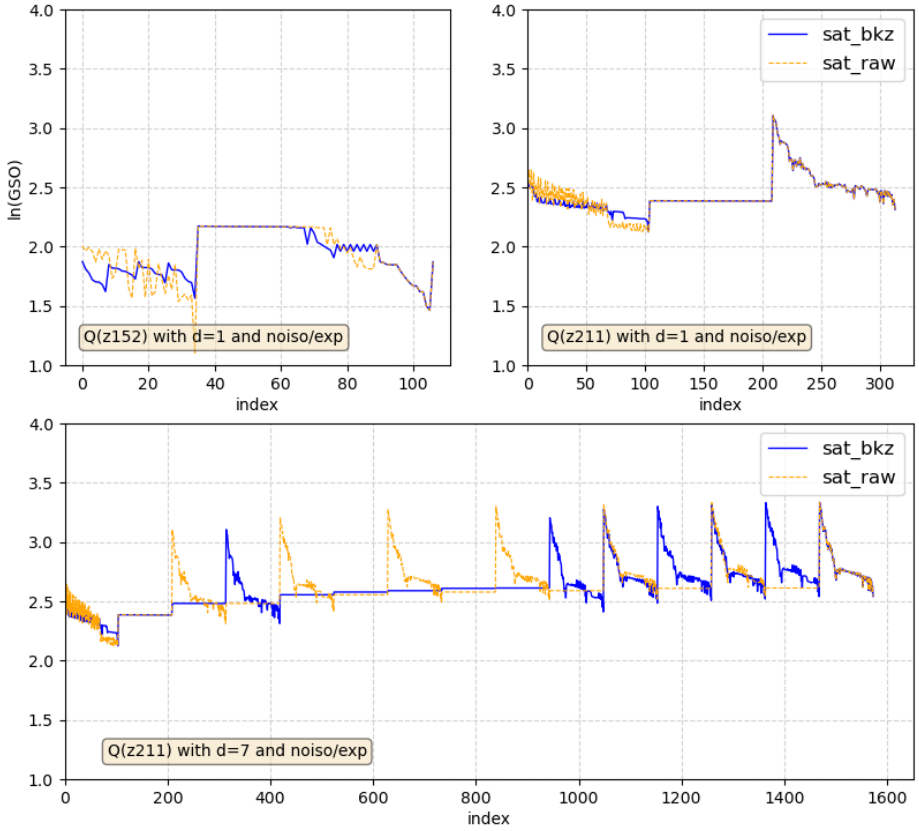

FIGURE 5.2 – L_{sat} lattices for $\mathbb{Q}(\zeta_{152})$ and $\mathbb{Q}(\zeta_{211})$: Gram-Schmidt log norms before and after reduction by BKZ_{40}.

lattice. These hypotheses have already been used in [DPW19, Hyp. 8] or Heur. 3.28, and are backed up by theoretical results in [BDPW20, Th. 3.3].

Drawing random elements modulo a lattice of rank k is done by following a Gaussian distribution of sufficiently large deviation. Concretely, we first choose a random split prime p in the range $[\![2^{97}, 2^{103}]\!]$. Then, for each $\mathfrak{L} \in \mathcal{S} \cap \mathcal{S}_0$, we pick random valuations $v_{\mathfrak{L}}(\alpha)$ modulo the lattice of class relations of rank $\sharp(\mathcal{S} \cap \mathcal{S}_0)$ and random elements $(u_\sigma)_{\sigma \in G_m^+} \in \mathbb{R}^{\varphi(m)/2}$ in the span of the log-unit lattice of rank $\frac{\varphi(m)}{2} - 1$. Finally, we simulate $(\ln|\sigma(\alpha)|)_\sigma$ by adding $\frac{\ln p + \sum_{\mathfrak{L} \in \mathcal{S}} v_{\mathfrak{L}} \ln \mathcal{N}(\mathfrak{L})}{\varphi(m)}$ to each coordinate u_σ, so that their sum is indeed $\frac{\ln|\mathcal{N}(\alpha)|}{2}$. For each field we thereby generate 100 random targets on which to test Twisted-PHS on all lattice versions.

Reconstruction of a solution.

For any simulated ClDLP solution α, given as a random vector $\left(\{\ln|\sigma(\alpha)|\}_{\sigma \in G_m^+}, \{v_{\mathfrak{L}}(\alpha)\}_{\mathfrak{L} \in \mathcal{S} \cap \mathcal{S}_0}\right)$, it is easy to compute $\varphi(\alpha)$ for any log-$\mathcal{S}$-embedding φ and to derive a target as in Eq. (3.25), including a drift parameterized by some β. Then, considering e.g., $L_{\mathsf{sat}} = \varphi(\mathfrak{F}_{\mathsf{sat}})$, given by the *BKZ*$_{40}$*-reduced* basis $U_{\mathsf{bkz}} \cdot \varphi(\mathfrak{F}_{\mathsf{sat}})$, we find a close vector $v = (y \cdot U_{\mathsf{bkz}}) \cdot \varphi(\mathfrak{F}_{\mathsf{sat}})$ to this target using Babai's Nearest Plane algorithm, and from y, U_{bkz} and $\mathfrak{F}_{\mathsf{sat}}$ we easily recover, in compact representation, $s \in \mathcal{O}_{K_m,\mathcal{S}}^\times$ s.t. $v = \varphi(s)$ and also α/s.

The purpose of the drift parameter β is to guarantee $v_{\mathfrak{L}}(\alpha/s) \geq 0$ on all finite places. As

mentioned in §3.4.3, the length of α/s is extremely sensitive to the value of β, so that they searched for an optimal value by dichotomy. However, this positiveness property actually does not seem to be monotonic in β, and in practice, using the same β on each finite place coordinate is too coarse when the dimension grows, which induces unnecessarily large approximation factors. We instead obtained best results using random drifts in ℓ_∞-norm balls of radius 1 centered on the $\mathbf{1}$ axis. A first sampling of $O(\varphi(m))$ random points $\beta \cdot \mathbf{1} + \mathcal{B}_\infty(1)$ for a wide range of random β's allows us to select a β_0 around which we found the best $\|\alpha/s\|_2$ with all $v_{\mathfrak{L}}(\alpha/s)$ being positive. Then we sample $O(\varphi(m))$ uniform random points in the neighbourhood of β_0, namely in $[0.9\beta_0, 1.1\beta_0] \cdot \mathbf{1} + \mathcal{B}_\infty(1)$, and output the overall optimal $\|\alpha/s\|_2$ having all $v_{\mathfrak{L}}(\alpha/s) \geq 0$.

Estimator of the approximation factor.

Since we do not have access to the shortest element of a challenge ideal, we cannot compute an exact approximation factor as is done in §3.4.3. Instead, we estimate the retrieved approximation factor using the inequalities implied by Eq. (2.42). We focus on the Gaussian Heuristic, which gives in small dimensions consistent results with the exact approximation factors found in §3.4.3. For each cyclotomic field, the plotted points are the means, over the 100 simulated random targets, of the minimal approximation factors obtained using options iso/noiso and exp/tw. For each family $\mathfrak{F}$, $\mathfrak{F}_{\mathsf{sat}}$ and $\mathfrak{F}_{\mathsf{su}}$, we chose to keep only the factor base that gives the best result. This systematically translated into using $d = 1$ G_m-orbit for $\mathfrak{F}$ and $\mathfrak{F}_{\mathsf{sat}}$, whereas we had to use $d = d_{\mathsf{max}}$ for $\mathfrak{F}_{\mathsf{su}}$, as predicted by the Twisted-PHS algorithm.

Figure 5.3 shows the approximation factor γ_{gh} obtained for all lattices L_{urs}, L_{sat} and L_{su} (when applicable) after BKZ_{40} reduction. Figure 5.4 is a zoom of Fig. 5.3 that focuses on L_{sat} and L_{su} on small dimensions.

First, we remark that using family $\mathfrak{F}$ from Eq. (5.13), the retrieved approximation factors are increasing rapidly. Using the 2-saturated family $\mathfrak{F}_{\mathsf{sat}}$ yields much better results, and looking closely at Fig. 5.4 shows that using a basis $\mathfrak{F}_{\mathsf{su}}$ of the full $\mathcal{S}$-unit group, when available, even improves the picture if $d_{\mathsf{max}} > 1$, in which case L_{su} is denser than L_{sat}. For L_{su}, we stress that we obtain estimated approximation factors very similar to the exact ones observed in §3.4.3.

More generally, we observe a very strong correlation between the density of our lattices and the obtained approximation factors — the denser, the better. As an important related remark, the variance seen for γ_{gh} in Fig. 5.3 for distinct fields of same degree follows the variations of the norm of the first split prime, thus of the reduced volume of the considered log-$\mathcal{S}$-unit sublattice. We expect this variance to be smoothed through conductors for the full log-$\mathcal{S}$-unit lattice.

Furthermore, considering $m = 211$, the $\mathfrak{F}$ family gives $\mathrm{Vol}^{1/314} L_{\mathsf{urs}} \approx 14.325$ and an estimated $\gamma_{\mathsf{gh}} \approx 13170$, for $\mathfrak{F}_{\mathsf{sat}}$ we get $\mathrm{Vol}^{1/314} L_{\mathsf{sat}} \approx 11.386$ and a much smaller estimated $\gamma_{\mathsf{gh}} \approx 16.4$, whereas the optimal number of orbits predicted by the Twisted-PHS Factor Base Choice Algorithm (Alg. 3.3) is $d_{\mathsf{max}} = 7$, which yields a full log-$\mathcal{S}$-unit lattice of reduced volume only $\mathrm{Vol}^{1/1574} L_{\mathsf{su}} \approx 9.635$.

Comparison to the CDW algorithm.

Using the same experimental setting, we compute the approximation factors obtained using the CDW algorithm as implemented in [DPW19] ("Naive version") with additional BKZ_{40} lattice reductions, as well as the experimentally derived *volumetric lower bound* from [DPW19, Eq. (5) and Tab. 1]. Those values are also represented in Fig. 5.3 and 5.4.

Our experimental results using the $\mathfrak{F}_{\mathsf{sat}}$ family clearly outperform the CDW algorithm over the experimental range considered, and are even comparable to its volumetric lower bound. Moreover, for some fields, e.g., in dimensions 96, 160, 168, 200, this lower bound is defeated by the (approximated version of the) Twisted-PHS algorithm. Note that this does not invalidate the

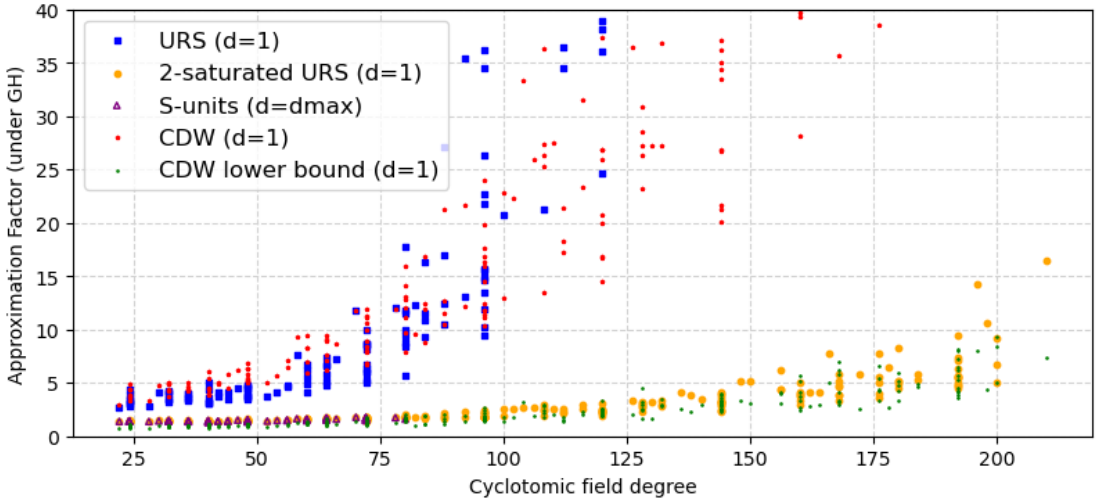

FIGURE 5.3 – Approximation factors, with Gaussian Heuristic, reached by Twisted-PHS for cyclotomic fields of degree up to 210, on lattices L_{urs}, L_{sat} and L_{su}.

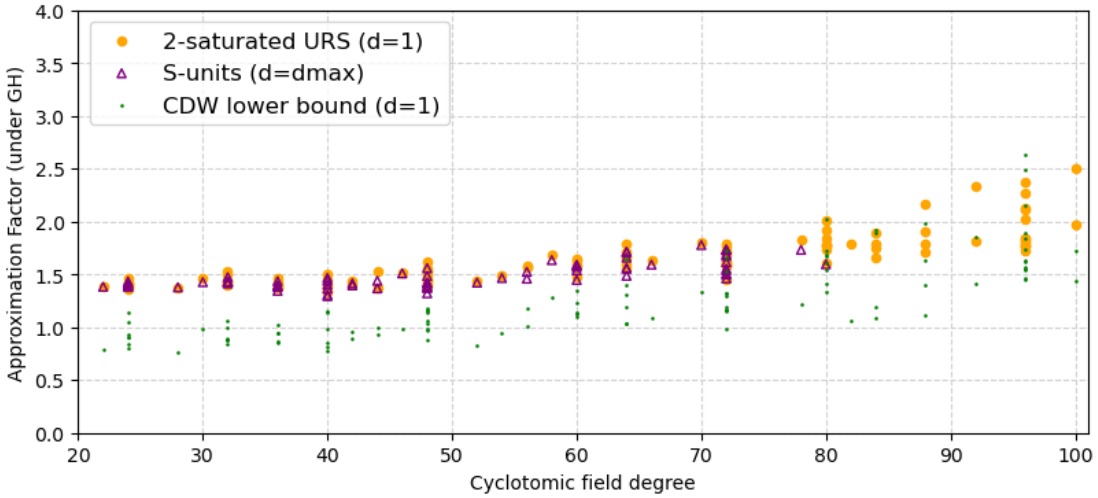

FIGURE 5.4 – Approximation factors, with Gaussian Heuristic, reached by Twisted-PHS for cyclotomic fields of degree up to 100, on lattices L_{sat} and L_{su}.

lower bound itself, which is stated for the two-phase CDW algorithm, but indicates the power of combining both steps in only one lattice as in the Twisted-PHS algorithm.

5.5 Supplementary Experimental Results

5.5.1 Geometry of log-$\mathcal{S}$-unit sublattices

In the following, we provide data regarding the geometry of the log-$\mathcal{S}$-unit sublattices L_{urs} and L_{sat} for additional cyclotomic fields.

m	d	set	k	$\mathrm{Vol}^{1/k}$	δ			$\max_{1\le i\le k}\|\mathbf{b}_i\|_2$		
					raw	LLL	bkz_{40}	raw	LLL	bkz_{40}
159	1	urs	155	11.291	2.177	1.702	1.686	71.228	62.253	60.096
		sat	155	8.989	6.143	1.898	1.921	3168.773	35.391	35.703
	2	urs	259	12.576	2.350	1.781	1.739	72.069	62.357	60.675
		sat	259	9.572	6.902	2.028	2.036	3168.773	36.062	35.703
	3	urs	363	13.364	2.419	1.798	1.750	75.913	65.973	63.701
		sat	363	9.978	7.602	2.066	2.066	3168.773	37.480	37.132
149	1	urs	221	12.192	2.828	2.091	1.999	74.637	71.073	68.291
		sat	221	9.697	12.473	2.305	2.244	12554.466	44.327	44.326
	2	urs	369	13.353	3.134	2.233	2.149	78.906	74.039	71.298
		sat	369	10.150	14.472	2.507	2.467	12554.466	47.719	46.438
	3	urs	517	13.962	3.269	2.271	2.190	80.529	76.289	76.007
		sat	517	10.410	22.211	2.569	2.531	85211.593	47.719	48.556
	4	urs	665	14.415	3.327	2.300	2.223	83.176	78.268	77.926
		sat	665	10.632	20.731	2.606	2.576	85211.593	47.768	48.556
516	1	urs	251	11.815	2.535	2.026	2.013	77.904	73.051	72.993
		sat	251	9.395	6.508	2.341	2.359	4850.233	44.290	43.783
	2	urs	419	12.921	2.833	2.156	2.129	82.452	76.629	75.586
		sat	419	9.818	8.208	2.550	2.565	5761.443	46.559	46.426
	3	urs	587	13.850	2.945	2.202	2.167	91.958	84.961	86.487
		sat	587	10.321	10.348	2.620	2.623	9544.834	49.096	49.971
	4	urs	755	14.445	2.998	2.222	2.188	93.457	86.198	87.794
		sat	755	10.650	12.682	2.652	2.652	26820.239	54.045	52.543
181	1	urs	269	12.855	2.747	2.308	2.146	81.230	79.924	79.204
		sat	269	10.220	7.486	2.537	2.499	5185.677	49.694	48.264
	2	urs	449	14.033	2.958	2.456	2.268	87.161	85.755	84.008
		sat	449	10.661	9.849	2.736	2.706	5185.677	50.406	51.466
	3	urs	629	14.823	3.064	2.508	2.311	92.620	90.665	88.578
		sat	629	11.045	12.340	2.801	2.778	9957.084	52.207	51.880
	4	urs	809	15.330	3.096	2.529	2.330	93.988	91.158	89.982
		sat	809	11.300	12.307	2.829	2.814	9957.084	53.598	53.519
209	1	urs	269	10.796	2.678	2.239	2.238	70.154	70.428	68.371
		sat	269	8.583	8.273	2.599	2.609	8920.663	42.887	42.683
	2	urs	449	12.651	2.921	2.320	2.300	92.739	89.996	88.251
		sat	449	9.612	14.860	2.729	2.722	45374.160	53.927	53.643
217	1	urs	269	12.110	2.608	2.137	2.115	83.336	76.670	76.186
		sat	269	9.629	6.814	2.420	2.410	4415.772	47.546	46.464
	2	urs	449	13.741	2.857	2.270	2.251	96.095	87.194	87.023
		sat	449	10.440	10.474	2.630	2.623	14735.404	56.381	56.328
	3	urs	629	14.646	2.941	2.319	2.313	99.437	89.912	93.209
		sat	629	10.913	11.667	2.696	2.696	14735.404	56.381	57.135
279	1	urs	269	12.059	2.573	2.080	2.064	81.546	76.724	84.960
		sat	269	9.588	11.575	2.391	2.397	12586.042	51.509	50.663

m	d	set	k	$\mathrm{Vol}^{1/k}$	δ			$\max_{1\le i\le k}\|\mathbf{b}_i\|_2$		
					raw	LLL	bkz_{40}	raw	LLL	bkz_{40}
279	2	urs	449	13.528	2.836	2.212	2.195	92.187	86.744	96.124
		sat	449	10.278	12.899	2.603	2.604	12586.042	57.098	57.696
	3	urs	629	14.378	2.965	2.263	2.250	96.095	89.520	96.124
		sat	629	10.713	16.966	2.677	2.683	25638.489	57.098	57.696
	4	urs	809	14.971	3.010	2.285	2.268	99.014	92.948	99.817
		sat	809	11.036	17.733	2.709	2.713	25638.489	58.977	58.807
	5	urs	989	15.396	3.053	2.302	2.280	100.238	93.692	99.817
		sat	989	11.271	18.878	2.729	2.731	26995.083	61.123	59.322
297	1	urs	269	12.331	3.169	2.074	2.005	86.980	81.006	81.451
		sat	269	9.804	21.668	2.308	2.319	94056.513	48.941	48.984
	2	urs	449	13.513	3.676	2.252	2.148	90.321	83.985	85.236
		sat	449	10.266	36.211	2.540	2.546	94056.513	50.795	51.447
	3	urs	629	14.165	3.895	2.327	2.196	92.913	86.090	85.236
		sat	629	10.555	37.241	2.645	2.640	94056.513	51.969	51.524
	4	urs	809	14.674	4.007	2.356	2.224	96.821	89.321	87.488
		sat	809	10.816	40.952	2.688	2.685	94056.513	52.120	53.167
235	1	urs	275	11.873	2.631	2.183	2.132	80.433	77.904	79.127
		sat	275	9.439	7.618	2.479	2.470	5297.502	47.586	46.684
	2	urs	459	13.287	2.936	2.347	2.275	91.190	87.506	82.926
		sat	459	10.094	12.645	2.706	2.699	28003.197	51.044	51.229
	3	urs	643	14.178	3.061	2.398	2.328	96.709	91.765	91.485
		sat	643	10.563	13.258	2.780	2.772	28003.197	52.348	52.334
	4	urs	827	14.743	3.099	2.423	2.349	98.093	93.292	92.979
		sat	827	10.867	13.861	2.815	2.807	28003.197	55.931	54.179
564	1	urs	275	12.264	2.551	2.035	2.061	82.573	77.166	76.021
		sat	275	9.750	14.624	2.390	2.370	39653.048	46.848	46.757
	2	urs	459	13.384	2.831	2.193	2.230	87.333	81.561	80.426
		sat	459	10.168	15.707	2.655	2.637	39653.048	50.285	49.290
	3	urs	643	14.393	2.984	2.240	2.274	98.851	90.926	90.825
		sat	643	10.724	17.342	2.727	2.714	39653.048	53.003	53.868
	4	urs	827	15.032	3.029	2.256	2.292	100.234	91.997	92.037
		sat	827	11.080	18.829	2.757	2.744	39653.048	55.358	55.921

TABLE 5.3 – Geometric characteristics of L_{urs}, L_{sat} and L_{su} for some cyclotomic fields using log-$\mathcal{S}$-embedding φ_{tw} (of type $\mathsf{iso/exp}$). For *all* bases, the root-Hermite factor verifies $|\delta_0 - 1| < 10^{-3}$.

m	d	φ_{tw}-type	k	$\mathrm{Vol}^{1/k}$	δ			$\max_{1\le i\le k}\|\mathbf{b}_i\|_2$		
					raw	LLL	bkz_{40}	raw	LLL	bkz_{40}
159	1	iso/exp	155	8.989	6.143	1.898	1.921	3168.773	35.391	35.703
		iso/tw	155	10.088	7.533	2.117	2.143	4481.257	38.437	37.421
		noiso/exp	155	8.989	6.143	1.894	1.905	3168.773	34.229	34.689
		noiso/tw	155	10.088	7.533	2.119	2.139	4481.257	37.723	38.596

m	d	φ_{tw}-type	k	$\mathrm{Vol}^{1/k}$	δ			$\max_{1\le i\le k}\|\mathbf{b}_i\|_2$		
					raw	LLL	bkz_{40}	raw	LLL	bkz_{40}
159	2	iso/exp	259	9.572	6.902	2.028	2.036	3168.773	36.062	35.703
		iso/tw	259	10.258	8.805	2.313	2.337	4481.257	38.437	37.670
		noiso/exp	259	9.572	6.902	2.024	2.024	3168.773	35.579	35.802
		noiso/tw	259	10.258	8.805	2.317	2.334	4481.257	37.723	38.596
	3	iso/exp	363	9.978	7.602	2.066	2.066	3168.773	37.480	37.132
		iso/tw	363	10.484	9.857	2.373	2.397	4481.257	39.327	39.938
		noiso/exp	363	9.978	7.602	2.064	2.064	3168.773	38.643	38.255
		noiso/tw	363	10.484	9.857	2.376	2.392	4481.257	39.286	41.548
149	1	iso/exp	221	9.697	12.473	2.305	2.244	12554.466	44.327	44.326
		iso/tw	221	10.883	15.626	2.672	2.602	17754.669	49.653	49.399
		noiso/exp	221	9.697	12.473	2.307	2.266	12554.466	43.736	45.013
		noiso/tw	221	10.883	15.626	2.668	2.612	17754.669	49.143	48.693
	2	iso/exp	369	10.150	14.472	2.507	2.467	12554.466	47.719	46.438
		iso/tw	369	10.878	18.958	2.982	2.936	17754.669	52.622	53.154
		noiso/exp	369	10.150	14.472	2.509	2.483	12554.466	48.576	47.820
		noiso/tw	369	10.878	18.958	2.982	2.949	17754.669	54.041	50.666
	3	iso/exp	517	10.410	22.211	2.569	2.531	85211.593	47.719	48.556
		iso/tw	517	10.938	29.658	3.084	3.050	120507.386	52.788	53.154
		noiso/exp	517	10.410	22.211	2.569	2.552	85211.593	48.576	48.778
		noiso/tw	517	10.938	29.658	3.085	3.058	120507.386	54.041	52.131
	4	iso/exp	665	10.632	20.731	2.606	2.576	85211.593	47.768	48.556
		iso/tw	665	11.050	27.968	3.149	3.117	120507.386	53.017	53.154
		noiso/exp	665	10.632	20.731	2.606	2.594	85211.593	48.576	48.778
		noiso/tw	665	11.050	27.968	3.149	3.128	120507.386	54.041	52.385
516	1	iso/exp	251	9.395	6.508	2.341	2.359	4850.233	44.290	43.783
		iso/tw	251	10.544	8.112	2.739	2.733	6859.195	49.680	50.548
		noiso/exp	251	9.395	6.508	2.342	2.354	4850.233	42.774	44.385
		noiso/tw	251	10.544	8.112	2.730	2.739	6859.195	52.260	50.964
	2	iso/exp	419	9.818	8.208	2.550	2.565	5761.443	46.559	46.426
		iso/tw	419	10.522	10.682	3.059	3.062	8147.832	51.931	53.538
		noiso/exp	419	9.818	8.208	2.549	2.557	5761.443	46.306	47.683
		noiso/tw	419	10.522	10.682	3.055	3.064	8147.832	52.534	51.448
	3	iso/exp	587	10.321	10.348	2.620	2.623	9544.834	49.096	49.971
		iso/tw	587	10.845	13.713	3.168	3.167	13498.373	56.763	56.892
		noiso/exp	587	10.321	10.348	2.617	2.615	9544.834	51.019	51.870
		noiso/tw	587	10.845	13.713	3.169	3.167	13498.373	54.998	57.177
	4	iso/exp	755	10.650	12.682	2.652	2.652	26820.239	54.045	52.543
		iso/tw	755	11.068	16.973	3.221	3.219	37929.528	58.551	56.892
		noiso/exp	755	10.650	12.682	2.649	2.650	26820.239	51.019	51.870
		noiso/tw	755	11.068	16.973	3.221	3.220	37929.528	57.437	57.177

Table 5.4 – Geometric characteristics of L_{sat} for some cyclotomic fields. Comparison between choices iso/noiso and exp/tw.

5.5.2 Gram-Schmidt logarithm norms

Here, we provide figures showing the Gram-Schmidt log norms for other cyclotomic fields and number of orbits, comparing values before and after reduction.

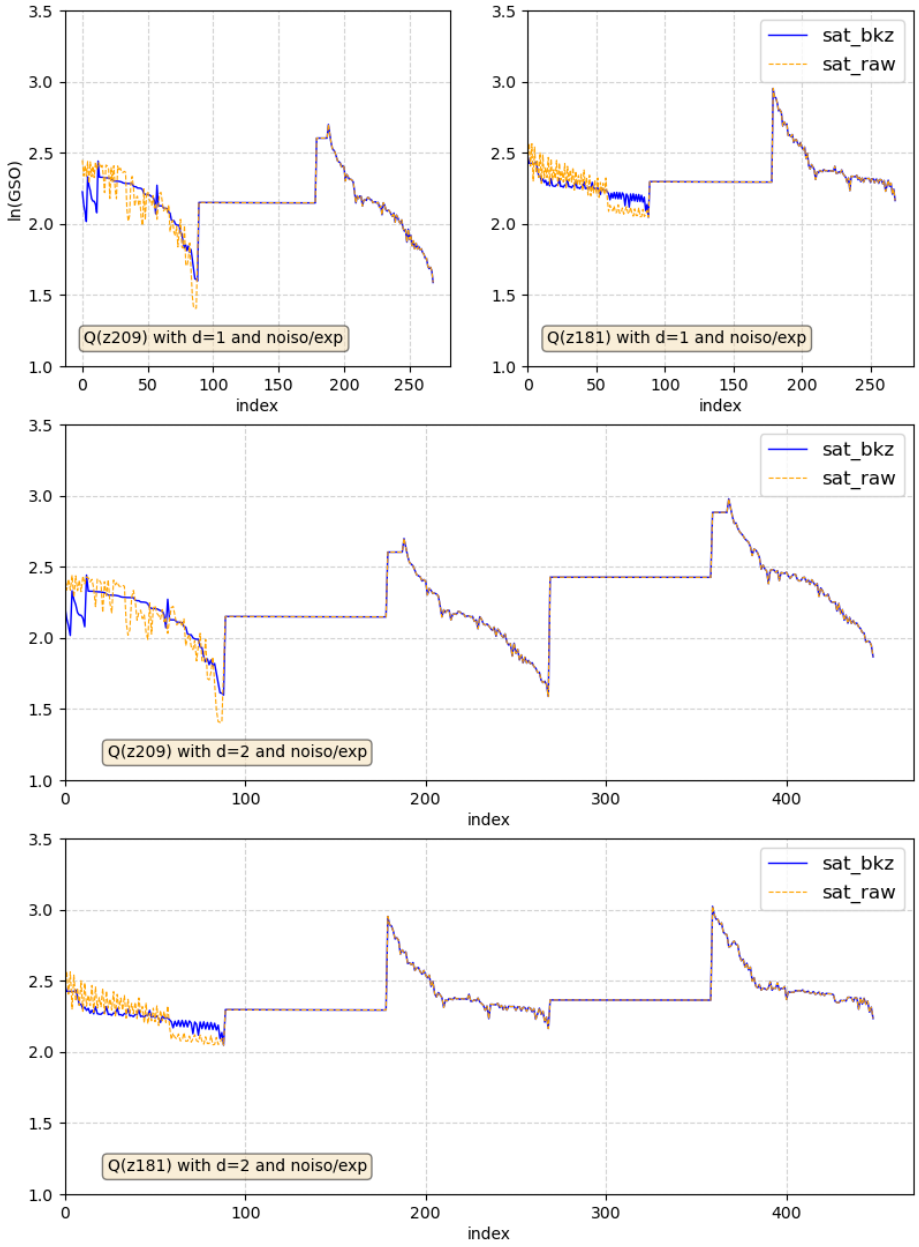

FIGURE 5.5 – L_{sat} lattices for $\mathbb{Q}(\zeta_{209})$ and $\mathbb{Q}(\zeta_{181})$: Gram-Schmidt log norms before and after reduction by BKZ_{40}, for $d = 1$ and $d = 2$ G_m-orbits.

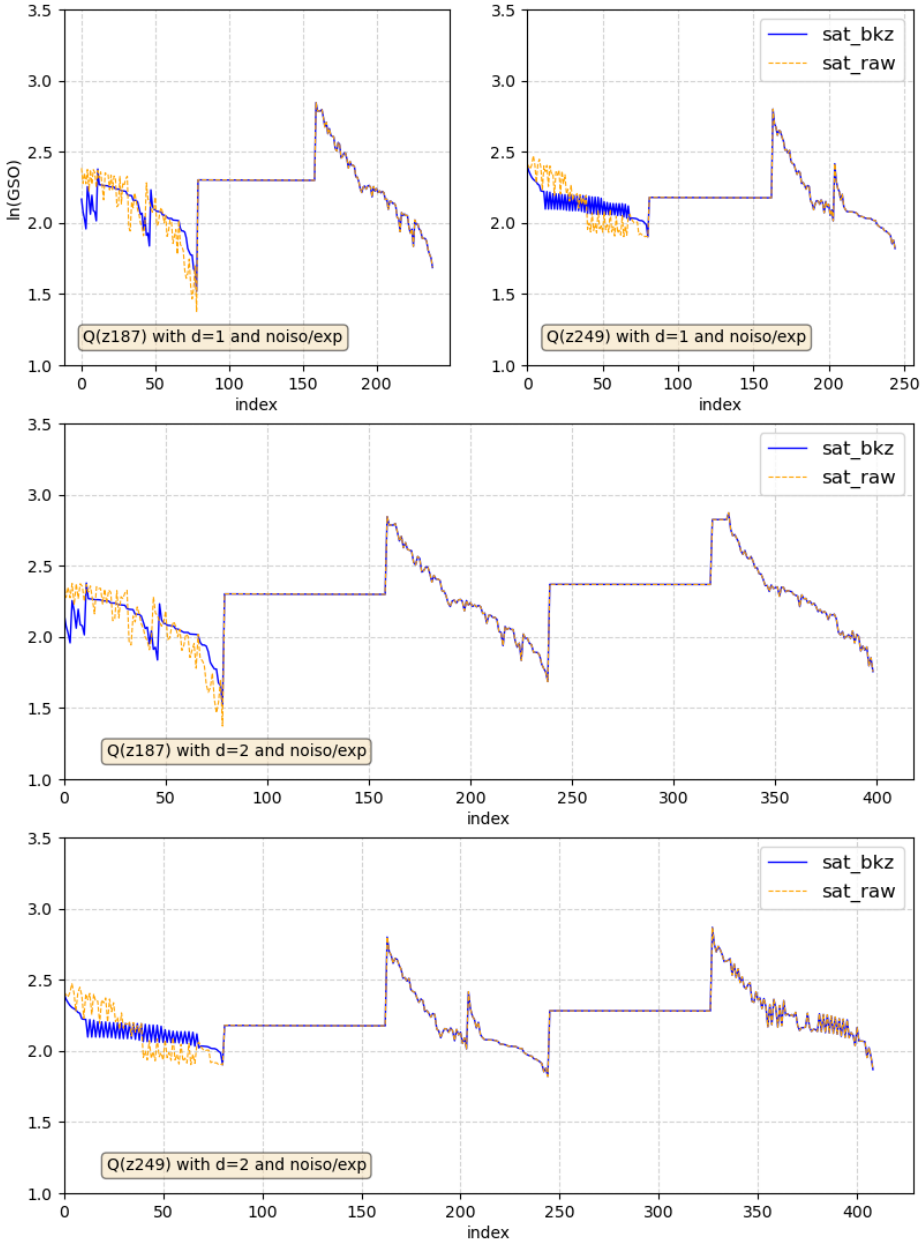

Figure 5.6 – L_{sat} lattices for $\mathbb{Q}(\zeta_{187})$ and $\mathbb{Q}(\zeta_{249})$: Gram-Schmidt log norms before and after reduction by BKZ_{40}, for $d = 1$ and $d = 2$ G_m-orbits.

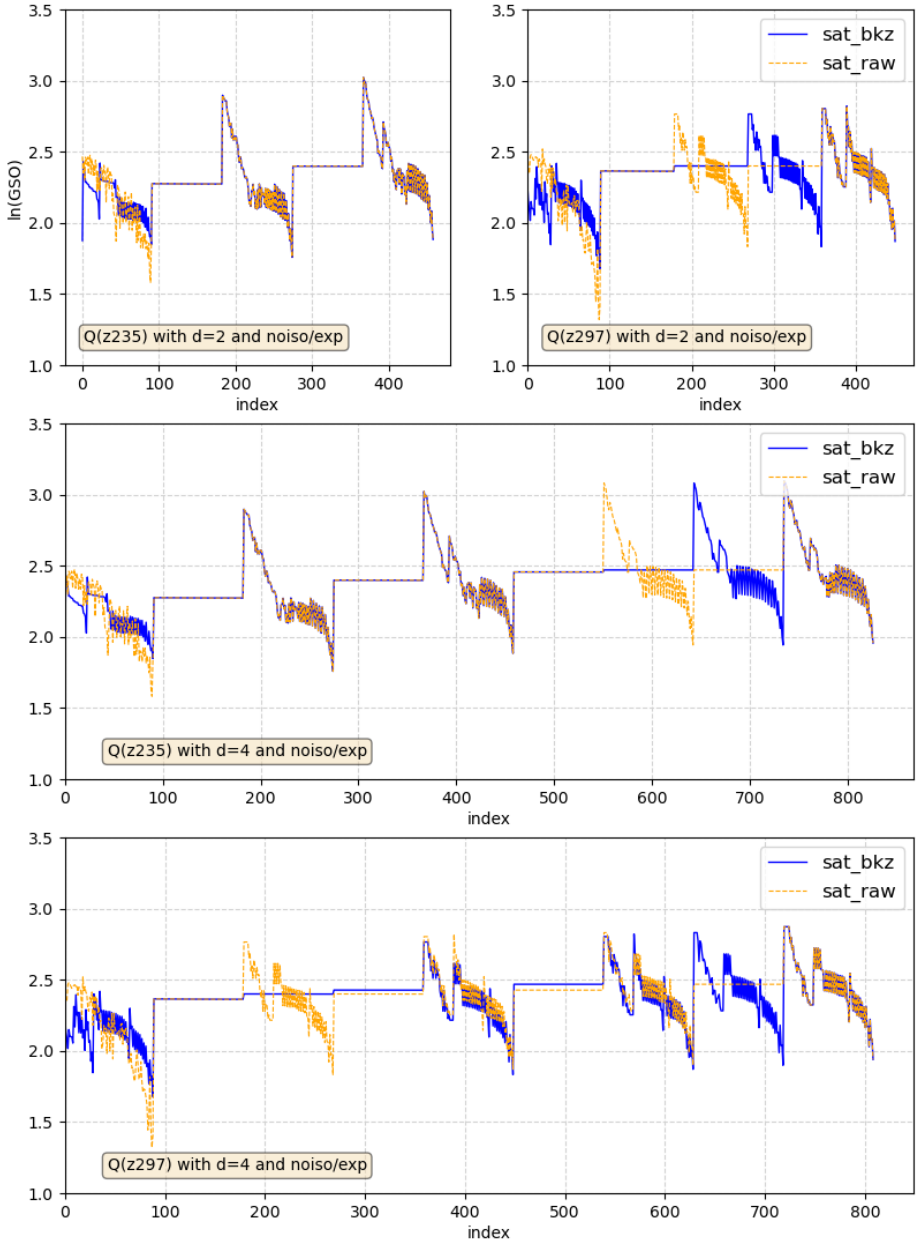

FIGURE 5.7 – L_{sat} lattices for $\mathbb{Q}(\zeta_{235})$ and $\mathbb{Q}(\zeta_{297})$: Gram-Schmidt log norms before and after reduction by BKZ_{40}, for $d = 2$ and $d = 4$ G_m-orbits.

Finally, Fig. 5.8 shows the impact of the four choices of log-$\mathcal{S}$-embedding on the Gram-Schmidt logarithm norms of the unreduced basis $\varphi(\mathfrak{F}_{\mathsf{sat}})$.

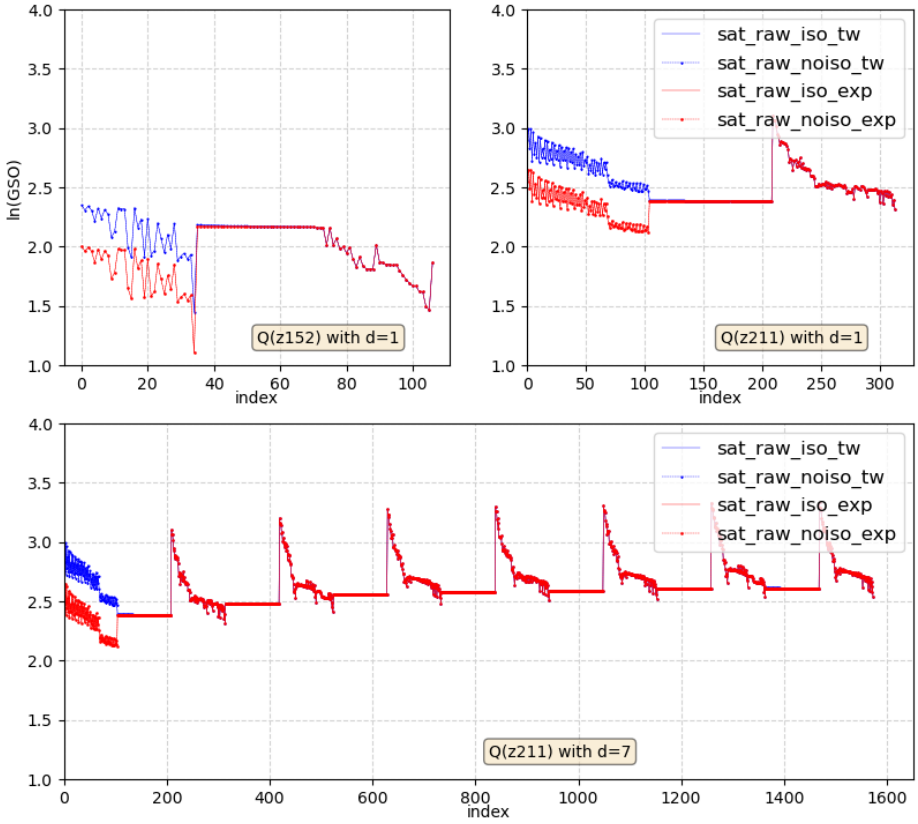

Figure 5.8 – L_{sat} lattices for $\mathbb{Q}(\zeta_{149})$ and $\mathbb{Q}(\zeta_{211})$: effect of the log-$\mathcal{S}$-embedding choices iso/noiso and exp/tw.

Conclusion and Perspectives

Eventually, the results of this thesis take place in the broad context of $\mathcal{S}$-unit attacks against the *Shortest Vector Problem* (Svp) in ideal lattices. First, we propose in Ch. 3 a new *Twisted* version of the PHS algorithm [PHS19a], using the $\mathcal{S}$-units formalism and the *Product Formula* to twist the log-$\mathcal{S}$-embedding with natural number-theoretic weights at finite places. This so-called Twisted-PHS algorithm provably reaches the same asymptotic trade-off between runtime and approximation factor than the PHS algorithm. On the practical side though, we provide the first experimental evidence that under this proper normalization, the log-$\mathcal{S}$-unit lattices at hand seem to behave much better in lattice reduction algorithms, as well as displaying very peculiar geometric characteristics, close to those of orthogonal lattices. The *exact* approximation factors obtained by our Twisted-PHS algorithm in small dimensions are strikingly small, hinting at a possible subexponential (or better) growth. In order to reach meaningful dimensions where asymptotic phenomena start to express, we exhibit in Ch. 4 a *short basis* of the Stickelberger ideal of *any* cyclotomic field, and show how the corresponding explicit algebraic generators are easily computed *via* Jacobi sums. Finally, using in Ch. 5 these extended Stickelberger techniques and the lattice of all real class group relations, we were able on one hand to remove almost all quantum steps in the CDW algorithm, and on the other hand to approximate the Twisted-PHS algorithm in all cyclotomic fields of degree up to 210. These large degree experiments confirmed our initial observations about the geometric peculiarities of the log-$\mathcal{S}$-unit lattices in *all* cyclotomic fields of degree up to 210, and allowed us to obtain an upper bound of the performance of Twisted-PHS in meaningful medium dimensions.

Nevertheless, at this point the obtained approximation factors are not sufficient to derive an asymptotic general behaviour for $\mathcal{S}$-unit attacks. Although our results do not show a catastrophic impact of $\mathcal{S}$-unit attacks, they do neither allow to dismiss this particular threat. Whereas gathering these experimental data is of utmost importance to better understand concretely $\mathcal{S}$-unit attacks and back up discussions on the hardness of id-Svp, further investigation is still needed to derive a sound asymptotic estimator.

Hence, we identify two axes for further research: the first is extending the range and scope of experiments to better circumscribe the relevant parameters needed to evaluate the performance of $\mathcal{S}$-unit attacks; the second is to work towards obtaining a sound asymptotic estimator relying on properly identified and verified heuristics.

Further Concrete Experimental Data

At the moment, our experiments constitute a first important step towards assessing the performance of $\mathcal{S}$-unit attacks. There are several directions to explore in order to strengthen our observations in the general case.

Other cyclotomic fields. The simplest improvement is to capture results for cyclotomic fields K_m such that $h_m^+ > 1$. Indeed, the lattice of real class group relations in the case $h_m^+ = 1$ is equivalent to the concatenation of two identity matrices and solving CVP in this lattice can be done in an optimal way. We therefore expect on one hand a greater gap between the CDW algorithm and the Twisted-PHS algorithm, and on the other hand a noticeable effect of the h_m^+ part on the approximation factor in both cases. At the moment, this case has not been dealt with for the only technical reason that the real class group relations matrix seems not to be easily accessible in SAGEMATH.

The second obvious improvement would be to make further progress to fields of higher degrees. This necessitates to tweak the parameters of the PARI/GP routine for computing class groups, and to work on a parallelized implementation of the Buchmann algorithm for the search of class group relations. A nice target would be to reach e.g., $\mathbb{Q}(\zeta_{512})$, which has degree 256.

Densify log-$\mathcal{S}$-unit sublattices. Due to the h_m^- part in the index of our full-rank families of independent $\mathcal{S}$-units, our log-$\mathcal{S}$-unit sublattices in medium dimensions are still far from the full log-$\mathcal{S}$-unit lattice. A regrettable consequence is that our log-$\mathcal{S}$-unit lattices reach maximal density for $d = 1$ orbit of split prime ideals, whereas the Twisted-PHS generally predicts a greater optimum $d = d_{\mathsf{max}}$ w.r.t. the log-$\mathcal{S}$-unit lattice density.

Hence, some significant effort should be put to improve the saturation step in order to capture as many prime factors of h_m^- as is reasonably possible. There are two bottlenecks: the computation of e-th-root characters, which at first glance costs $O(\sqrt{e})$, and the e-th-root computation itself. The latter could be virtualized by considering only the absolute values of the complex embeddings while discarding their complex arguments. However, it seems illusory to expect to capture prime factors of h_m^- significatively larger than 64 bits. For our range of computations, this might still be interesting in a non-negligible proportion of the fields, but note e.g., that h_{197}^- is divisible by 939830268487086665622561154 9, a 93-bit prime!

This would *in fine* allow to verify in medium dimensions the behaviour of the Twisted-PHS algorithm when the density of the log-$\mathcal{S}$-unit lattice is increasing with $d > 1$.

Guarantee that all finite valuations are positive. A very painful requirement of the algorithm is that the solution vector corresponds to an element of the challenge ideal. In Ch. 3, this was done by applying a diagonal drift on finite place coordinates, searching for an optimum value by dichotomy. As seen in Ch. 5, this method does not scale properly, and we had to use a randomized strategy applying $O(n)$ random drifts in ℓ_∞-norm balls of radius 1 centered on guessed diagonal values. This is quite costly and still not satisfactory as we observe a noticeable variance between two random runs on the same target.

In an ideal world, one could hope for a specifically designed algorithm, like a modified or backtracking Babai's Nearest Plane Algorithm, that guarantees that the close solution which is returned lies in the correct cone w.r.t. the target.

Obtain verifiable examples in medium dimensions. The exact approximation factors shown in Ch. 3 for Twisted-PHS are completely verifiable examples, as the ClDLP step is concretely performed. However, this still represents in practice the main bottleneck in high dimensions, where obtaining a single relation involving a challenge prime ideal of big norm and many small prime ideals in the factor base is significantly harder in practice than obtaining many relations involving only ideals of the factor base. This is why in Ch. 5 we use random targets simulating the output of the ClDLP step. Hence, obtaining verifiable concrete examples would be very useful in order to confirm that the approximation factors estimated *via* random targets still match the reality beyond the small dimensions reached in Ch. 3.

We could think of two possible ways of performing this ClDlp step, that can even be combined together. The first would be to use a special-$\mathfrak{q}$ strategy, like in the *General Number Field Sieve* (GNFS) context. However, this requires to sieve in large dimensions, and the algebraic norm of the elements grows extremely rapidly. The second way aims at reducing the dimension of the ClDlp by generalizing the Gentry-Szydlo algorithm to this problem. This Gentry-Szydlo method would probably allow to double the reachable dimension for explicit ClDlp computations, hence this would currently reach dimensions 100 to 120, since we did not yet succeed in performing explicit ClDlp computations beyond dimensions 50 to 60.

Towards an $\mathcal{S}$-unit Attack Asymptotic Simulator

The works of this thesis altogether allowed to reach a state where it becomes possible to extract the meaningful properties of log-$\mathcal{S}$-unit lattices that are not bound to small dimensions pathological phenomena. At this point, we need to address two theoretical questions: first, explain the striking ease of reduction and orthogonality defects of our obtained log-$\mathcal{S}$-unit lattices, then obtain an heuristic estimation of the final approximation factor. The former would allow to establish the concrete running time of the preprocessing phase for these specific lattices, and the latter should give, at least, a *lower* bound of the performance of the Twisted-PHS algorithm.

The first question can be explored as follows. On the one hand, it seems possible to obtain sensible estimations of the size of the log-$\mathcal{S}$-embeddings of $\mathcal{S}$-units from easily computable number-theoretic values. Further, the Gram-Schmidt orthogonalization matrix has a very specific structure that can be made explicit, due to the special shape of all the vectors of the basis. Indeed, the basis vectors have two fixed balanced parts whose sum is equal to the logarithm of the algebraic norm of the corresponding $\mathcal{S}$-units: on the infinite places, this weight is borne evenly by at most n coordinates; on finite places, this weight is borne by $k \geq n$ distinct places, each non-zero coordinates being weighted by some $\ln \mathcal{N}(\mathfrak{p})$, a particular distortion which possibly has noticeable consequences.

A pending question is to determine whether this particular shape of the basis vectors is sufficient to explain the geometric behaviour of the log-$\mathcal{S}$-unit lattices, or whether a deeper number-theoretic explanation is mandatory. Simulating the distribution of the coordinates could give important insights. If the approximation factors obtained with these virtual lattices coincide with those of Ch. 5 in medium dimensions, this would allow to extrapolate the performance of $\mathcal{S}$-unit attacks in cryptographically relevant dimensions. This would also offer much freedom to test a wide variety of parameters. In particular, it would allow to test whether extending the factor base beyond the point of maximum density is indeed helpful, as claimed by Bernstein in his talk at SIAM Conference on 20-th August 2021.

Bibliography

[Ajt96] M. Ajtai: *Generating hard instances of lattice problems.* In *STOC*, pp. 99–108, ACM, 1996.

[Ajt98] M. Ajtai: *The Shortest Vector Problem in L_2 is NP-hard for randomized reductions.* In *STOC*, pp. 10–19, ACM, 1998.

[Bab86] L. Babai: *On Lovász' lattice reduction and the nearest lattice point problem.* Combinatorica, **6**(1), pp. 1–13, 1986.

[Bac90] É. Bach: *Explicit bounds for primality testing and related problems.* Math. Comp., **55**(191), pp. 355–380, 1990.

[BBV+17] J. Bauch, D. Bernstein, H. de Valence, T. Lange, C. van Vredendaal: *Short generators without quantum computers: the case of multiquadratics.* In *EUROCRYPT (1)*, vol. 10210 of *LNCS*, pp. 27–59, Springer, 2017.

[BCLV17] D. J. Bernstein, C. Chuengsatiansup, T. Lange, C. van Vredendaal: *NTRU Prime: Reducing attack surface at low cost.* In *SAC*, vol. 10719 of *LNCS*, pp. 235–260, Springer, 2017.

[BCP97] W. Bosma, J. Cannon, C. Playoust: *The Magma algebra system. I. The user language.* J. Symbolic Comput., **24**(3-4), pp. 235–265, 1997, Computational algebra and number theory (London, 1993).

[BDF08] K. Belabas, F. Diaz y Diaz, E. Friedman: *Small generators of the ideal class group.* Math. Comp., **77**(262), pp. 1185–1197, 2008.

[BDPW20] K. d. Boer, L. Ducas, A. Pellet-Mary, B. Wesolowski: *Random self-reducibility of Ideal-SVP via Arakelov random walks.* In *CRYPTO (2)*, vol. 12171 of *LNCS*, pp. 243–273, Springer, 2020.

[BEF+17] J. Biasse, T. Espitau, P. Fouque, A. Gélin, P. Kirchner: *Computing generator in cyclotomic integer rings.* In *EUROCRYPT (1)*, vol. 10210 of *LNCS*, pp. 60–88, Springer, 2017.

[BEHC21] P. Browne, R. Egan, F. Hegarty, P. Ó. Catháin: *A survey of the Hadamard maximal determinant problem.* ArXiv pre-print arXiv:2104.06756 [math.CO], 2021.

[Bel04] K. Belabas: *A relative van Hoeij algorithm over number fields.* J. Symb. Comput., **37**(5), pp. 641–668, 2004.

[BF14] J. Biasse, C. Fieker: *Subexponential class group and unit group computation in large degree number fields.* LMS J. Comp. Math., **17**(A), pp. 385–403, 2014.

[BF15] K. Belabas, E. Friedman: *Computing the residue of the Dedekind zeta function.* Math. Comp., **84**, pp. 357–369, 2015.

[BFHP21] J. Biasse, C. Fieker, T. Hofmann, A. Page: *Norm relations and computational problems in number fields.* ArXiv pre-print arXiv:2002.12332v3 [math.NT], 2021.

[BK21] O. Bernard, R. Kučera: *A short basis of the Stickelberger ideal of a cyclotomic field.* ArXiv pre-print arXiv:2109.13329 [math.NT], 2021.

[BL21] D. J. Bernstein, T. Lange: *Non-randomness of S-unit lattices.* Cryptology ePrint Archive, Report 2021/1428, 2021, https://ia.cr/2021/1428.

[BLNR21] O. BERNARD, A. LESAVOUREY, T.-H. NGUYEN, A. ROUX-LANGLOIS: *Log-S-unit lattices using explicit Stickelberger generators to solve Approx Ideal-SVP*. Cryptology ePrint Archive ePrint:2021/1384, 2021.

[BLP93] J. BUHLER, H. LENSTRA, C. POMERANCE: *Factoring integers with the Number Field Sieve*, vol. 1554 of *Lecture Notes in Math.* Springer, 1993.

[BMT15] D. W. BOYD, G. MARTIN, M. THOM: *Squarefree values of trinomial discriminants*. LMS J. Comput. Math., **18**(1), pp. 148–169, 2015.

[BPR04] J. BUHLER, C. POMERANCE, L. ROBERTSON: *Heuristics for class numbers of prime-power real cyclotomic fields*. Fields Inst. Commun., **41**, pp. 149–157, 2004.

[BR20] O. BERNARD, A. ROUX-LANGLOIS: *Twisted-PHS: Using the product formula to solve Approx-SVP in ideal lattices*. In *ASIACRYPT*, vol. 12492 of *LNCS*, pp. 349–380, Springer, 2020.

[BS16] J.-F. BIASSE, F. SONG: *Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields*. In *SODA*, pp. 893–902, SIAM, 2016.

[BV18] J. BIASSE, C. VAN VREDENDAAL: *Fast multiquadratic S-unit computation and application to the calculation of class groups*. In *ANTS-XIII*, vol. 2 of *The Open Book Series*, pp. 103–118, Mathematical Sciences Publisher, 2018.

[CDPR16] R. CRAMER, L. DUCAS, C. PEIKERT, O. REGEV: *Recovering short generators of principal ideals in cyclotomic rings*. In *EUROCRYPT (2)*, vol. 9666 of *LNCS*, pp. 559–585, Springer, 2016.

[CDW17] R. CRAMER, L. DUCAS, B. WESOLOWSKI: *Short Stickelberger class relations and application to Ideal-SVP*. In *EUROCRYPT (1)*, vol. 10210 of *LNCS*, pp. 324–348, Springer, 2017.

[CDW21] R. CRAMER, L. DUCAS, B. WESOLOWSKI: *Mildly short vectors in cyclotomic ideal lattices in quantum polynomial time*. J. ACM, **68**(2), 2021.

[CGS14] P. CAMPBELL, M. GROVES, D. SHEPHERD: *Soliloquy: A cautionary tale*, 2014, available at http://docbox.etsi.org/Workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf.

[Che13] Y. CHEN: *Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe*. Ph.D. thesis, Paris 7, 2013.

[CN11] Y. CHEN, P. Q. NGUYEN: *BKZ 2.0: Better lattice security estimates*. In *ASIACRYPT*, vol. 7073 of *LNCS*, pp. 1–20, Springer, 2011.

[Coh93] H. COHEN: *A course in computational algebraic number theory*, vol. 138 of *Graduate texts in mathematics*. Springer, 1993.

[Con] K. CONRAD: *Ostrowski for number fields*. In *Expository papers on Algebraic Number Theory*, available at https://kconrad.math.uconn.edu/blurbs/gradnumthy/ostrowskinumbfield.pdf.

[DPW19] L. DUCAS, M. PLANÇON, B. WESOLOWSKI: *On the shortness of vectors to be found by the Ideal-SVP quantum algorithm*. In *CRYPTO (1)*, vol. 11692 of *LNCS*, pp. 322–351, Springer, 2019.

[EHKS14] K. EISENTRÄGER, S. HALLGREN, A. Y. KITAEV, F. SONG: *A quantum algorithm for computing the unit group of an arbitrary degree number field*. In *STOC*, pp. 293–302, ACM, 2014.

[FGW92] G. FUNG, A. GRANVILLE, H. C. WILLIAMS: *Computation of the first factor of the class number of cyclotomic fields*. J. Number Theory, **42**(3), pp. 297–312, 1992.

[FpL16] FPLLL DEVELOPMENT TEAM: *fplll, a lattice reduction library*, 2016, available at https://github.com/fplll/fplll.

[Gal12] S. D. GALBRAITH: *Mathematics of Public Key Cryptography*. Cambridge University Press, 2012.

[Gél17] A. GÉLIN: *Calcul de groupes de classes d'un corps de nombres et applications à la cryptologie*. Ph.D. thesis, UPMC Paris 6, 2017.

[GK89] R. GOLD, J. KIM: *Bases for cyclotomic units*. Compos. Math., **71**(1), pp. 13–27, 1989.

[GM16] L. GRENIÉ, G. MOLTENI: *Explicit versions of the prime ideal theorem for Dedekind zeta functions under GRH*. Math. Comp., **85**(298), pp. 889–906, 2016.

[GN08] N. GAMA, P. Q. NGUYEN: *Predicting lattice reduction*. In *EUROCRYPT*, vol. 4965 of *LNCS*, pp. 31–51, Springer, 2008.

[HPS98] J. HOFFSTEIN, J. PIPHER, J. H. SILVERMAN: *NTRU: A ring-based public key cryptosystem*. In *ANTS*, vol. 1423 of *Lecture Notes in Computer Science*, pp. 267–288, Springer, 1998.

[HW38] G. H. HARDY, E. M. WRIGHT: *An Introduction to the Theory of Numbers*. Oxford University Press, 1938, Fourth Edition.

[HWB17] P. HOLZER, T. WUNDERER, J. A. BUCHMANN: *Recovering short generators of principal fractional ideals in cyclotomic fields of conductor $p^\alpha q^\beta$*. In *INDOCRYPT*, vol. 10698 of *LNCS*, pp. 346–368, Springer, 2017.

[Kom75] K. KOMATSU: *Integral bases in algebraic number fields*. Journal für die reine und angewandte Mathematik, **1975**(278-279), pp. 137–144, 1975.

[Kuč86] R. KUČERA: *On a certain subideal of the Stickelberger ideal of a cyclotomic field*. Archivum Mathematicum, **22**(1), pp. 7–19, 1986.

[Kuč92] R. KUČERA: *On bases of the Stickelberger ideal and of the group of circular units of a cyclotomic field*. J. Number Theory, **40**(3), pp. 284–316, 1992.

[Kuč96] R. KUČERA: *On the Stickelberger ideal and circular units of a compositum of quadratic fields*. J. Number Theory, **56**(1), pp. 139–166, 1996.

[Laa16] T. LAARHOVEN: *Sieving for closest lattice vectors (with preprocessing)*. In *SAC*, vol. 10532 of *LNCS*, pp. 523–542, Springer, 2016.

[Lan03] E. LANDAU: *Neuer Beweis des Primzahlsatzes und Beweis des Primidealsatzes*. Math. Ann., **56**, pp. 645–670, 1903.

[Les21] A. LESAVOUREY: *Usability of structured lattices for a post-quantum cryptography: practical computations, and a study of some real Kummer extensions*. Ph.D. thesis, University of Wollongong, 2021.

[LLL82] A. K. LENSTRA, H. W. LENSTRA, L. LOVÁSZ: *Factoring polynomials with rational coefficients*. Math. Ann., **261**, pp. 515–534, 1982.

[LM06] V. LYUBASHEVSKY, D. MICCIANCIO: *Generalized compact knapsacks are collision resistant*. In *ICALP*, vol. 4052 of *LNCS*, pp. 144–155, Springer, 2006.

[Lou00] S. LOUBOUTIN: *Explicit bounds for residues of Dedekind zeta functions, values of L-functions at $s = 1$, and relative class numbers*. J. Number Theory, **85**(2), pp. 263–282, 2000.

[Lou14] S. LOUBOUTIN: *Upper bounds on relative class number of cyclotomic fields*. Math. Slovaca, **64**(1), pp. 21–26, 2014.

[LPR10] V. LYUBASHEVSKY, C. PEIKERT, O. REGEV: *On ideal lattices and learning with errors over rings*. In *EUROCRYPT*, vol. 6110 of *LNCS*, pp. 1–23, Springer, 2010.

[LPS20] A. LESAVOUREY, T. PLANTARD, W. SUSILO: *Short principal ideal problem in multicubic fields*. J. Math. Cryptol., **14**(1), pp. 359–392, 2020.

[LS15] A. LANGLOIS, D. STEHLÉ: *Worst-case to average-case reductions for module lattices*. Des. Codes Cryptogr., **75**(3), pp. 565–599, 2015.

[MG02] D. MICCIANCIO, S. GOLDWASSER: *Complexity of Lattice Problems*, vol. 671 of *The Kluwer International Series in Engineering and Computer Science*. Springer, 2002.

[Mil14] J. C. MILLER: *Class numbers of real cyclotomic fields of composite conductor*. LMS J. Comput. Math., **17**, pp. 404–417, 2014.

[Nar04] W. NARKIEWICZ: *Elementary and Analytic Theory of Algebraic Numbers*. Springer Monographs in Mathematics, Springer, 3 edn., 2004.

[Neu99] J. NEUKIRCH: *Algebraic Number Theory*, vol. 322 of *Grundlehren des mathematischen Wissenschaften*. Springer, 1999.

[NS06] P. Q. NGUYEN, D. STEHLÉ: *LLL on the average*. In *ANTS*, vol. 4076 of *LNCS*, pp. 238–256, Springer, 2006.

[NV10] P. Q. NGUYEN, B. VALLÉE (eds.): *The LLL Algorithm*. Information Security and Cryptography, Springer, 2010.

[Pei16] C. PEIKERT: *A decade of lattice cryptography*. Foundations and Trends in Theoretical Computer Science, **10**(4), pp. 283–424, 2016.

[PHS19a] A. PELLET-MARY, G. HANROT, D. STEHLÉ: *Approx-SVP in ideal lattices with pre-processing*. In *EUROCRYPT (2)*, vol. 11477 of *LNCS*, pp. 685–716, Springer, 2019.

[PHS19b] A. PELLET-MARY, G. HANROT, D. STEHLÉ: *Published code of "Approx-SVP in ideal lattices with pre-processing"*, 2019, available at https://apelletm.pages.math.cnrs.fr/page-perso/code/code-approx-ideal-svp.zip.

[PR06] C. PEIKERT, A. ROSEN: *Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices*. In *TCC*, vol. 3876 of *LNCS*, pp. 145–166, Springer, 2006.

[PZ89] M. POHST, H. ZASSENHAUS: *Algorithmic Algebraic Number Theory*. Encyclop. Math. Appl., Cambridge University Press, 1989.

[Reg05] O. REGEV: *On lattices, learning with errors, random linear codes, and cryptography*. In *STOC*, pp. 84–93, ACM, 2005.

[Sag20] SAGE DEVELOPERS: *SageMath, the Sage Mathematics Software System (Version 9.0)*, 2020, available at https://www.sagemath.org.

[Sch87] C. SCHNORR: *A hierarchy of polynomial time lattice basis reduction algorithms*. Theor. Comput. Sci., **53**, pp. 201–224, 1987.

[Sch03] R. SCHOOF: *Class numbers of real cyclotomic fields of prime conductor*. Math. Comp., **72**(242), pp. 913–937, 2003.

[Sch08] R. SCHOOF: *Catalan's Conjecture*. Universitext, Springer, 2008.

[SE94] C. SCHNORR, M. EUCHNER: *Lattice basis reduction: Improved practical algorithms and solving subset sum problems*. Math. Program., **66**, pp. 181–199, 1994.

[Sho97] P. W. SHOR: *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*. SIAM J. Comput., **26**(5), pp. 1484–1509, 1997.

[Sin78] W. SINNOTT: *On the Stickelberger ideal and the circular units of a cyclotomic field*. Ann. Math., **108**(1), pp. 107–134, 1978.

[Sin80] W. SINNOTT: *On the Stickelberger ideal and the circular units of an abelian field*. Invent. Math., **62**, pp. 181–234, 1980.

[SSTX09] D. STEHLÉ, R. STEINFELD, K. TANAKA, K. XAGAWA: *Efficient public key encryption based on ideal lattices*. In *ASIACRYPT*, vol. 5912 of *LNCS*, pp. 617–635, Springer, 2009.

[Swa62] R. G. SWAN: *Factorization of polynomials over finite fields*. Pacific J. Math., **12**(3), pp. 1099–1106, 1962.

[Tho12] E. THOMÉ: *Square root algorithms for the Number Field Sieve*. In *WAIFI*, vol. 7369 of *LNCS*, pp. 208–224, Springer, 2012.

[Was97] L. C. WASHINGTON: *Introduction to Cyclotomic Fields*, vol. 83 of *Graduate Texts in Mathematics*. Springer, 2 edn., 1997.

[Wes18] B. WESOLOWSKI: *Generating subgroups of ray class groups with small prime ideals*. In *ANTS-XIII*, vol. 2 of *The Open Book Series*, pp. 461–478, Mathematical Sciences Publisher, 2018.

[Xu13] P. XU: *Experimental quality evaluation of lattice basis reduction methods for decorrelating low-dimensional integer least squares problems*. EURASIP J. Adv. Signal Process., **2013**, pp. 137–165, 2013.

DOCTORAT BRETAGNE LOIRE / MATHSTIC

UNIVERSITÉ DE RENNES 1

Titre : Cryptanalyse Algébrique du Problème du Plus Court Vecteur dans les Réseaux Idéaux

Mots-clés : Réseaux idéaux, Problème du Plus Court Vecteur, S-unités, Stickelberger, Algorithme Twisted-PHS

Résumé : Les travaux de cette thèse portent sur les attaques par $\mathcal{S}$-unités contre le *Problème du Plus Court Vecteur* (SVP) dans les réseaux idéaux. Tout d'abord, nous proposons une version *Tordue* de l'algorithme PHS, utilisant le formalisme des $\mathcal{S}$-unités et la *Formule du Produit* pour pondérer le $\mathcal{S}$-plongement logarithmique avec les poids standards de théorie des nombres sur les places finies. Sur le plan théorique, cet algorithme nommé Twisted-PHS réalise le même compromis temps-facteur d'approximation que l'algorithme PHS. Sur le plan pratique, nous fournissons la première preuve expérimentale qu'avec cette pondération, les réseaux log-$\mathcal{S}$-unités ont des caractéristiques géométriques très particulières. De plus, les facteurs d'approximation *exacts* obtenus en petite dimension sont remarquablement petits, potentiellement sous-exponentiels ou mieux.

Afin d'atteindre des dimensions où les phénomènes asymptotiques s'expriment, nous exhibons une *base courte* de l'idéal de Stickelberger pour *tout* corps cyclotomique, et montrons comment calculer les générateurs explicites correspondants *via* des sommes de Jacobi.

Finalement, grâce à ces résultats avancés sur l'idéal de Stickelberger, et à l'aide de toutes les relations du groupe de classe réel, nous supprimons presque toutes les étapes quantiques de l'algorithme CDW, et approximons l'algorithme Twisted-PHS pour tous les corps cyclotomiques de degré jusqu'à 210. Cela a permis de confirmer les particularités géométriques des réseaux log-$\mathcal{S}$-unités pondérés, ainsi que d'obtenir une borne supérieure sur les performances de notre algorithme Twisted-PHS en moyenne dimension.

Title: Algebraic Cryptanalysis of the Shortest Vector Problem in Ideal Lattices

Keywords: Ideal lattices, Shortest Vector Problem, Stickelberger, S-units, Twisted-PHS Algorithm

Abstract: The results of this thesis take place in the broad context of $\mathcal{S}$-unit attacks against the *Shortest Vector Problem* (SVP) in ideal lattices.

First, we propose a new *Twisted* version of the PHS algorithm, using the $\mathcal{S}$-units formalism and the *Product Formula* to twist the log-$\mathcal{S}$-embedding with standard number-theoretic weights at finite places. This so-called Twisted-PHS algorithm provably reaches the same asymptotic trade-off between runtime and approximation factor than the PHS algorithm. On the practical side, we provide the first experimental evidence that using this normalization, the log-$\mathcal{S}$-unit lattices have very peculiar geometric characteristics. *Exact* approximation factors obtained in small dimensions are strikingly small, in a way that could be subexponential or better.

In order to reach dimensions where asymptotic phenomena start to express, we exhibit a *short basis* of the Stickelberger ideal of *any* cyclotomic field, and show how its explicit algebraic generators can be computed *via* Jacobi sums.

Finally, using these extended Stickelberger techniques and all real class group relations, we were able to remove almost all quantum steps in the CDW algorithm, and to approximate the Twisted-PHS algorithm in all cyclotomic fields of degree up to 210. This allowed us to confirm the geometric peculiarities of twisted log-$\mathcal{S}$-unit lattices, and to obtain an upper bound of the performance of the Twisted-PHS algorithm in medium dimensions.