



HAL
open science

A new model of greenhouse gas transport in the global atmosphere adapted to the evolution of high-performance computing resources

Zoé Lloret

► **To cite this version:**

Zoé Lloret. A new model of greenhouse gas transport in the global atmosphere adapted to the evolution of high-performance computing resources. Ocean, Atmosphere. Université Paris-Saclay, 2024. English. NNT: 2024UPASJ021 . tel-04768119

HAL Id: tel-04768119

<https://theses.hal.science/tel-04768119v1>

Submitted on 5 Nov 2024

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

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

A new model of greenhouse gas transport in the global atmosphere adapted to the evolution of high-performance computing resources

*Un nouveau modèle du transport des gaz à effet de serre dans
l'atmosphère mondiale adapté à l'évolution des moyens de calcul à
haute performance*

Thèse de doctorat de l'université Paris-Saclay

École doctorale n° 129, Sciences de l'environnement d'Île-de-France (SEIF)
Spécialité de doctorat : Géosciences
Graduate School : Géosciences, climat, environnement et planètes.
Réfèrent : Université de Versailles -Saint-Quentin-en-Yvelines

Thèse préparée dans la unité de recherche **LSCE (UVSQ, CNRS, CEA, université Paris Saclay)**, sous la direction de **Frédéric CHEVALLIER**, Directeur de recherche, et le co-encadrement de **Anne COZIC**, Ingénieure de recherche.

Thèse soutenue à Gif-sur-Yvette, le 4 octobre 2024, par

Zoé LLORET

Composition du jury

Membres du jury avec voix délibérative

Marielle SAUNOIS Professeure, LSCE-UVSQ	Présidente
Sylvie MALARDEL Ingénieur des travaux de la météorologie Hors Classe, HDR, Météo-France	Rapporteuse & Examinatrice
Arthur VIDARD Chargé de recherche, HDR, INRIA	Rapporteur & Examineur
Laurent LI Directeur de recherche, LMD	Examineur

Titre: Un nouveau modèle du transport des gaz à effet de serre dans l'atmosphère mondiale adapté à l'évolution des moyens de calcul à haute performance

Mots clés: Calcul haute performance, Modélisation, Gaz à effet de serre

Résumé: La compréhension du changement climatique en cours implique une bonne connaissance des émissions et des absorptions des gaz à effet de serre sur le globe. Les modèles atmosphériques inverses y contribuent de manière croissante, mais la résolution des modèles planétaires reste grossière, en particulier au regard de l'accroissement de la densité des réseaux de mesure fournissant les observations qu'ils assimilent. Les nouvelles architectures des supercalculateurs, soutenues par les processeurs "Many-core" ou graphiques, permettent une parallélisation massive des calculs et ouvrent des perspectives importantes pour les modèles inverses, que nous explorons dans cette thèse.

Nous considérons le système d'inversion CAMS/LSCE. Il dérive son modèle de transport atmosphérique d'un modèle de circulation générale (GCM), et est piloté par un ensemble de variables météorologiques préalablement générées par ce dernier. Ce GCM a récemment intégré un nouveau cœur dynamique fonctionnant sur un pavage du globe avec des hexagones, DYNAMICO, plus adapté au parallélisme et conçu pour monter efficacement en résolution. Nous évaluons le nouveau GCM en détail pour le transport du dioxyde de carbone (CO_2) à l'aide de mesures sur le globe et dans l'atmosphère, et d'une simulation de référence utilisant la grille classique faite de rectangles. Nous montrons des différences faibles entre les deux versions, qu'une comparaison aux observations disponibles ne permet pas de hiérarchiser. Nos tests vérifient aussi l'accélération des calculs sur la nouvelle grille, en particulier aux résolutions les plus élevées. Nous préparons la génération des variables météorologiques pour le modèle inverse, mais la complexité du code de DYNAMICO ne nous permet pas de l'adapter aux spécifications du modèle inverse (modularité des composantes, développement des codes tangentielles et adjoints) dans cette thèse. En parallèle, nous étudions la possibilité de porter le mod-

èle de transport sur des cartes graphiques. Nous montrons que l'algorithme sur la grille cartésienne classique peut être exprimé de manière suffisamment parallèle pour fonctionner efficacement sur de tels matériels, sans aucune simplification malgré les échanges continus de masse entre les cellules. Le temps de calcul dans la version directe du modèle inverse devient même inférieur au temps de lecture des données météorologiques d'entrée. Nous appuyant sur nos développements avec les cartes graphiques, nous réalisons une inversion pluriannuelle à une résolution mondiale autour du degré, assimilant des observations satellitaires. Par comparaison, la résolution du système avant nos développements était autour de 3 degrés sur l'horizontale, avec deux fois moins de niveaux sur la verticale : ses temps de calcul étaient comparables (ici de l'ordre de la semaine) malgré le nombre bien plus faible de points de grille à traiter. Avec la nouvelle résolution, nous montrons une légère amélioration dans la représentation du transport horizontal et vertical du CO_2 , en comparaison avec des mesures variées, sans pour autant pouvoir attester d'une meilleure estimation des flux de surface. L'ensemble des résultats démontre qu'il n'existe plus aujourd'hui de frein à une augmentation massive de la résolution des modèles atmosphériques inverses planétaires pour le CO_2 , permettant ainsi d'aborder sur l'ensemble du globe des échelles spatiales jusqu'ici réservées aux modèles à aire limitée. Ils ouvrent la perspectives de produits d'inversion à l'échelle de pays, pas uniquement les plus grands, comme les produits attendus dans le cadre des accords internationaux sur le climat. En revanche, la modestie des améliorations constatées sur la qualité des flux et des champs 3D de traceur et obtenues par l'augmentation de la résolution ou par l'amélioration du cœur dynamique, suggère un examen critique des ressources énergétiques et écologiques mobilisées pour cette montée en résolution au regard des besoins des utilisateurs en aval.

Title: A new model of greenhouse gas transport in the global atmosphere adapted to the evolution of high-performance computing resources

Keywords: High-performance computing, Modeling, Greenhouse gases

Abstract: To understand current climate change, a thorough knowledge of global greenhouse gas emissions and sinks is needed. Atmospheric inverse models play an increasingly significant role in this, but the resolution of global models remains coarse, particularly compared to the increasing density of measurement networks supplying their assimilated observations. New supercomputer architectures, supported by "Manycore" or graphics processors, enable massive parallelization of computations and open up important prospects for inverse models, which we explore in this thesis.

We consider the CAMS/LSCE inversion system. It derives its atmospheric transport model from a general circulation model (GCM) and is driven by a set of meteorological variables pre-generated by it. This GCM has recently integrated a new dynamical core, DYNAMICO, based on a hexagonal tiling of the globe, which is better adapted to parallelism and has been designed to efficiently scale in resolution. We evaluate the new GCM in detail for the transport of carbon dioxide (CO₂) using measurements around the globe's surface and in the atmosphere, and a reference simulation using the classical grid made of rectangles. We show small differences between the two versions, which a comparison to the available observations does not allow us to rank. Our tests also verify the acceleration of calculations on the new grid, particularly at higher resolutions. We prepare to generate the meteorological variables needed for the inverse model, but the complexity of the DYNAMICO code does not allow us to adapt it to the requirements of the inverse model (modularity of components, development of a tangent-linear and adjoint code) in this thesis. In parallel, we investigate the possibility of porting the transport model to graphics cards. We show that the algo-

rithm on the classical Cartesian grid can be expressed in a sufficiently parallel way to run efficiently on such hardware without any simplification despite the continuous exchange of mass between cells. The computation time itself in the forward version of the inverse model becomes less than the time required to read the input meteorological data. Based on our developments with graphic cards, we produce a multi-year inversion at a global resolution of around one degree, assimilating satellite observations. By comparison, the resolution of the system before our developments was around three degrees on the horizontal, with half as many levels on the vertical: its computation times were comparable (here of the order of a week) despite the much smaller number of grid points to be processed. With the new resolution, we show a slight improvement in the representation of horizontal and vertical CO₂ transport, compared with diverse measurements, without being able to attest to a better estimation of surface fluxes.

Taken together, the results demonstrate that there are no longer any limitations on a massive increase in the resolution of global inverse atmospheric models for CO₂, enabling spatial scales reserved until now to limited-area models to be used globally. They open up the prospect of inversion products on the scale of countries, not only the largest, such as the products required by international climate agreements. However, the modest improvement observed in the quality of fluxes and 3D tracer fields obtained by increasing resolution or improving the dynamic core suggests a critical examination of the energy and ecological resources mobilized for this increase in resolution in relation to the needs of downstream users.



Acknowledgements

Au fond de ma révolte contre les forts, je trouve du plus loin qu'il me souvienne l'horreur des tortures infligées aux bêtes. (...) On m'a souvent accusée de plus de sollicitude pour les bêtes que pour les gens : pourquoi s'attendrir sur les brutes quand les êtres raisonnables sont si malheureux ? C'est que tout va ensemble, depuis l'oiseau dont on écrase la couvée jusqu'aux nids humains décimés par la guerre. La bête crève de faim dans son trou, l'homme en meurt au loin des bornes. Et le cœur de la bête est comme le cœur humain, son cerveau est comme le cerveau humain, susceptible de sentir et de comprendre. On a beau marcher dessus, la chaleur et l'étincelle s'y réveillent toujours.

Mémoires – *Louise Michel*

Je tiens avant tout à exprimer ma profonde gratitude à mes deux encadrants, Frédéric Chevallier et Anne Cozic. Lorsque j'ai commencé cette thèse il y a quatre ans, j'étais ingénieure ; aujourd'hui, grâce à votre enseignement, je suis devenue chercheuse. Ce chemin n'a pas été de tout repos, et je vous en suis d'autant plus reconnaissante. Merci de m'avoir offert un environnement de travail où j'ai pu m'épanouir tant sur le plan professionnel que personnel.

Frédéric, ta rigueur et ton efficacité au travail m'ont énormément inspirée et ont été une source d'apprentissage précieuse. Tu as su maintenir un équilibre délicat, en me guidant sans jamais imposer trop de pression, ce qui est loin d'être facile en tant que directeur de thèse. Tu as réussi cet exercice avec brio, et je te remercie profondément pour cela. Je suis fière du travail que nous avons accompli ensemble qui j'espère servira à l'équipe ces prochaines années.

Anne, je te remercie chaleureusement pour ton soutien sans faille et ton aide tout au long de ces années. Merci déjà pour ton expertise technique et ta patience, aborder un modèle comme LMDZ est très intimidant, mais grâce à ton aide, je ne me suis jamais sentie seule ou débordée. Et tu auras beau continuer à répéter que tu n'as pas de doctorat et que la science n'est pas ta spécialité, beaucoup de nos résultats viennent directement de nos discussions scientifiques toujours passionnantes et n'auraient jamais aboutis sans toi.

Enfin, merci à tous les deux pour votre patience et votre flexibilité durant la dernière phase de ma thèse, souvent à distance. Je pense que le résultat final en valait la peine, et je vous suis infiniment reconnaissante pour votre accompagnement jusqu'à ma soutenance.

Merci également à tous mes collègues avec qui j'ai eu la chance de travailler durant ces années. Le travail scientifique est un travail collaboratif, et c'est tout particulièrement le cas quand on utilise des modèles climatiques. Sakina, ça a été un réel plaisir de travailler avec toi. Ton arrivée et notre collaboration sur DYNAMICO a été un vrai regain de motivation pour moi, de même que nos discussions ont ajouté une convivialité à ce travail qui aurait pu être sinon un peu trop solitaire. Un grand merci également à Yann Meurdesoif, pour ton aide technique et ton expertise de DYNAMICO. Sans ton travail préalable cette thèse n'aurait tout simplement pas eu lieu, et sans ton aide elle n'aurait sans doute pas été menée à terme. Merci d'avoir trouvé du temps pour m'aider malgré ton emploi du temps chargé. Merci beaucoup, Marine, pour ton aide précieuse, particulièrement au début de ma thèse. Ton soutien et tes conseils ont été d'une grande aide pour m'orienter et démarrer sur de bonnes bases. Je remercie également les membres de mon comité de thèse, Ehouarn Millour et Jean-Luc Attié, pour leur travail et leur conseils tout au long de ce parcours. Vos retours m'ont permis d'améliorer mes recherches et de progresser après chacun de nos entretiens. Merci à tous mes autres collègues et ami-e-s du labo que je ne peux pas citer d'avoir créé une telle ambiance, surtout après les confinements. Vous m'avez tous et toutes énormément motivé, souvent en partageant nos galères de doctorat.

Un grand merci à ma famille, vous m'avez toujours soutenu dans tous mes projets, toutes mes décisions. Je sais que je pourrais toujours compter sur vous, même quand tout le reste change. Je ne pourrais jamais assez vous remercier pour ça.

Un mot pour Les p'tits lapins. Garder un groupe de potes comme ça pendant plus de 10 ans c'est pas rien, merci à vous d'être toujours là. On est devenus des adultes ensemble, mais malgré le temps et la distance parfois, à chaque fois qu'on se retrouve, c'est comme si rien n'avait changé. Merci au moins pour ça puisque vous remercier pour tout le reste serait sans aucun doute trop long.

Et comment ne pas remercier ma colloc ! Merci pour toutes ces soirées passées à rire et refaire le monde, souvent avec un cocktail d'extrême qualité ou un burger aux rösti de vrai chef. Je ne pense pas que j'aurai survécu à la région parisienne sans ce petit îlot de calme.

Enfin, puisque les jours de grève et de manif sont tout autant nécessaires à un travail bien fait que ceux passés au labo, merci à tous les camarades de lutte. D'ici ou d'ailleurs, dans la rue ou en ligne, de Nanterre à Gaza. Car la lutte écologique ne se fera jamais sans lutte sociale.

And most of all I want to thank my wife, Sophie. It is no exaggeration to say that without you none of this would have been possible. You have been with me every single step of the way, and I will never thank you enough for that. You are my everything, and the home we are building together with our beautiful little bunnies makes all of this worth it. I love you.

In loving memory of Michel, who left us too soon.



Contents

Acknowledgements	i
Contents	iii
List of Figures	viii
List of Tables	x
Introduction	2
Part I General Introduction	3
I Climate change and carbon dioxide	5
1 Climate change and the role of atmospheric CO ₂	6
2 Carbon cycle	8
2.1 Natural carbon cycle	8
2.2 Anthropogenic interference in the carbon cycle	10
II Estimates of carbon dioxide fluxes	13
1 Bottom-up approach	14
2 Top-down approach	15
3 Observations	16
3.1 Direct measurements	16
3.2 Remote sensing	18
3.2.1 TCCON	19
3.2.2 Satellites	19
III Improving atmospheric models	21
1 Why we need better horizontal resolution	22
2 Scaling up models and the need for speed	23
3 Choosing the right grid	25
IV A strategy for the evolution of our atmospheric inverse system	29
1 The state of our current inversion system	30

2	Initial strategy	31
2.1	Initial assessment of the master GCM	31
2.2	Initial step-by-step plan	32
3	Final content of the thesis	32
3.1	Final step-by-step plan	32
Part II Methods		35
V	Atmospheric general circulation models (GCM)	37
1	What makes a general circulation model?	38
1.1	Atmospheric modeling	38
1.2	Eulerian approach	39
2	Coupled configurations of the master GCM	40
3	Direct atmospheric model: LMDZ	41
3.1	Grid of LMDZ	41
3.1.1	Horizontal grid	41
3.1.2	Vertical grid	42
3.2	Dynamics of LMDZ	43
3.3	Physics of LMDZ	44
3.4	Nudging of atmospheric variables	44
3.5	Parallelization of LMDZ	45
4	Icosahedral dynamical core: DYNAMICO	46
4.1	Description of DYNAMICO	46
4.2	Icosahedral grid	47
4.2.1	Horizontal resolution	47
4.3	Parallelization of DYNAMICO	48
5	Chemistry model: INCA	49
6	Land model: ORCHIDEE	49
7	Input/Output module: XIOS	50
VI	Atmospheric inverse modeling: theoretical concept and description of the CAMS/LSCE inverse system	53
1	Concept of atmospheric inverse modeling	54
1.1	Definition of the variables and parameters	54
1.2	Error types	55
1.2.1	Prior error	55
1.2.2	Observation error	55
2	Solving the inverse problem	56
2.1	Bayesian formulation	56
2.2	The Gaussian assumption	56
2.3	Cost function	57
2.4	Adjoint operator	57
2.5	Methods for solving the inverse problem	58
2.5.1	Analytical method	58
2.5.2	Ensemble method	58
2.5.3	Variational method	59
3	The CAMS/LSCE inverse system	59

Part III Scientific findings	61
VII Preliminary developments	63
1 Initial state of the models	64
1.1 Master GCM	64
1.2 Inverse system	64
1.3 Technical challenges	65
1.3.1 Computational cost	65
1.3.2 File size	66
2 Developments	66
2.1 Nudging	66
2.2 Mass fluxes	67
2.3 Tracers	68
2.4 Using XIOS in the inverse system	68
3 Dealing with the icosahedral grid	69
3.1 NetCDF files	69
3.2 Comparison to a regular grid	69
3.3 Drawing maps	70
VIII Direct modeling with DYNAMICO	71
1 Motivation	73
2 Outline of the study	73
2.1 Setup of the study	73
2.2 Summary of the results	74
2.3 Conclusion of the study	75
3 Introduction	76
4 Presentation of the model and experiments	78
4.1 Configuration description	78
4.1.1 General description of the two models	79
4.1.2 ICO - New icosahedral grid configuration	81
4.2 Description of the simulations	82
4.3 Observational data	83
4.4 Evaluation methodology	84
4.4.1 Surface stations	84
4.4.2 Annual gradient between stations	85
4.4.3 Seasonal cycle	85
4.4.4 Synoptic variability	85
4.4.5 ERA5	85
4.4.6 Aircraft measurements	85
4.4.7 AirCore measurements	85
5 Results and discussions	86
5.1 Mass conservation	86
5.2 Computational efficiency	86
5.2.1 Computational setup	86
5.2.2 Computational gains	88
5.3 Vertical temperature profiles	88
5.4 Seasonal analysis	91
5.4.1 Annual gradient	91
5.4.2 Seasonal cycle	94
5.4.3 Synoptic variability	95
5.5 Vertical profiles of CO ₂ mole fractions	96
5.5.1 Troposphere	96

5.5.2	Low stratosphere	96
6	Conclusion	99
7	Code and data availability	99
8	References	100
IX	Prerequisite for high-resolution modeling	107
1	Motivation	108
2	Outline of the developments	108
2.1	Preliminary developments	108
2.2	Summary of the developments	109
3	Introduction	110
4	Model description	111
4.1	Variational framework	111
4.2	Transport model	111
4.3	Direct model	112
4.4	Tangent-linear and adjoint versions	113
4.5	Parallelization strategies	113
5	Exploiting GPU hardware	114
6	Evaluation	114
7	Discussion and conclusions	117
8	Data availability statement	117
9	References	118
X	High-resolution inversions	119
1	Motivation and strategy for the evolution of the inverse system	120
2	Outline of the study	121
2.1	Setup of the study	121
2.2	Summary of the results	122
2.3	Conclusion of the study	123
3	Key points	124
4	Abstract	125
5	Plain Language Summary	125
6	Introduction	125
7	Model and inversion setup	127
7.1	Inversion system	127
7.2	Inversion configuration	128
7.3	Evaluation	129
7.3.1	CO ₂ data for evaluation	129
7.3.2	Processing of the surface stations	131
7.3.3	Processing of the column-averaged CO ₂ and vertical profiles	132
7.3.4	Processing of the surface flux estimates	132
8	Results and discussion	133
8.1	Surface stations	133
8.2	TCCON observations	136
8.3	Vertical profiles	138
8.4	Regional fluxes	142
8.5	Local fluxes	145
9	Conclusion	146
10	Data availability statement	148
11	References	148
12	Appendix A : Observation datasets	156

Part IV Conclusions and perspectives	169
Conclusion	171
Perspectives	173
1 The future of the CAMS/LSCE inverse system	173
1.1 Inverse system using DYNAMICO	173
1.1.1 Current state of the inverse system	173
1.1.2 Possible optimization	174
1.2 Bypassing the master GCM	175
2 Lessons on increasing resolution	175
2.1 GPU: an opportunity not to be missed	175
2.2 How to choose the grid?	176
2.3 Technical challenges	176
3 Could AI's reign on global atmospheric transport be over before it even started?	177
4 Is increasing the resolution worth the pain?	177
4.1 Modest benefits so far	178
4.2 A non-negligible ecological cost	178
4.3 A paradigm change in the future?	178
Bibliography	181
Bibliography	181
Appendices	190
A Acronyms and common notations	193
B List of communications	197

List of Figures

I.1	Monthly mean carbon dioxide measured at Mauna Loa Observatory	7
I.2	Observed global mean surface temperature change	8
I.3	IPCC global carbon budget	9
I.4	Combined components of the global carbon budget	11
II.1	Methods of CO ₂ fluxes estimation	15
II.2	CO ₂ fluxes per latitude band	17
II.3	Map of the TCCON network	19
III.1	Pole convergence	26
IV.1	Architecture of the CAMS/LSCE inverse system	31
V.1	Coupled configurations of the master GCM	40
V.2	Staggered Arakawa C-grid	42
V.3	Vertical grid of LMDZ	43
V.4	Parallelization of the dynamics of LMDZ	46
V.5	Icosahedral grid construction	47
V.6	Mesh of DYNAMICO	48
V.7	Tiles for parallelizing DYNAMICO	48
V.8	Architecture of XIOS	50
VI.1	Schematic representation of the use of Bayes' formula to solve the inverse problem	57
VII.1	Comparison of a CO ₂ mixing ratio field on a native and regular grid	70
VIII.1	Schematic of the structure of the two coupled configurations REG and ICO	79
VIII.2	Icosahedral grid of the ICO configuration	81
VIII.3	Subdivision of the primary mesh	82
VIII.4	Vertical profile of zonal temperatures	89
VIII.5	Time series of the average difference in zonal temperature at the stratopause	90
VIII.6	Annual gradient of CO ₂ mole fraction compared to SPO	91
VIII.7	Normalized standard deviation of the annual gradient	92
VIII.8	Bias per station and average bias per latitude band of the annual gradient of CO ₂ mole fraction	93
VIII.9	Seasonal cycle of the CO ₂ mole fraction	94
VIII.10	Pearson correlation and normalized standard deviation of the seasonal cycle	95

VIII.11	Pearson correlation and normalized standard deviation of the daily average residue . . .	95
VIII.12	Seasonal and annual means of CO ₂ vertical profiles compared to aircraft measurements	96
VIII.13	Difference in CO ₂ mole fraction vertical profiles compared to AirCore measurements . . .	98
IX.1	Difference in simulations of CO ₂ between MPI and GPU versions	116
X.1	Map of the location of selected surface stations, TCCON sites, and AirCore flights	131
X.2	Pearson correlation coefficient and average normalized peak-to-peak amplitude of the seasonal cycle for the two resolutions of the inverse model	133
X.3	Pearson correlation coefficient and average normalized peak-to-peak amplitude of the daily average residue for the two resolutions of the inverse model	134
X.4	Correlation and normalized standard deviation of the difference between the modeled XCO ₂ and TCCON data for the two resolutions of the inverse model	136
X.5	Pearson correlation and average normalized peak-to-peak amplitude of the seasonal cycle at TCCON stations for the two resolutions of the inverse model	137
X.6	Vertical profile of the CO ₂ mole fraction for the two resolutions of the model compared to AirCore measurements	138
X.7	Vertical profile of the difference in CO ₂ mole fraction between the HR and LR models averaged over the two years and per longitude band, with corresponding zonal XCO ₂	140
X.8	Maps of XCO ₂ difference after inversion between the two resolutions, averaged per year and season	141
X.9	Total annual surface emissions for the LR and HR inverse model for each Transcom3 region	143
X.10	Monthly averaged surface flux for the LR and HR inverse model for four significant Transcom3 regions	144
X.11	Surface flux increments between the prior and posterior state of the inversion for both resolutions of the model	145
X.12	Total monthly surface flux for the LR and HR inverse model, zoomed near the LAN station in China	146



List of Tables

III.1	Models used in the Global Carbon Budget 2022	24
III.2	Characteristics of satellite missions	25
VII.1	File size of the GCM simulations.	66
VIII.1	Computational setup and results of the simulations	87
IX.1	Four transformations of the code for an efficient use by the GPU	115
X.1	Notable stations identified by seasonal and synoptic variability performance	135
X.2	Total annual estimates of natural CO ₂ fluxes for the LR and HR inverse model at the global scale	142
X.3	Appendix A1: List of datasets used from Obspack for surface stations	156
X.4	Appendix A2: List of TCCON sites used and their locations	167
A.1	Acronyms and abbreviations used in this report	193
A.2	Chemical & mathematical symbols and physical units	195



Introduction

Global climate change poses a severe threat to human societies and biodiversity, necessitating urgent mitigation efforts to limit its impacts. The rapid increase of global temperatures, in particular, can largely be attributed to the rising concentrations of greenhouse gases of anthropogenic origin in the atmosphere. Among those greenhouse gases, carbon dioxide (CO₂) is the most significant, despite constituting only about 0.04% of the Earth's atmosphere by volume. Its concentration has increased by more than 50 % since the Industrial Revolution, and has already caused the global average surface temperature to rise by about 1.05 °C. Atmospheric carbon dioxide represents only a small fraction of a complex natural carbon cycle, and monitoring the sources and sinks of carbon at the surface of the Earth can help us better understand climate change and monitor the efficacy of mitigation efforts.

There are two types of methods for estimating the surface fluxes of greenhouse gases: the bottom-up approach and the top-down approach. The bottom-up approach extrapolates observations and models of local fluxes to a larger scale, relying on detailed process-based models and inventories. This method compiles data from various sources, such as fossil fuel emission records, land use changes, and ecosystem models, to estimate the overall fluxes. The top-down approach, which is the focus of this work, uses inverse models to estimate surface fluxes. This method involves statistical techniques that combine prior estimates of surface fluxes with atmospheric observations of greenhouse gas concentrations. The inverse models optimize the flux estimates to better match the observed concentration data depending on their uncertainties. This approach can provide a more integrated view of the carbon cycle by constraining flux estimates with actual atmospheric measurements. However, their relatively low resolution at the global scale hinders their ability to capture the spatial variability of fluxes accurately and to make better use of high-resolution satellite observations.

Supercomputers used for atmospheric inversions are rapidly evolving, with architectures increasingly incorporating "Manycore" processors and Graphics Processing Units (GPUs). These advancements enable massive parallelization of computations, which should be leveraged to enhance our models. The use of GPUs, in particular, can significantly accelerate the computations required for high-resolution inverse models, making it feasible to run these models at much finer scales than previously possible.

In this work, we investigated two complementary approaches to improve the inverse system used to generate global CO₂ inversion products for the Copernicus Atmosphere Monitoring Service (CAMS). The first approach involves using DYNAMICO, a dynamical core based on an hexagonal tiling of the globe. DYNAMICO offers several advantages over a regular latitude-longitude grid, particularly around the poles. The second approach focuses on increasing the spatial resolution of the inverse model itself. Higher resolution inversions can provide more detailed and accurate estimates of surface fluxes

by capturing finer-scale processes and variability. These methods will ensure that this inverse system scales efficiently in the future and stays up to date with technological innovations. By leveraging the capabilities of modern supercomputers and advanced modeling techniques, we aim to produce more accurate and detailed global estimates of CO₂ fluxes.

This manuscript is divided into three main parts:

The first part is a general introduction in four chapters. The first one gives a general overview of the causes of modern global climate change, with particular emphasis on the role of atmospheric carbon dioxide, as well as a presentation of the global carbon cycle. The second chapter presents the different ways that can be used to estimate surface fluxes of carbon dioxide and the observations that make these methods possible. The third chapter presents an overview of the reasons for improving the resolution of an atmospheric model, the challenges this creates and how they can be mitigated. The final chapter presents our strategy for improving the atmospheric carbon dioxide inverse system studied in this work.

The second part presents in two chapters a detailed description of the models used in this work. The first chapter focuses on a general circulation model, its different components, and the new configuration created for this work. The second chapter presents our atmospheric inverse system, its overall architecture as well as the data products and observations that make it operational.

The third part presents all the scientific findings and developments that have been carried out as part of this thesis. It consists of four chapters, the first of which describes in detail the preliminary developments necessary to carry out the subsequent scientific studies. Following a brief description of the state of the models at the beginning of this thesis, the technical challenges that came up during the course of this work are also discussed. The developments that were required in the models are described, as well as some of the tools that have been used. The second chapter consists of a study looking at the impact on atmospheric transport of carbon dioxide when using a new dynamical core in our general circulation model. To do so, we performed a comparison of two configurations of our model in a 40-year-long simulation, as well as other computational tests. The third chapter presents a study on the porting of our inverse system to graphical processing units, a work that was crucial for the following chapter. The fourth and final chapter presents a study of the effects of increasing the spatial resolution of our inverse system, using a 2-year inversion assimilating satellite data.

In the final chapter of this thesis, the main conclusions are summarized and an outlook on future developments of our model is given, including a general takeaway of this work.

Part I

General Introduction

Climate change and carbon dioxide

Pour le peuple colonisé la valeur la plus essentielle, parce que la plus concrète, c'est d'abord la terre: la terre qui doit assurer le pain et, bien sûr, la dignité.

Les damnés de la terre – *Frantz Fanon*

Contents

1	Climate change and the role of atmospheric CO ₂	6
2	Carbon cycle	8
2.1	Natural carbon cycle	8
2.2	Anthropogenic interference in the carbon cycle	10

CARBON DIOXIDE (CO₂) is a molecule at the heart of global climate change. It has been the subject of numerous scientific studies over the last few decades and is the motivation behind the present study.

This first chapter introduces the basic physical processes that cause climate change on Earth and the role that atmospheric CO₂ plays in this process. The characteristics of atmospheric CO₂ and its impact on the global climate are described in [section 1](#). [Section 2](#) introduces the Earth's natural carbon cycle and how human activities have affected it.

1 Climate change and the role of atmospheric CO₂

Carbon dioxide (CO₂) holds historical significance as the first gas discovered by a scientist, Jan Baptist van Helmont, around 1640. Initially thought to be exclusive to charcoal combustion, it was later isolated by physicist Joseph Black in the 18th century. He achieved this by reacting calcium carbonate with acids and explained its occurrence from animal and microbial respiration. Only during the 18th century was its composition into basic elements (carbon and oxygen) clarified, and the term "carbon dioxide" became commonly used in the 20th century.

The CO₂ molecule is composed of one carbon atom (C) and two oxygen atoms (O) arranged linearly with double covalent bonds. At atmospheric pressures and temperatures, CO₂ exists as a gas. Despite its relatively low concentration in the Earth's atmosphere, only 0.04% of the atmosphere's volume, CO₂ plays a vital role in regulating the planet's temperature.

The presence of certain greenhouse gases, including water vapor (H₂O), carbon dioxide (CO₂), methane (CH₄), and nitrous oxide (N₂O), is essential in maintaining a habitable temperature on Earth. Without them, the average surface temperature would be as low as $-18\text{ }^{\circ}\text{C}$ (Lacis et al., 2010), making life as we know it impossible. The majority of the energy received by the Earth comes in the form of short-wave radiation from the sun: on average around $340\text{ W}\cdot\text{m}^{-2}$ annually. Approximately 30% of the short-wave radiation is directly reflected back into space, while the remaining $240\text{ W}\cdot\text{m}^{-2}$ is absorbed by the atmosphere and the Earth's surface (Wild et al., 2013). The heated surface emits this energy back in the form of outgoing long-wave radiation. Some of it is radiated directly into space, and some is reabsorbed by greenhouse gases in the atmosphere. These greenhouse gases re-emit part of this thermal radiation downward, contributing to a so-called greenhouse effect and further warming the Earth's surface.

Approximately 75% of the greenhouse effect is attributable to water vapor and clouds, which are primarily influenced by natural processes on a global scale. In contrast, CO₂ contributes to around 20% of the greenhouse effect in the present day (Schmidt et al., 2010).

Ice core measurements have shown that CO₂ concentrations have fluctuated significantly over the past 800,000 years, ranging from approximately 180 parts per million (ppm) at its lowest point 740,000 years ago to 300 ppm around 350,000 years ago (Lüthi et al., 2008). The pre-industrial levels in the 1700s were around 280 ppm and human activity at large scales has always impacted atmospheric CO₂ concentrations. For example, events such as the colonization of the Americas, the resulting deaths, and the subsequent large-scale reorganization of land use even lead to rapid declines in CO₂ concentrations of several ppm over the course of a century (Koch et al., 2019). However, since the Industrial Revolution (usually dated as the year 1750 for climate studies), the combustion of fossil fuels (coal, oil, natural gas), deforestation, and land use changes have substantially increased global CO₂ concentrations. As of March 2024, the global monthly mean CO₂ concentration reached 423.16 ppm (Figure I.1, page 7, Lan, Tans, and Kirk, 2024b).

In addition to CO₂, human activity has also led to notable increases in the concentrations of other greenhouse gases. As of March 2024, CH₄ concentrations have doubled, while N₂O concentrations have increased by more than 38% compared to pre-industrial levels (NOAA, Lan, Tans, and Kirk, 2024a).




Figure I.1: Time serie of atmospheric carbon dioxide mole fraction (ppm) from measurements acquired at Mauna Loa since 1957 Lan, Tans, and Kirk, 2024b

While aerosols and their interaction with clouds can have a negative impact on radiative forcing, they have also contributed to limiting it at both global and local scales. Aerosols, whether of natural (e.g., volcanic) or anthropogenic origin, can significantly influence the climate for extended periods (Schurer, Tett, and Hegerl, 2014). This has prompted research into large-scale geoengineering methods to reduce total radiative forcing (Kravitz et al., 2015). However, such approaches carry potential risks of unforeseen and adverse effects on climate, weather patterns, and human health (Richter et al., 2017, Tilmes et al., 2022). Consequently, the preferred approach of intergovernmental bodies to mitigate these changes remains the reduction of anthropogenic greenhouse gas emissions.

The escalating rise in greenhouse gas concentrations leads to a rapid increase in global temperatures, posing severe threats to human life, society, and biodiversity. According to the Intergovernmental Panel on Climate Change (IPCC), the global mean surface temperature of Earth during the period 2011-2019 was approximately 1.07°C higher than the pre-industrial era (1850-1900) (Figure I.2, page 8).

The international community adopted the Paris Agreement in 2015 to respond to this escalating climate crisis. This treaty aims to limit the global temperature rise to below 2°C and preferably below 1.5°C by 2100. At the time of writing, 195 entities have signed and ratified the agreement.

To achieve the ambitious goal of limiting the temperature increase to 1.5°C, countries would need to significantly reduce their greenhouse gas emissions, aiming for a 50% cut by 2030. Such an achievement necessitates transformative changes in human societies and production methods, as it requires a concerted effort across the globe to transition towards sustainable and low-carbon practices.




Figure I.2: Observed global mean surface temperature from four datasets, relative to the average temperature of 1850–1900 in each dataset (Figure 1.12 in Intergovernmental Panel On Climate Change, 2023)

2 Carbon cycle

Carbon dioxide (CO₂) serves as the primary driver of anthropogenic greenhouse gas emissions, making it essential to comprehend both its natural cycle and the influence of human activities on its concentration. Achieving a comprehensive understanding of these factors is critical in gaining insights into climate change dynamics on a global scale and formulating effective mitigation strategies.

2.1 Natural carbon cycle

While we focus on atmospheric CO₂, carbon follows a long cycle and takes different forms depending on where it is stored on Earth. In the atmosphere, it is mainly a gas; on land, it is primarily stocked in solid form, whereas in the ocean, it is primarily in a dissolved form. Stocks of carbon are often expressed either in gigatons of carbon (GtC) or petagrams (10¹⁵g) of carbon (PgC), which are equivalent units. Carbon fluxes at the global scale for the natural carbon cycle are usually expressed in PgC per year (PgC/yr).

Recent estimates of pre-industrial carbon stocks (Figure I.3) show that the oceans harbor the vast majority of Earth's carbon, with 37,100 PgC mainly dissolved in the intermediate and deep sea, manifesting as CO₂, carbonate, and bicarbonate. Additional carbon stocks exist in ocean floor sediment (1750 PgC), dissolved organic carbon from the decomposition of dead organic matter (700 PgC), and the surface of the oceans (900 PgC). Comparatively, the living ocean biosphere, including flora and fauna, accounts for only 3 PgC.

The second largest reservoir of carbon is found in terrestrial ecosystems, distributed among soils (1700 PgC), permafrost (1200 PgC), and vegetation (450 PgC). Additionally, a portion of the carbon is stored as fossil fuel reserves, encompassing coal, oil, and gas (928 PgC).

Lastly, a relatively small proportion of the total carbon on Earth (591 PgC) is confined to the atmosphere, primarily as CO₂, with smaller contributions from CO and CH₄.




Figure I.3: Global carbon (CO₂) budget (2010–2019). Yellow arrows represent annual carbon fluxes (in PgC yr⁻¹) associated with the natural carbon cycle, estimated for the time prior to the industrial era, around 1750. Pink arrows represent anthropogenic fluxes averaged over the period 2010–2019. Circles with yellow numbers represent pre-industrial carbon stocks in PgC. Circles with pink numbers represent anthropogenic changes to these stocks (cumulative anthropogenic fluxes) since 1750. (Figure 5.12 in Intergovernmental Panel On Climate Change, 2023)

The observed carbon stocks in various reservoirs result from exchanges between these reservoirs occurring at different timescales. The pre-industrial natural fluxes are represented by yellow arrows in Figure I.3 (Intergovernmental Panel On Climate Change, 2023). All the natural fluxes discussed in the following correspond to the values evaluated for the pre-industrial period.

On a yearly timescale, direct fluxes between land and ocean are relatively low, amounting to 0.8 PgC/yr, with most carbon fluxes at a rapid rate occurring at the atmospheric interface. At the land-atmosphere interface, gross photosynthesis stands out as the most substantial land sink of carbon, sequestering approximately 113 PgC/yr and supporting plant and tree growth. However, respiration from plants, animals, and natural forest fires act as a carbon source, emitting about 111.1 PgC/yr into the atmosphere. Additionally, freshwater areas contribute as a net carbon source, releasing approximately 1.5 PgC/yr to the atmosphere.

The ocean-atmosphere interface is responsible for large-scale gas exchange through diffusion, facilitated by differences in the partial pressure of CO₂ between the ocean and the atmosphere. This exchange rate depends on various variables, such as ocean surface wind speed, temperature, salinity, and atmospheric CO₂ concentrations. This process results in a natural net flux of 0.6 PgC/yr from the ocean to the atmosphere. Understanding these natural fluxes and their dynamics is fundamental for comprehensively assessing the global carbon cycle and its implications for the Earth's climate system.

2.2 Anthropogenic interference in the carbon cycle

The estimates of carbon stocks and fluxes mentioned earlier were based on the conditions around 1750 before significant human activities altered them. Over hundreds of millions of years, vast amounts of carbon from dead organic matter have been sequestered in fossil fuels such as coal (300 million years ago) and oil and gas (over dozens of millions of years) (Dai et al., 2022). This created substantial carbon reservoirs that remained isolated from the atmosphere and the fast carbon cycle relevant to our human timescale. However, with the burning of fossil fuels, the carbon they contain is released into the atmosphere, mainly as CO₂, increasing atmospheric carbon levels. This, in turn, impacts the fluxes between land, ocean, and atmosphere. The changes in carbon stocks due to human activities are depicted as pink circles in Figure I.3, page 9, along with the corresponding changes in fluxes shown by pink arrows.

As of 2019 (Intergovernmental Panel On Climate Change, 2023), approximately 445 PgC have been extracted from fossil fuel reserves, and 173 PgC have been absorbed by the ocean, while an additional 279 PgC has accumulated in the atmosphere as evidenced by atmospheric CO₂ amounts in Figure I.3, page 9. Human-induced land-use changes, such as deforestation and agriculture, also affect the carbon cycle. As of 2019, these activities have contributed an extra net flux of 1.6 PgC/yr to the atmosphere, resulting in a total depletion of 25 PgC from vegetation and soils. While net fluxes between the atmosphere, land, and oceans have generally increased, providing a more effective sink for atmospheric carbon (3.4 PgC/yr for the net land flux and 2.5 PgC/yr for the net ocean flux), this increase is insufficient to offset the fluxes from fossil fuels (9.4 PgC/yr) and land use.

The intricate interplay between natural and anthropogenic sources and sinks of carbon has highlighted the urgency of comprehending the carbon cycle to control our emissions effectively. The Global Carbon Project (<https://www.globalcarbonproject.org/>) was established in 2001 in response to this critical need. Since 2007, the Global Carbon Budget (<https://globalcarbonbudget.org/>) has been producing yearly estimates of natural and anthropological carbon emissions and natural sinks for monitoring purposes. The primary objective of this annual budget is to gain insights into the patterns and variability of global carbon fluxes, understand the feedback mechanisms of the carbon cycle, and identify opportunities for human management of this vital cycle. Figure I.4, page 11 from the Global Carbon Budget 2022 (Friedlingstein et al., 2022) illustrates the annual variability of carbon emissions and sinks. Panel (a) presents the yearly estimates of each flux, highlighting their considerable inter-annual variability and the increasing share of fossil emissions in more recent years. Panel (b) presents the partitioning of the cumulative fluxes since 1850 and shows that the fossil CO₂ emissions have already become the most important source of atmospheric CO₂. Nevertheless, it is essential to acknowledge that the values of the total amount of carbon stored and the different fluxes are subject to uncertainties that can be quite significant, up to ± 0.7 GtC/yr (1σ) for the land use for example. Improving the accuracy of these estimates is crucial in enhancing our understanding of human impact on the carbon cycle and devising effective strategies for mitigating climate change.




Figure I.4: Combined components of the global carbon budget as a function of time for fossil CO₂ emissions (E_{FOS}), emissions from land-use change (E_{LUC}), as well as their partitioning among the atmosphere (G_{ATM}), ocean (S_{OCEAN}), and land (S_{LAND}). Panel (a) shows annual estimates of each flux, and panel (b) shows the cumulative flux (the sum of all prior annual fluxes) since the year 1850. All data are in GtC per year (a) and GtC (b) (Friedlingstein et al., 2022)

Estimates of carbon dioxide fluxes

Contents

1	Bottom-up approach	14
2	Top-down approach	15
3	Observations	16
3.1	Direct measurements	16
3.2	Remote sensing	18
3.2.1	TCCON	19
3.2.2	Satellites	19

AS SEEN IN THE PREVIOUS CHAPTER, precisely estimating CO₂ emissions both in time and space helps to understand anthropogenic climate change and to drive policy-making decisions. Monitoring the sensitivity of natural sources and sinks of carbon in response to our changing climate is also needed, as it aids in refining models and predicting future emission pathways.

There are two generic ways of estimating CO₂ surface fluxes: the bottom-up and the top-down approaches. Despite inherent uncertainties, these approaches offer distinct yet complementary perspectives that together offer a powerful means to cross-validate results and gain a more robust understanding of carbon fluxes. Integrating the strengths of bottom-up and top-down approaches allows for a more comprehensive assessment of the complex dynamics driving CO₂ emissions.

The following chapter presents the different ways of estimating CO₂ surface fluxes and the observations used to make it possible. In [section 1](#), the bottom-up approach is described. The top-down approach is then explained in [section 2](#). Finally, the observations used for inverse modeling are described in [section 3](#). A schematic summary of these two approaches is presented in [Figure II.1](#) on page 15.

1 Bottom-up approach

The bottom-up approach encompasses various estimation techniques, ranging from straightforward assessments of fossil fuel usage and land use to complex models that simulate intricate chemical and biological processes. To obtain regional or global estimates, the bottom-up approach often involves extrapolating local observations or processes to larger scales, which may introduce uncertainties. This is necessary when direct measurements or comprehensive data at broader scales are not feasible. The Global Carbon Budget 2022 (Friedlingstein et al., 2022), as presented in [chapter I, subsection 2.2](#), uses a comprehensive bottom-up approach by combining these different methods. It compiles data from multiple sources and models to estimate carbon emissions and sinks on a global scale.

Their estimation of global and national fossil CO₂ emissions (E_{FOS} in [Figure I.4](#), page 11) involves the integration of multiple sources of economic and energy usage activity from numerous countries, with specific emissions factors assigned to each. Special care is taken to avoid double counting from various sources (Andrew, 2020), and the calculations account for processes such as CO₂ capture through cement carbonation. A similar method is used to get CO₂ fluxes from land use (E_{LUC} in [Figure I.4](#), page 11) such as deforestation or cultivation. Bookkeeping approaches were used to calculate sources and sinks and find the overall net CO₂ flux from these processes (Hansis, Davis, and Pongratz, 2015, Gasser et al., 2020, Houghton and Nassikas, 2017). The terrestrial land sink (S_{LAND}) is obtained from the mean value of 16 dynamic global vegetation models (DGVMs), which account for the impact of climate variability and CO₂ concentrations on land. This value does not include the land sinks from land use already counted in E_{LUC} , but differentiating them can be challenging (Erb et al., 2013). For the estimation of the ocean CO₂ sink (S_{OCEAN}), two values are averaged. The first value is derived from the mean of 10 global ocean biochemistry models, which simulate anthropogenic and natural CO₂ cycles and carbon transport from the atmosphere to the ocean interior. These models were rigorously evaluated against independent observations for validation (Hauck et al., 2020). The second value is obtained from the mean of seven so-called data products based on measurements of the fugacity of CO₂ (partial pressure corrected for non-ideal gas behavior) in the ocean (Pfeil et al., 2013). The assessment of sources and sinks must be compared to the growth rate of atmospheric CO₂ concentration (G_{ATM}), which is provided by the US National Oceanic and Atmospheric Administration Global Monitoring Laboratory (NOAA/GML). G_{ATM} is derived from the average of multiple marine stations situated in the atmosphere's boundary layer across the globe. These stations provide measurements distant from local emissions sources and represent well-mixed air. The estimation of G_{ATM} is considered highly reliable due to the extensive and consistent observations worldwide (Bal-

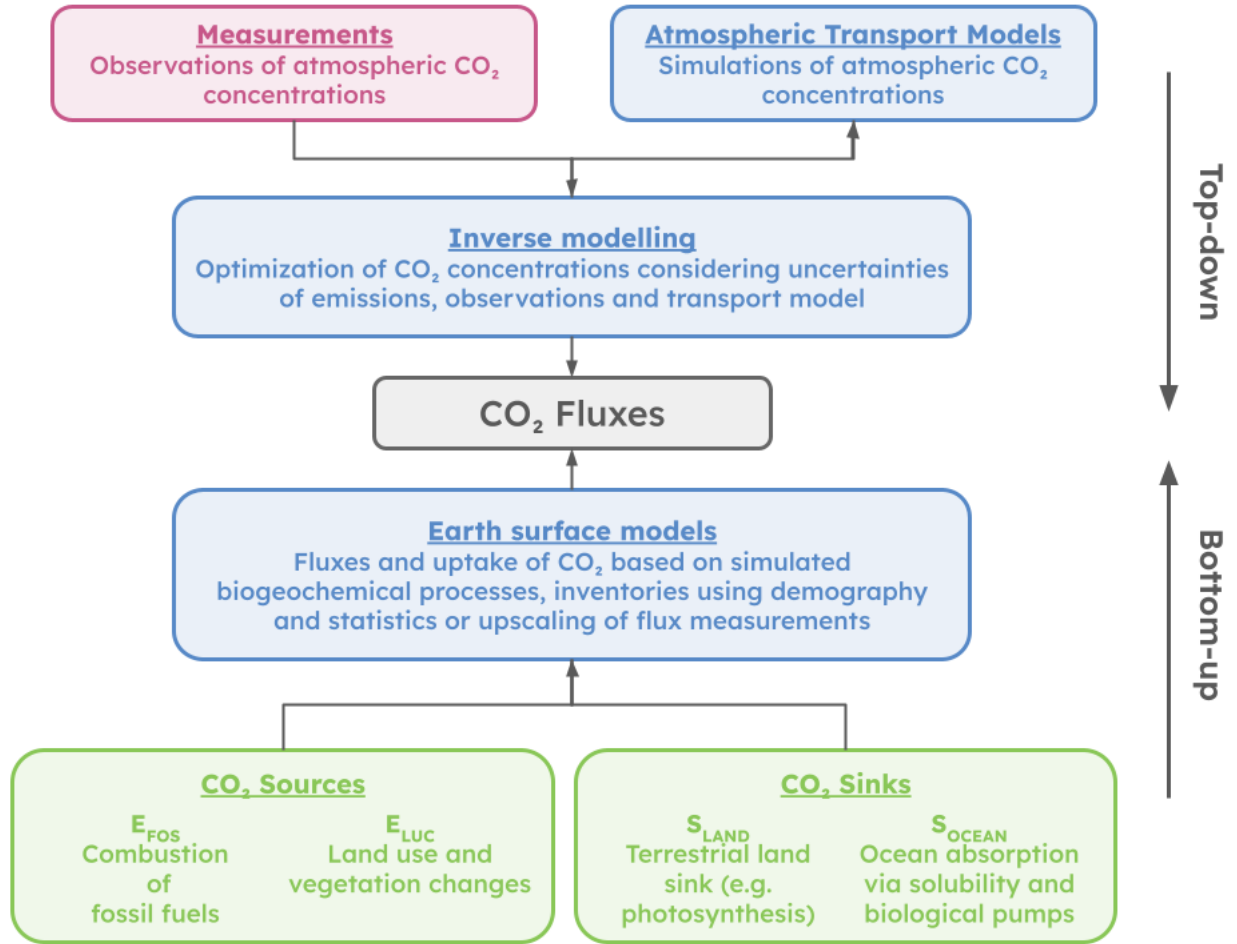


Figure II.1: Schematic illustration of top-down and bottom-up approaches to estimate CO₂ fluxes.

lantyne et al., 2012). Uncertainties for the values of these fluxes are derived from the range of values of the biochemical models or observations.

Over the entire 1850-2021 period, the cumulative budget imbalance (BIM, see Equation II.1) amounts to 15 GtC or approximately 2% of total emissions. Most of this imbalance likely stems from overestimates in land and fossil fuel use during the 1920-1960. However, for the more recent 1960-2021 period, the imbalance was close to zero in fluxes, averaging at 0.07 GtC/yr. Notably, since the atmospheric growth rate did not constrain these estimates, this suggests a good understanding of emissions and their partitioning over these longer time scales.

$$B_{IM} = E_{FOS} + E_{LUC} - (G_{ATM} + S_{OCEAN} + S_{LAND}) \tag{II.1}$$

Nonetheless, the variability at lower time scales remains significant, and further efforts are required to precisely attribute these variations to specific errors in the estimates of different carbon sources and sinks.

2 Top-down approach

A top-down approach, in our case atmospheric inverse modeling, improves prior estimates of CO₂ fluxes by using statistical methods combining these prior estimates with simulations and observations of atmospheric CO₂ concentrations together. Instead of modeling the intricate anthropogenic and biological processes governing CO₂ sources and sinks on the Earth’s surface, these models try to find

the optimal simulated fluxes of CO_2 that best explain the observations used as input (Ciais et al., 2010). The observations used consist of atmospheric CO_2 concentrations, either in fixed points from in-situ measurements or from remote sensing using columns of CO_2 from satellites for example (Chevallier et al., 2019). Inversions start from “prior” states that consist of fields of CO_2 fluxes generally derived from a bottom-up approach. These fields are then used to drive a simulation using an atmospheric transport model in order to give a first estimation of atmospheric CO_2 concentrations. A numerical scheme then uses the differences between this first simulated concentration field and the observations to be assimilated to adjust the prior fluxes. This inversion process results in a state that best agrees with the prior fluxes and the observations, considering their respective uncertainties. This final posterior state of CO_2 fluxes has a lower uncertainty than the original prior fluxes.

Atmospheric inversions can be done at a global scale (Chandra et al., 2022), restricted to certain regions (Monteil et al., 2020), or even at a local scale (Bréon et al., 2015). Their use and study have led to a better understanding of the sources and sinks of atmospheric carbon and their changes over time. By design, global atmospheric inversions close the budget imbalance discussed above ($B_{IM} = 0$). The growth rate of the atmospheric CO_2 concentration (G_{ATM}) is very well constrained by the observation network, and E_{FOS} is assumed to be well known, the distribution of fluxes is then solved across the land and ocean. They can thus give another view of these fluxes and their uncertainties. At a global scale, studies of multiple inverse models like in the Global Carbon Budget 2022 show good agreements with the bottom-up approach for total atmosphere-to-land and atmosphere-to-ocean fluxes (Figure II.2, page 17). However, this is not the case when looking at different regions or timescales. For example, inverse models suggest higher estimates of the total land flux in northern latitudes ($> 30^\circ \text{N}$) for 2012 - 2021 than the DGVMs provide (0.6 to 2.0 GtC/yr for inversions versus 1.0 ± 0.4 GtC/yr). There are also differences in the results given by different inverse models since they differ in their transport model, prior fluxes, and assimilated observations (Basu et al., 2018, Schuh et al., 2019). Therefore, improving these models can help to better understand the global carbon cycle.

3 Observations

The global inverse models briefly described in the last section necessitate comprehensive measurements of atmospheric CO_2 concentrations at the global scale. For atmospheric inversions these measurements serve three primary purposes: first, for running the model and optimizing the fluxes, second, for validating the outcomes of these inversions and finally for validating the transport models themselves. Two distinct measurement methods are employed: direct measurements, which directly assess the CO_2 concentration at specific observation points, and remote sensing, which involves analyzing wavelengths of light that are then correlated to CO_2 concentrations.

Remote sensing techniques can be implemented using either satellite-based platforms or ground-based stations. Unlike direct measurements that provide localized data, remote sensing offers the advantage of capturing entire vertical columns of CO_2 in the atmosphere at once. This broader scope of observation facilitates a more comprehensive understanding of the spatial distribution and variations in CO_2 concentrations across the globe. We will first present direct observations and ground remote sensing methods used to validate our models before looking at the specificities of satellite-based CO_2 remote sensing.

3.1 Direct measurements

Atmospheric CO_2 concentrations vary both in space and time around the globe, at large scale, or locally when close to cities or other emission sources. They follow seasonal but also diurnal cycles and vary in the vertical space. Therefore, an extensive and robust observation network is required to cover these variations. In-situ measurements are varied and can be taken from static surface stations, aircraft, or even instruments launched by high-altitude balloons. These measurements are often col-

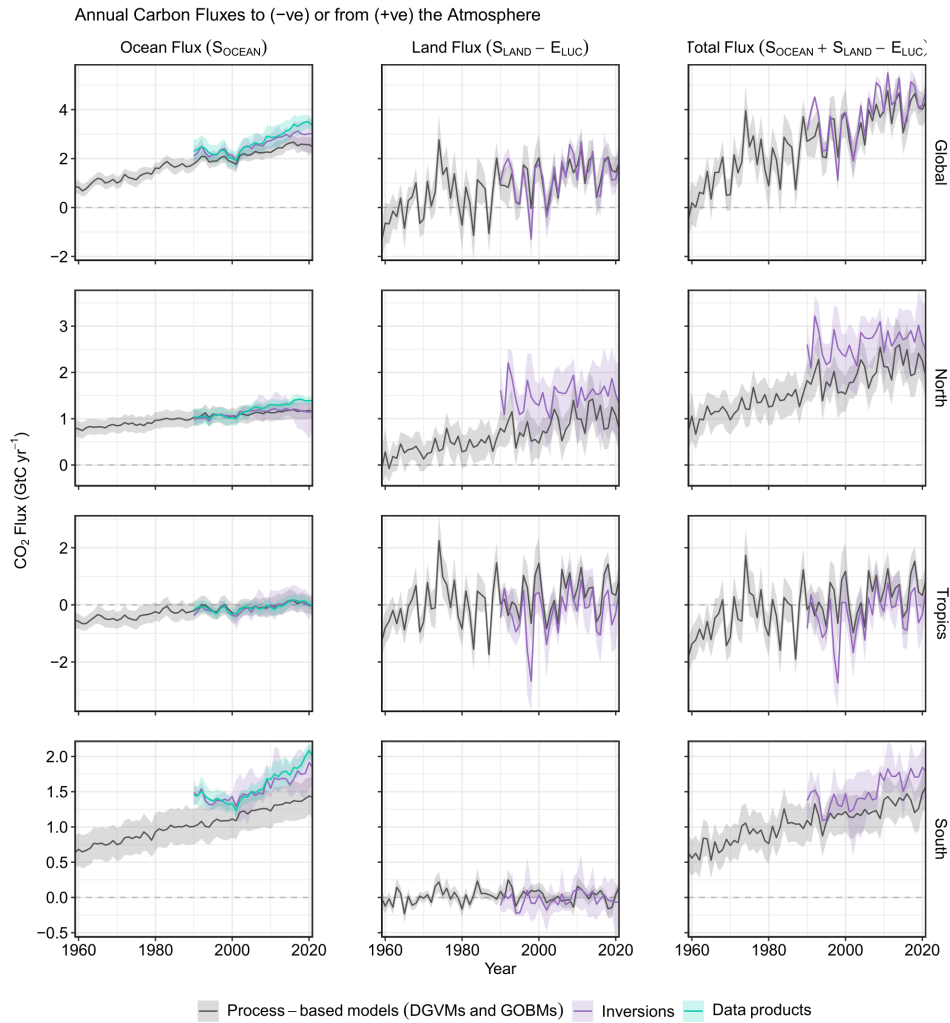


Figure II.2: CO₂ fluxes between the atmosphere and the Earth’s surface separated between land and oceans globally and in three latitude bands. The ocean flux is S_{OCEAN} , and the land flux is the net of atmosphere–land fluxes from the DGVMs. The latitude bands are (top row) global, (second row) north ($> 30^\circ \text{ N}$), (third row) tropics (30° S – 30° N), and (bottom row) south ($< 30^\circ \text{ S}$), showing values over ocean (left column) and land (middle column) and in total (right column). Estimates are shown for process-based models, inversion systems, and $f\text{CO}_2$ -based data products (Friedlingstein et al., 2022).

lected and aggregated by organizations such as the NOAA and its Observation Package (ObsPack) data products for easier use by the scientific community.

Surface stations can measure atmospheric CO₂ concentrations with two methods: either from sites where instruments provide continuous measurements or flask samples taken at the site and analyzed at a later time in a central laboratory. Having such observation sites all around the globe is necessary, but installing and maintaining these sites in some regions is often complicated. This leads to a disparity in the distribution of these sites, with the majority being in the northern hemisphere and, most notably, Western Europe and North America, with only a limited number in the tropics and the southern hemisphere. Most of these instruments use the technique of cavity ring-down spectroscopy, which has a very high precision: better than 0.1 ppm(1 σ) for hourly mean data (Wu et al., 2016). When properly calibrated every 2 weeks or 2 months, these instruments do not show any significant bias. However this type of instruments typically cost around EUR 50 k per unit and this cost limits the size of observation networks. Some instruments can measure concentrations of atmospheric CO₂ for much cheaper, but the accuracy of these CO₂ sensors will be lower. The minimum suggested accuracy for sensors to be used in an effective urban observation network is 1 ppm on hourly mean measurement (Wu et al., 2016). This could be achieved with low-cost medium precision sensors for around EUR 1 k per unit, and creation and deployments of such sensors are currently underway (Lian et al., 2024).

Measurements taken on aircraft can give information on the spatial distribution of CO₂ that surface stations cannot provide, in particular in the vertical. Specific scientific campaigns can also be carried out more quickly in regions with poor coverage of stations, either through short-term campaigns or by using air sampling equipment on commercial aircraft (programs such as the Comprehensive Observation Network for TRace gasses by AirLiner: CONTRAIL Machida et al., 2008). These vertical profiles provide valuable information for validating both forward and inverse models and better understanding their behavior in the troposphere. These aircraft measurements still have altitude limitations, with commercial airliners only reaching around 11 km and scientific aircraft usually reaching a maximum of around 14 km.

Stratospheric balloons are the most used method to obtain CO₂ concentration measurements at higher altitudes than this. They can carry different instruments, one in particular used in the greenhouse gas monitoring community is AirCore. Initially developed at the NOAA (Karion et al., 2010), it is now used by many different laboratories and agencies all over the globe (Membrive et al., 2017). It consists of a coiled tube that evacuates air while the balloon ascends, then collects ambient air when descending. It is then sealed on recovery and analyzed. The profile of the trace gases concentrations along the length of the tube then corresponds to a vertical profile of the concentrations in the atmosphere where the balloon descended. This system allows air sampling up to 30 km, which is extremely useful when comparing data to columns of CO₂ obtained through remote sensing.

3.2 Remote sensing

Remote sensing is the practice of observing the properties of an object without making physical contact with it. In our case, for observing greenhouse gas concentrations in the atmosphere, the most common method consists of using spectrometers to measure radiation emitted by the atmosphere and analyzing the absorption bands of the signal to deduce the concentrations. For CO₂, two main domains in the infrared can be used: near-infrared (780 nm to 2500 nm) and thermal infrared (2500 nm to 15 500 nm). They do not have the same vertical sensitivity and have different altitudes where the retrieval is the most accurate (Boesch et al., 2011). The domain used by the instruments in the next section is the near-infrared, which is sensitive to all of the atmosphere's height and can therefore give results on the so-called total dry air column of CO₂ (XCO₂). This remote sensing does carry information on the vertical distribution of the trace gas concentrations, but it is not enough to recreate an accurate vertical profile; instead, the information is integrated on the whole height to give a result on the whole column of dry air. Estimates of XCO₂ are derived by taking the retrieved vertical pro-




Figure II.3: Map of the location of current and future TCCON sites (<https://tcccon-wiki.caltech.edu/Main/TCCON>).

file of CO_2 concentrations x and multiplying it by a function dependent on the vertical profile of the atmosphere pressure h .

$$X\text{CO}_2 = h \cdot x \quad (\text{II.2})$$

These remote sensing instruments can be used on the ground at fixed stations or on moving ships (Warneke et al., 2005), but can also be used from space on satellites.

3.2.1 TCCON

The Total Carbon Column Observing Network (TCCON) is a network of ground stations that uses Fourier transform infrared spectrometric measurements (Petersen et al., 2008). The GGG2020 algorithm then analyzes these near-infrared measurements to obtain dry air columns of trace gas such as CH_4 , N_2O , or CO_2 . The network consists of 28 operational stations at the time of writing, with at least four future sites planned. This network is used extensively as a reference to calibrate and validate satellite remote sensing instruments (Crisp et al., 2017, Hong et al., 2022) and can also be used for inverse modeling for the same purpose (Byrne et al., 2023). Smaller versions of the instrument can be transported and moved to designated spots to be observed directly by satellites or even brought on-board ships to measure in the middle of the ocean. The growing number of groups using the smaller EM27/SUN instrument has led to the creation of the Collaborative Carbon Column Observing Network (COCCON, Frey et al., 2019) to complement the current TCCON network.

3.2.2 Satellites

Satellites equipped with appropriate spectrometers are the most important providers of remote sensing data on greenhouse gases. Several missions were launched with such instruments, with SCIAMACHY being the first capable of measuring the total air column in 2002. However, the first mission dedicated to measuring the total columns of greenhouse gases ($X\text{CO}_2$ and $X\text{CH}_4$) was the Japanese

satellite GOSAT, launched in 2009 and still in operation at the time of writing. Its measurements allowed at the time unprecedented precision of XCO₂ values and have been used in inverse modeling to get better estimates of CO₂ fluxes (Basu et al., 2013, Houweling et al., 2015). In 2014, the American satellite OCO-2 was launched with one of its goals being retrieving XCO₂ with a precision of 1 ppm in order to improve even further CO₂ fluxes estimates (Crisp et al., 2017). These two missions are only some examples of satellites currently monitoring the Earth's atmosphere, and their characteristics will be described in more detail in [chapter III, section 1](#). As soon as SCIAMACHY was able to provide the first estimates of XCO₂ from space, the challenges in obtaining accurate, spatially, and temporally desirable data became evident (Houweling et al., 2005) and started being improved for the missions that came after.

The retrieval of the concentrations of CO₂ in a column is highly dependent on the presence of clouds, aerosols, and variation in the orography (O'Brien and Rayner, 2002). One way of correcting some of these issues is to do a comparative absorption measurement using atmospheric oxygen (O₂). Molecular oxygen's concentration remains constant, well-known, and uniformly distributed throughout the atmosphere, making it an ideal measurement reference. These observations, for example, allow the OCO-2 mission to infer total atmospheric pressure and the concentrations of CO₂ in the air column. The different absorption bands of CO₂ do not all have the same sensitivity to the presence of aerosols, and comparing them can also give information on their presence and concentration. All these methods and calibrations against TCCON data have greatly improved the accuracy of XCO₂ estimates. Nevertheless, when used in inverse modeling, systematic errors as low as a few tenths of a ppm can lead to significant variations in the regional estimates of CO₂ fluxes (Chevallier, Bréon, and Rayner, 2007, Chevallier et al., 2005).

Another area for improvement with satellite data is their coverage, dependent on their orbits and instruments. For example, GOSAT and OCO-2 are on sun-synchronous orbits with a 3 and 16-day repeat cycle, which means that they recapture data at a given point on the surface of the Earth at this frequency. Many applications in the study of greenhouse gases, especially those of anthropogenic origin, could benefit from more frequent and dense observations, but numerical models must also be improved to better use the already existing data.

Improving atmospheric models

Contents

1	Why we need better horizontal resolution	22
2	Scaling up models and the need for speed	23
3	Choosing the right grid	25

INVERSE MODELS serve as a tool for comprehending the intricacies of the carbon cycle and the extent of anthropogenic influence on surface carbon fluxes. However, it must be noted that the estimates of these fluxes continue to exhibit substantial uncertainties, especially at regional levels and lower time scales. Enhancing the horizontal resolution of the atmospheric models used in inversions is imperative to refine these estimates, particularly when conducting inversions on a global scale.

Nonetheless, augmenting the resolution of these models entails heightened computational demands, which can present challenges regarding available resources, financial costs, and real-time limitations. The advancement of these models necessitates significant development efforts to address these complexities. In the subsequent sections, various strategies for managing these challenges will be presented, offering approaches to scale up these models while mitigating the associated computational costs.

In [section 1](#) the advantages and necessities of improving the horizontal resolution of these models is explained. Then, the necessary changes in computing resources to make these new resolutions possible are presented in [section 2](#). Finally, in [section 3](#), we explain the impact the grid choice has on models and how it can be leveraged when trying to reach higher horizontal resolution.

1 Why we need better horizontal resolution

Projects such as the Global Carbon Budget 2022 combine results from many different inverse models to get better estimates of surface carbon fluxes and can give a good overview of the state of the art of the resolution of these models when used for long-term inversions at the global scale. This data is reproduced in [Table III.1](#) (page 24). We can see that for the eight models that use a regular latitude-longitude grid at the global scale, the average horizontal resolution is 3.54° longitude and 3.53° latitude. At the equator, this is equivalent to a resolution of 393 km by 392 km. This is the size of a small country, and without efficient downscaling or other inversions done at regional scales (like Carbon-Tracer Europe) a lot of information is lost. Satellite missions can observe greenhouse gases concentrations at a much higher resolution. The characteristics and measurement resolution of the Greenhouse Gases Observing Satellite (GOSAT) and Orbiting Carbon Observatory-2 (OCO-2) satellite instruments are shown in [Table III.2](#) on page 25.

Missions like OCO-2 are equipped with instruments with a much greater resolution than previous ones, recording soundings over a 10 km-wide or less swath of the surface and with each pixel in that swath measuring 1.3 km by 2.3 km. Assimilating and aggregating that data to a grid with individual cells 51,000 times bigger leads to a loss of information. This data is often averaged in time which indirectly results in a loss of resolution, for example the 10-second average used in the OCO-2 model intercomparison project corresponds to about 67 km along the orbit track (Schuh et al., 2019). This difference in resolution between observations and models is called representativeness error. It can be substantial when a whole cell at a very coarse resolution is assumed to be representative of a point observation, be it from a station or a satellite observation. The representativeness error of CO₂ in transport models has been studied at different scales (Tolk et al., 2008), and while some studies at global scales of CO₂ show that simply increasing the model's resolution is not enough to decrease this error (Lin et al., 2018, Remaud et al., 2018) this is not the case at lower scales. When studied over certain regions with kilometer-scale simulations, the importance and impact of high resolution has been demonstrated in particular for regions with complex terrain (Hedelius et al., 2017). This is the case because representativeness error is not only due to the resolution mismatch of the assimilated observations. The resolution of the prior fluxes, orography maps, coastal boundaries as well as the meteorological fields used to drive the simulations also play an essential role. The weather forcing used in inverse models (shown in [Table III.1](#) on page 24) is usually at a higher resolution, 0.25° for the ERA5 (ECMWF Re-Analysis 5) reanalysis for example, and therefore could be better exploited by

high resolution global scale inversions. This could be very beneficial since uncertainty on the winds is one of the biggest source of error of CO₂ variability near emission hotspots (Agustí-Panareda et al., 2019). Prior fluxes can be more challenging to accurately obtain at higher resolution since they rely on many different acquisition methods, chemical and biological models. The orography is generally easier to get at very high resolution, and satellite data acquisition for example already requires such maps with a high degree of accuracy. However, these maps must also be accurate in time since man-made change can rapidly affect terrain height for example with large-scale mining operations, and these can affect local boundary layer height or XCO₂ values (Shi et al., 2000). All of these factors show that when increasing the resolution of the transport models, the resolution of the input data must also be taken into account. Furthermore, increasing the resolution of inverse models cannot be assumed to create linear improvement on the modeling of CO₂ (Agustí-Panareda et al., 2019). Sensitivity studies that find a small impact of horizontal resolution for coarse resolutions therefore cannot be extrapolated to very high resolutions. Different metrics will also scale differently to increased resolution. For example, flux distribution in regions with strong emission hotspots or complex orography will benefit more from resolution increase compared to XCO₂ values since they have a larger footprint (Agustí-Panareda et al., 2019).

Evidence points to the representation of clouds being the most significant cause of uncertainty in climate projections (Schneider et al., 2017), and at the current scale of horizontal resolutions for global simulations clouds can only be parametrized, making it almost impossible to properly reduce this uncertainty. Explicitly modeling deep convective clouds would significantly improve the representation of precipitation but would require a kilometer-scale horizontal resolution (Fuhrer et al., 2018). This provides an extremely ambitious goal for atmospheric modeling that comes with particularly challenging computing requirements.

2 Scaling up models and the need for speed

Weather and climate numerical modeling have been at the forefront of supercomputer applications since the 1960s (Manabe and Bryan, 1969), following advances in technology to provide ever more accurate, fast and lengthy simulations. Nowadays, most simulations for climate modeling and inverse modeling of CO₂ at global scales are run on such supercomputers, with tens of thousands or even hundreds of thousands of central processing unit (CPU) cores. Many models are legacy codes based on millions of lines of Fortran code that are highly parallelized for use with many CPUs using the Message Passing Interface (MPI) and Open Multiprocessing (OpenMP) libraries to divide the model space into different domains and timescale for optimization. This approach was mainly chosen to make best use of the constant upgrade of computing systems and hardware. Moore's law and Dennard scaling have observed that transistor count doubled every two years, ensuring faster processors while the growth of their energy consumption remained limited, allowing ever more CPUs to be used in parallel. However, the use of these resources by models raises questions of efficiency. Indeed, on present-day machines, climate models only achieve around 5% sustained floating-point performance and are barely exploiting the exponential growth of computing power that happened (Bauer et al., 2021). The traditional approach to accelerate climate models consists generally of refactoring portions of the code to better handle data communication, memory usage or making better use of parallelism. However, the basic algorithms as such are rarely changed. However, relying only on measuring floating-point performance can lead to misleading diagnostics of computing speed and resource allowance. Since most high-level climate models are actually coupling various models (atmosphere, land, ocean...), balancing the load between them so that there is a minimal amount of idle time is crucial for optimization. A more useful metric to see how a model fares on different computers is to look at the Simulated Year Per Day (SYPD) and Core Hours per Simulated Year (CHSY) to see how the model actually scales as well as giving an estimate of how much energy it consumes. The actual SYPD (ASYPD) is often lower than the theoretical one because of queuing time, work-flow, or system interruptions (Balaji et al., 2017). Scaling of these models is most often non-linear, and a big

	SCIAMACHY	GOSAT	OCO-2
Time of operation	2002 - 2012	2009 - ?	2014 - ?
Orbit	800 km sun-synchronous	666 km sun-synchronous	705 km sun-synchronous
Repeat cycle	6 days	3 days	16 days
Horizontal resolution per pixel	30 km × 60 km	10.5 km diameter	1.3 km × 2.3 km
CO ₂ spectral band	970 nm to 1772 nm	1560 nm to 1720 nm 1920 nm to 2080 nm	1591 nm to 1622 nm 2043 nm to 2083 nm

Table III.2: Characteristics of the SCIAMACHY, GOSAT and OCO-2 satellite missions (Dogniaux, 2021)

reason for this is because data movement is around a hundred times more costly than floating point operations (Shalf, Dosanjh, and Morrison, 2011). With increased resolutions also come challenges in simply storing the resulting data, performing simulations at kilometer-scale would increase the total data volume by around three orders of magnitude (Schär et al., 2020). Approaches based on minimal data storage, almost direct analysis, and keeping restart files for future reruns as needed may become more prevalent.

The reliance on mass parallelization of CPUs is also accompanied by a growing use of Graphics Processing Units (GPUs). GPUs are receiving widespread attention and investment from many different research fields (machine learning, big data), and the ratio of GPUs to CPUs in the biggest supercomputers tends to increase to adapt to their new uses. Atmospheric modeling should strive to stay at the forefront of these advances, and many research teams around the world have invested resources to meet this goal. The Nonhydrostatic ICosahedral Atmospheric Model (NICAM) was the first atmospheric model to be run at a large scale on GPUs (Yashiro et al., 2016) but in recent years many other climate or inverse models have started porting at least parts of their code to run on GPUs (Fuhrer et al., 2018, Giorgetta et al., 2022, Chevallier et al., 2023).

Parallelization of code to run on GPUs relies on different paradigms than the traditional approach of MPI and OpenMP, and it can seem like a daunting task to rework so many lines of code in already existing models. But most of the work can be facilitated by porting only certain routines (in the case of NICAM, starting only with the dynamical core) to run on GPUs, and by using tools such as OpenACC to minimize the amount of coding required and instead rely on compiler directives to handle parallelism of loops. This results in a parallel execution of the model on CPUs and GPUs, with GPUs performing the bulk of the calculations.

3 Choosing the right grid

Almost all the inverse models used in the Global Carbon Budget and presented in Table III.1, page 24, are running on a regular latitude-longitude spherical grid, meaning that each cell of the 3D grid is of a rectangular shape with evenly spaced lines of constant latitude and longitude. Figure III.1 illustrates such a latitude-longitude grid. Their rectangular shape, orthogonality, and symmetry make them well-suited to fulfill many of the requirements of weather and climate modeling. However, when approaching the pole, the convergence of the meridians makes the grid size in terms of distance approach zero and eventually reach a singularity. In practice, most of the models presented in Table III.1 (page 24) modify their grid near the poles to mitigate or bypass the problems created by this singularity. When using an explicit time integration scheme with finite difference, models have a restriction on the time step that is related to wind speed and grid size. This restriction is called the Courant-Friedrichs-Lewy (CFL) condition. This condition means that with an explicit time integration scheme

LATITUDE-LONGITUDE GRID




Figure III.1: A latitude-longitude grid consisting of equally spaced lines of constant latitude and longitude showing convergence at the poles (Williamson, 2007).

when the distance between grid cells decreases near the pole, the time step goes toward zero since wind speed is a physical quantity that should not change. Simply reducing the time step near the poles creates an enormous computing cost, with most of the computing time going toward simulating only a few longitude bands near the poles (Williamson, 2007). The development of semi-implicit time integration schemes allowed models to keep a single time step across latitudes, but this came at the cost of a lot of data communication across grid points near the poles (Staniforth and Thuburn, 2012). This greatly enhanced the viability of latitude-longitude grids for many years of model development. However, we saw above that models now tend to run on hundreds of processors in parallel which also requires spending a lot of time exchanging data between them. This most often becomes the real bottleneck in terms of computing performance and scaling to higher resolution. This problem is expected to become even more crucial to solve with future Exascale supercomputers since floating point operations are relatively fast and cheap compared to data movement (Shalf, Dosanjh, and Morrison, 2011, Schulthess et al., 2019).

One way to deal with the singularity at the pole is to simply bypass it. This can be done while staying with a general latitude-longitude grid by combining several projections to cover the sphere; this approach is called a composite mesh or overset grid. The other way to bypass it is to use a non-quadrilateral grid instead of a latitude-longitude grid, this is the case for example of NISMON-CO₂ in Table III.1 on page 24. This approach solves the singularity problem without having to deal with larger data movement near the poles, which makes these kinds of grids especially well suited for scaling to kilometer-scale simulations (Schulthess et al., 2019). The grid choice is also particularly important for spectral models: since the fields are represented in both the grid-point space and the spectral space, the transformation from one to the other depends on the grid properties. The grid not only determines the distance between each grid point but also how many grid points are used to represent the smallest wavelength of the spectral space. The limitations of linear grids have already been addressed in some spectral models, for example by replacing the grid with a cubic-octahedral grid (Malardel et al., 2016).

These unstructured grids are not new, and shallow-water models based on icosahedral grids for example have been developed since as far back as 1997 (Thuburn, 1997), but they have recently become

the new focus of many teams around the world. A good way to review the advances of models toward this goal is to look at the design of their dynamical cores, that is the component of the model that solves the adiabatic and frictionless equations of fluid motion in atmospheric dynamics, the Navier-Stokes equations. Out of the eleven reviewed dynamical cores in the 2016 Dynamical Core Model Intercomparison Project (DCMIP2016, Ullrich et al., 2017) six were running on non-quadrilateral grids. One was run on a centroidal Voronoi mesh: MPAS (Skamarock et al., 2012). One was run on an icosahedral triangular grid: ICON (Zängl et al., 2015). The last 4 were run on a geodesic or icosahedral hexagonal grid: OLAM (Walko and Avissar, 2008), CSU (Heikes, Randall, and Konor, 2013), DYNAMICO (Dubos et al., 2015), and NICAM (Niwa et al., 2017). Not all of these dynamical cores have been fully implemented into climate models or inverse models.

As noted above, dynamical cores need to satisfy some essential conditions for proper modeling of the climate. These prerequisites are easier to satisfy when employing regular latitudinal-longitudinal grids due to their rectangular structure, orthogonality, and symmetry. However, fulfilling these criteria can prove more challenging when dealing with alternative grid types (Staniforth and Thuburn, 2012). Some highly desirable conditions include preserving mass conservation for tracers and dry air, especially for long-term climate simulations. Balanced propagation of Rossby waves is essential, and model accuracy should approach a second-order level. Additionally, it is imperative to minimize grid imprinting, ensuring that points within the grid, such as the twelve pentagons in the geodesic grid, do not introduce spurious signals into the numerical solution. Finally, the discretization scheme of the continuous equations should possess mimetic properties to maintain the integrity of the modeling process.

In this manuscript, we will use the DYNAMICO dynamical core in one configuration of a global circulation model, and the icosahedral hexagonal grid will be described in details in [section 4](#) of [chapter V](#).

Outputs of climate simulations or inversions on these grids can also be difficult to exploit since most tools used routinely for diagnostics only work on regular latitude-longitude grids. Post-treatment of the data therefore has to be done either manually for each grid or using one of the few software capable of handling unstructured grids such as *Psyplot* (Sommer, 2017), or *UXarray* (Chmielowiec et al., 2023). Since many different types of grids co-exist, attempts are being made to propose new conventions to unify their use. The Unstructured Grid (UGRID) convention has been created to store unstructured (or flexible mesh) model data in the Unidata Network Common Data Form (NetCDF) file. This convention is becoming the new standard to allow easier use of these grids across different teams and models, but a lot of effort is still required from the community to make these open-source projects widely used.

Modifying already existing atmospheric models to use one of these grid is also a significant endeavor, as would be the other option of starting from scratch. For the case of our atmospheric CO₂ inverse system, the rapid evolution of the general circulation model from which it is derived forced our hand but also gave us an opportunity to improve its spatial resolution.

A strategy for the evolution of our atmospheric inverse system

Contents

1	The state of our current inversion system	30
2	Initial strategy	31
2.1	Initial assessment of the master GCM	31
2.2	Initial step-by-step plan	32
3	Final content of the thesis	32
3.1	Final step-by-step plan	32

THE INVERSE SYSTEM for atmospheric CO₂ at the Laboratory for Sciences of Climate and Environment (LSCE) had a newfound opportunity to quickly evolve and improve. In this chapter, we will present the context in which our CO₂ atmospheric inverse system was created and what strategies we initially envisioned to improve its horizontal spatial resolution. This plan that we set out at the beginning of our work continuously evolved as setbacks and unexpected successes occurred during development and the scientific study of our results. In this chapter we will first briefly present our atmospheric inverse system and how it fits with the other models at the LSCE in [section 1](#). Then in [section 2](#) we will explain the original plan and steps we wanted to follow to increase the resolution of our inverse system. Finally in [section 3](#) we present the steps that we ended up taking and how each stage of our scientific study shaped the modifications of our original plan. We also explain how each part of this manuscript fits into the approach we have described.

1 The state of our current inversion system

The inverse system which we want to improve the spatial resolution of is used to generate global CO₂ atmospheric inversion products for the Copernicus Atmosphere Monitoring Service (CAMS). It was created by LSCE at the start of the Global Earth-system Monitoring using Space and in-situ data research project (GEMS) that was the first precursor of CAMS (Hollingsworth et al., 2008). It has been continuously operated within CAMS by LSCE for CO₂ and the Norwegian Institute for Air Research (NILU) for N₂O since its start in 2015 through the successive service contracts CAMS 73, CAMS 73 P2, CAMS2 55 and CAMS2 55bis. We will call it the CAMS/LSCE inverse system in the following. It has a particular relation to other models developed at the LSCE and IPSL (Institut Pierre-Simon Laplace). In [Figure IV.1](#) (page 31) we present a schematic view of how the inverse system fits with the evolution across time of the other models.

The CAMS/LSCE inverse system is based on an atmospheric transport model directly developed from the LMDZ (Laboratoire de Météorologie Dynamique Zoom) atmospheric general circulation model: certain routines from the code of LMDZ were directly extracted to create an offline transport model. This offline model uses pre-generated mass fluxes from a simulation of a general circulation model (**master GCM**) to transport atmospheric CO₂ for a very minimal computational cost. This offline transport model is associated to its tangent linear and adjoint model to perform inversions of atmospheric CO₂.

The master GCM is a coupled configuration that is composed of an atmospheric model (LMDZ), a land model (ORCHIDEE) and an atmospheric chemistry model (INCA). We describe these models and configurations in detail in [chapter V](#). The development of this GCM is mainly carried out by the IPSL and its primary concern has to do with climate simulations, not with atmospheric transport of CO₂. This highlights a problem with our inverse system: it is dependent on models developed by other teams that have different goals.

The offline transport model was originally created in the 2000s from an old version of LMDZ and was later improved multiple times following the developments of LMDZ. Nonetheless it always had to play catchup with LMDZ, at the risk of becoming obsolete if not changed in time with each major update of LMDZ. At the start of this work, the CAMS/LSCE inverse system was using inputs of mass fluxes from the master GCM at a resolution that was slowly being phased out for other applications. A first update of the inverse system to run at the new resolution was already looking necessary, but the associated computational costs were expected to be high.

The biggest update however was the introduction of DYNAMICO to the master GCM: a new dynamical core running on an icosahedral grid created to replace the dynamical core of LMDZ. Given that the long-term plan for the master GCM was for it to entirely switch to DYNAMICO, we had the opportunity to be some of the first users of this new configuration and to prepare the inverse system to this change ahead of time. Also notable is the introduction of XIOS (Xml Input/Output Server) in

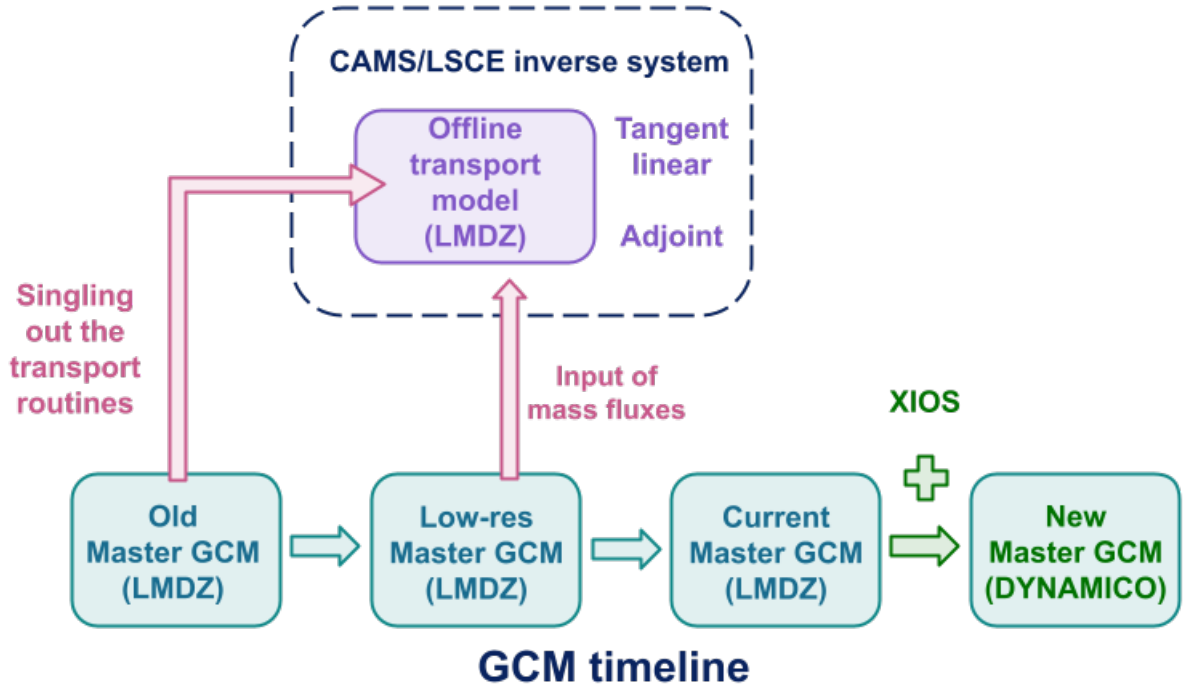


Figure IV.1: Overview of the architecture of the CAMS/LSCE inverse system and its position relative to the evolution of the master GCM.

the master GCM: a new input/output server designed to scale well to high-resolutions. This tool could perhaps also be used with the inverse system.

2 Initial strategy

As presented previously in [chapter III](#), one of the main advantage of an icosahedral grid, and therefore of DYNAMICO, is to remove the singularity at the poles allowing for a better parallelization of the model. The new grid would therefore provide computational gains but we also hoped that the removal of the polar filter might improve the representation of the atmospheric transport near the poles. The advantages of this new grid and dynamical core would allow for an increase in the horizontal resolution of the CAMS/LSCE inverse system and to leverage the rapid evolution of high-performance computing.

2.1 Initial assessment of the master GCM

The initial assessment of the evolution of the master GCM and how we could adapt the inverse system to these changes lead to starting this thesis with a first plan of action. The changes brought by DYNAMICO in the representation of atmospheric transport of CO₂ had to first be evaluated in the new configuration of the master GCM and compared to a reference configuration using the dynamics of LMDZ on a regular latitude-longitude grid. This comparison would also allow us to generate mass fluxes to drive CO₂ inversions at a higher resolution on the regular grid. We could then focus on developing a new version of the CAMS/LSCE inverse system using an offline version of DYNAMICO as its transport model. This new version would be the key to further increase the spatial resolution of the inverse system and ultimately aim to reach a sub 1° target resolution.

2.2 Initial step-by-step plan

To sum up, our initial step by step plan was:

1. Evaluate the representation of the atmospheric transport of CO₂ with the master GCM using a coupled configuration with DYNAMICO.
2. Test a first increase in horizontal resolution of the inverse system with the regular grid of LMDZ.
3. Modify the CAMS/LSCE inverse system to use DYNAMICO.
4. Use the DYNAMICO version of the inverse system to perform and test the first high-resolution global inversions of CO₂.

Throughout this work, our initial strategy to improve the inverse system required frequent adjustments, as each step often yielded unexpected results. This flexibility enabled us to quickly pivot to new techniques, circumvent setbacks, and take advantage of unforeseen solutions, all while maintaining our ultimate goal of increasing spatial resolution and reach a sub 1° target.

3 Final content of the thesis

The central aim of this thesis was to explore and evaluate different ways of preparing our CO₂ atmospheric transport model and inverse system for future computing infrastructures and to increase its horizontal resolution.

The initial plan outlined above only gave us a rough guideline of which leads to follow but was based on assumptions that often proved incorrect. We will now briefly present the steps that we took during this work, the motivations behind our choices and where each of these steps is detailed in the manuscript.

3.1 Final step-by-step plan

We first present the models we are working with in [Part II](#), starting with the master global circulation model we use for CO₂ tracer transport in [chapter V](#). This model then feeds input into our inverse system which we present in details in [chapter VI](#). The general concept of these two methods are also described.

After this, [Part III](#) presents the major developments carried out during this thesis, the scientific evaluation of their implementation, and the takeaways from these studies.

The first step of our initial plan:

1. Evaluate the representation of the atmospheric transport of CO₂ with the master GCM using a coupled configuration with DYNAMICO.

was already complicated by the initial state of the coupled configuration of the master GCM using DYNAMICO. It was not entirely ready to be used in a similar way as the regular configuration and the exploitation of the new icosahedral grid also proved difficult. In [chapter VII](#) we first present the initial state of the models and tools we had at our disposal, as well as some preliminary work that was necessary for the next chapters.

Then in [chapter VIII](#) we compare the new configuration of the master GCM running with DYNAMICO to our reference configuration running on a regular latitude-longitude grid. We are looking to answer two main questions:

Does using an icosahedral grid with DYNAMICO improve the performance of our atmospheric transport model of CO₂, in particular near the poles?

Does this configuration have a better computational performance than the reference?

The new configuration of the master GCM using DYNAMICO only showed modest improvements in terms of the modeling of atmospheric tracer transport but did show promise in terms of computational performance at higher resolutions. The generation of the mass fluxes by the master GCM requires a computationally intensive simulation and at high-resolution, the configuration using DYNAMICO would parallelize much better and allow us to generate them faster. However, it quickly became clear that creating an offline transport model based on DYNAMICO would have a high development cost and creating an adjoint model of it to use in our inverse system even more so. The decision was made to delegate this work to a new hire, Sakina Takache, who would focus on this particular goal and work in parallel to this thesis. Since the computational advantage of DYNAMICO in the master GCM was mainly useful at higher resolutions for which our inverse system was in any case not yet prepared for, this work could be delayed until the inverse system was ready to handle inputs of high-resolution mass fluxes.

We could however still continue with the second step of the initial plan, and focus first on increasing the horizontal resolution of the inverse system:

2. Test a first increase in horizontal resolution of the inverse system with the regular grid of LMDZ.

However, this required a lot of preliminary work to achieve, namely compressing the input mass fluxes and parallelizing the code on GPUs. This work is presented in [chapter IX](#).

GPU parallelization of the regular latitude-longitude offline transport model was not initially envisioned as a solution to increase the resolution of our inverse system but proved to be extremely efficient for a relatively small engineering cost compared to using DYNAMICO. At the same time, the development of an offline version of DYNAMICO was facing a lot of difficulties and proved more complicated than first expected. Given these results, we chose to modify the next step to become:

3. Directly reach the goal of high resolution inversions with the regular latitude-longitude grid of LMDZ.

In [chapter X](#), we carry out and study the impact of this increase in horizontal resolution on global inversions of atmospheric CO₂. The last study of this thesis was successful in its goal of resolution increase and showed some improvements brought by the high-resolution inversion.

The last step:

4. Develop the DYNAMICO version of the CAMS/LSCE inverse system.

is still ongoing in the SATINV (INVerse modeling for atmospheric and SATellite measurements) team at LSCE. We conclude this manuscript with [Part IV](#), first by summing up the work carried out so far and then by discussing the final configurations of the CAMS/LSCE inverse system that resulted from this work, including those still under development. DYNAMICO was eventually successfully integrated in the inverse system and the specifics of these first results will also be discussed.

Part II

Methods

Atmospheric general circulation models (GCM)

Nous sommes dans une situation où les méthodes des sciences expérimentales ne nous servent pratiquement à rien.

Parce que, finalement, une planète Terre, il y en a une seule et une situation comme une situation de crise où nous sommes maintenant n'a lieu qu'une seule fois dans l'histoire de l'évolution.

Allons-nous continuer la recherche scientifique ?

– *Alexandre Grothendieck*

Contents

1	What makes a general circulation model?	38
1.1	Atmospheric modeling	38
1.2	Eulerian approach	39
2	Coupled configurations of the master GCM	40
3	Direct atmospheric model: LMDZ	41
3.1	Grid of LMDZ	41
3.1.1	Horizontal grid	41
3.1.2	Vertical grid	42
3.2	Dynamics of LMDZ	43
3.3	Physics of LMDZ	44
3.4	Nudging of atmospheric variables	44
3.5	Parallelization of LMDZ	45
4	Icosahedral dynamical core: DYNAMICO	46
4.1	Description of DYNAMICO	46
4.2	Icosahedral grid	47
4.2.1	Horizontal resolution	47
4.3	Parallelization of DYNAMICO	48
5	Chemistry model: INCA	49
6	Land model: ORCHIDEE	49
7	Input/Output module: XIOS	50

TO UNDERSTAND THE BEHAVIOR OF CO₂ in the atmosphere, we are required to use models to represent complex processes.

In this chapter, we will first present the general concepts governing atmospheric modeling, with a particular emphasis on Eulerian models. In [section 2](#), we will present the specific coupled configurations used later in [Part III](#). Then, in [section 3](#), we will describe the LMDZ global circulation model used throughout this work. After that, in [section 4](#), we will describe DYNAMICO, a modern dynamical core designed to be the cornerstone of modern versions of LMDZ. Finally, the last three sections present the land model, atmospheric chemistry model and input/output server used in the GCM.

1 What makes a general circulation model?

Mathematical models are a way to represent the physical world we exist in, using mathematical equations, and using these to make predictions about it. While the use of mathematical models has been evidenced for much of recorded written history, the way we have thought about the relation between models and reality has varied across time and cultures. According to the Cambridge History of Science (Schank and Twardy, 2009) we can broadly separate the answers to the model-reality dichotomy into three kinds:

- **Aristotelianism** asserts that the quantities studied by mathematics literally occur in the objects subject of these models. But since few if any real world object perfectly match the *true* mathematical object, they can only be studied as being close enough to them.
- **Pythagoreanism** take an opposite approach and describe nature as being governed itself by mathematical principles. Physical processes follow hidden mathematical structures which mathematical models try to discover and approach.
- **Instrumentalism** on the other hand sees the mathematical objects modeled only as *symbolic* representations of facts about real world objects. Mathematical models are not descriptions of the world but merely instruments of predictions.

Whichever of these approaches is subscribed to, the mathematical models in our possession only give us a glimpse in the complexity of the physical world and allow us to make imperfect predictions about it.

The most common way to judge if a mathematical model is “good” or not is to judge the quality of the past and future predictions it makes of the physical world. This can only be judged in light of our observations of the physical world, which are themselves imperfect. The closer the predictions of the model are to the observations of the reality, the “better” the model. For complex models, trying to predict complex systems, this task of improving the predictions of many processes to better match vast amounts of observations seems daunting.

This is what atmospheric modeling aims to achieve.

1.1 Atmospheric modeling

General circulation models (GCMs) are numerical models that represent the physical processes of the atmosphere, ocean, land and cryosphere. **Atmospheric general circulation models** refer to numerical models that only model the atmosphere while imposing and receiving boundary conditions from other domains such as the ocean. In the rest of this text, we will refer to atmospheric general circulation models interchangeably as **GCMs** even though no modeling of the ocean circulation will be taken into account or simply as **atmospheric models**.

Atmospheric models describe the evolution in time and space of physical variables of the atmosphere (temperature, winds, pressure...), according to a set of equations governing the motion of the

atmospheric flows. They are generally based on a set of three so-called primitive nonlinear differential equations:

- A **continuity equation** describing the conservation of mass.
- A form of the **Navier-Stokes equations** corresponding to the conservation of momentum.
- A **thermodynamic equation** describing the conservation of energy.

These primitive equations are supplemented by parameters used to simplify small-scale and complex processes that cannot be explicitly resolved by the model.

Atmospheric models can be used to predict states of the atmosphere at different space and time scales, to predict the weather tomorrow in a specific city or to predict the average global temperature 200 years from now. Today these mathematical models are resolved numerically: the equations are discretized and we find approximate solutions to these nonlinear differential equations.

In this work we are primarily focused on studying the evolution in time and space of CO₂ concentrations in the atmosphere. In our atmospheric model we represent CO₂ as a passive **tracer**. The intra-annual variations of its very low concentrations hardly affect the flow or density of the air mass. It is also non-reactive: atmospheric CO₂ is mainly affected by sources and sinks from the land and ocean, with marginal contributions from the oxidation of reduced carbon compounds in the atmosphere. We call the parts of the atmospheric model that govern the flow of this tracer in the air masses the **transport model**.

1.2 Eulerian approach

There are two ways to look at the motion of a fluid, with the fluid in this case being the atmosphere:

The **Lagrangian** specification of a flow field consists of following individual infinitesimally small volumes of fluid and tracking their position across time and space. In a numerical atmospheric model, it means that *particles* representing small air parcels are transported in the atmosphere. These particles are independent of the computational grid and can have in theory infinitesimally small resolution. The main advantage of this approach is that a Lagrangian model does not exhibit any numerical diffusion. An example of a Lagrangian atmospheric particle dispersion model that can also be used for atmospheric inversions of CO₂ is FLEXPART (FLEXible PARTicle dispersion model, Stohl et al., 2005, Pisso et al., 2019).

The second approach, chosen in the model used in this work, is the **Eulerian** specification of the flow field. It consists in fixing a frame of reference and looking at specific locations in which the fluid flows. For a tracer i of mass density ρ_i , the continuity, or mass conservation equation means that in every point of the field we have the following:

$$\frac{\partial \rho_i}{\partial t} + \rho_i \nabla v = 0 \quad (\text{V.1})$$

Where v is the wind speed. The flux of air entering and leaving a given volume is calculated by $\rho_i \nabla v$. For numerical atmospheric modeling it means that this continuous equation is discretized on a finite spatial grid and then this equation is solved in each grid cell. In addition, at the surface, sources and sinks for each of these tracers modify the tracer mass in each cell.

Back in [chapter II \(Part I\)](#), we briefly described CO₂ atmospheric inverse modeling: models that estimate the sources and sinks of CO₂ from prior information about these fluxes and from information gathered by observations of CO₂ concentrations. Forward or **direct modeling** of atmospheric CO₂ is the opposite of inverse modeling and the more straightforward way of studying the same phenomenon: given a certain initial atmospheric state and prescribed surface fluxes of CO₂, how does the atmospheric concentration of CO₂ evolve in time and space?

2 Coupled configurations of the master GCM

The direct modeling of atmospheric CO₂ in this work was carried out using a GCM that couples different models. We call it **Master GCM** since it drives the inputs of the inverse system. This was already briefly explained in [Part I, chapter IV](#) (see [Figure IV.1](#) on page 31) and will be described further in [chapter VI](#).

We study two coupled configurations in particular:

- **REG**, which stands for *regular*. Coupling the dynamics and physics of LMDZ for the atmospheric model with the atmospheric chemistry model INCA (Interaction between Chemistry and Aerosol) and the land model ORCHIDEE (Organising Carbon and Hydrology In Dynamic Ecosystems). It runs on a regular latitude-longitude grid.
- And **ICO**, standing for *icosahedral*. A new configuration that couples the dynamical core of DYNAMICO with the physics of LMDZ and also uses INCA and ORCHIDEE models. It runs on the icosahedral grid of DYNAMICO.

Both of these configurations use XIOS for input/output (except for the inputs of the LMDZ and ORCHIDEE models in the REG configuration) and a schematic view of these configurations is shown in [Figure V.1](#).

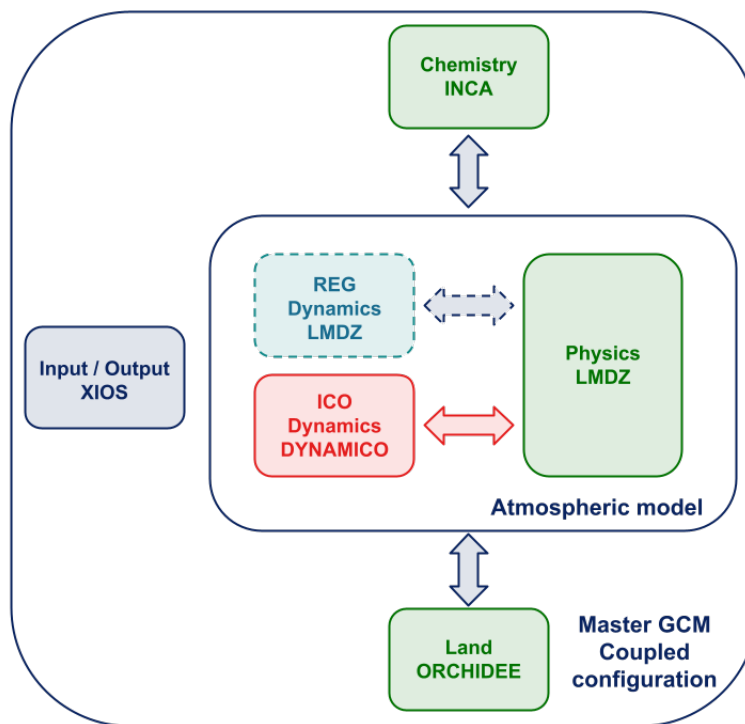


Figure V.1: Structure of the two coupled configurations of the master GCM. They differ by which dynamical core they use and therefore which grid they run on. Either the regular LMDZ dynamical core and latitude-longitude grid (REG configuration). Or the DYNAMICO core on an icosahedral grid (ICO configuration).

These two configurations differ by which dynamical core they use and by which underlying grid they are run on. We compare their representation of the atmospheric transport of CO₂ in [chapter VIII \(Part III\)](#). We now present LMDZ, the Eulerian atmospheric general circulation model used throughout this work.

3 Direct atmospheric model: LMDZ

To carry out direct simulations of CO₂ atmospheric transport in this work we use different configurations of the LMDZ atmospheric model. LMDZ takes its name from the Laboratoire de Météorologie Dynamique where it was created in Paris, France, the “Z” standing for “Zoom” capability, denoting the ability to stretch its grid. LMDZ is the second version of the LMDZ climate model first developed in the 1980s. It is still under active development at the Laboratoire de Météorologie Dynamique (LMD) and in use at the Institut Pierre-Simon Laplace (IPSL), and many other laboratories. Of particular note, it is the atmospheric component of the IPSL climate model and has been used in the international Coupled Model Intercomparison Project (CMIP). The LMDZ model is divided in two parts: one handling the dynamics (also called the dynamical core) described in [subsection 3.2](#). The second handling the parametrized sub-grid physical processes described in [subsection 3.3](#). The method used for nudging atmospheric variables to pre-determined fields is described in [subsection 3.4](#). The atmospheric transport of tracers is done in a few specific sections of the dynamical core. These sections can be singled out to create a so-called “offline model”: a stripped down version of the model using pre-computed dynamical and physical variables to simulate the atmospheric transport of tracers at minimal computing cost. This offline model forms the basis of the CAMS/LSCE inverse system described later in [chapter VI](#).

The LMDZ atmospheric model can be freely coupled with oceanic (Nucleus for European Modelling of the Ocean, NEMO), land (ORCHIDEE) or atmospheric chemistry (INCA) models, or use forced conditions instead. The dynamical core is flexible and can also be used to model non-terrestrial planetary atmospheres such as Mars (Pottier et al., 2017) or Venus (Navarro, Schubert, and Lebonnois, 2018). When used in that field of study, the convention now is to call this model the “Planetary Climate Model” or PCM, preceding it with the name of the planet studied.

We will now present the horizontal and vertical grid of the LMDZ atmospheric model.

3.1 Grid of LMDZ

The LMDZ atmospheric model is run on a regular latitude-longitude Arakawa C-grid (Kasahara, 1977) with a hybrid $\sigma - p$ vertical coordinate system.

3.1.1 Horizontal grid

The dynamical equations are discretized on the sphere on a regular latitude-longitude staggered Arakawa C-grid. For each cell, the variables such as pressure or temperature are evaluated at the center of each grid (i, j) but vectors such as wind velocities or mass fluxes are evaluated at the center of the grid faces. Moreover, east-west (u) and north-south (v) velocities for example are evaluated at different grid faces (see [Figure V.2](#)).

Given that the size of the rectangular cells of the grid diminishes when approaching the poles, to avoid having to reduce the time step to respect the CFL condition (see [chapter III](#) in [Part I](#)), a longitudinal filter is used for latitudes after 60° in both hemispheres: close to the poles, the meteorological fields are filtered so that only wavelengths superior to half the size of the cells at the equator in the longitudinal direction in km are retained (Hourdin et al., 2013a).




Figure V.2: Representation of a cell in a staggered Arakawa C-grid. Variables can be evaluated at the center $((i, j))$ or at the corners (q) . Vectors $(u$ or $v)$ are evaluated at the center of the grid faces.

3.1.2 Vertical grid

The vertical discretization is done according to a classical hybrid $\sigma - p$ coordinate system. For a given vertical layer l , the pressure p_l depends on the surface pressure p_s such that at the interface with the previous layer:

$$p_l = A_l + B_l \cdot p_s \quad (\text{V.2})$$

The values of the coefficients A_l and B_l determine how influential the topography of the surface is on the vertical layers. It is chosen so that for $l = 1$, at the surface, $A_1 = 0$ and $B_1 = 1$. Then as the layer number increases toward the maximum L : $\lim_{l \rightarrow L} B = 0$ and A_l dominates and they become equivalent to a pressure coordinate.

All the configurations of LMDZ in this work use a setup with **79 vertical layers**. It was previously run with only 39 vertical layers for use in CMIP5, extending into the stratosphere but only giving a very coarse representation of it Hourdin et al., 2013a. The 79 layer discretization developed for CMIP6 was primarily aimed at improving stratospheric circulation though it also increased the vertical resolution near the surface. The first layer is centered at 10 m above the surface and 25 layers are used to model the first 2 km with the thickness of each layer varying in a linear manner. The size of each vertical layer then increases faster until the final one at an altitude of around 80 km. A representation of the vertical grid is shown in Figure V.3. The future version of LMDZ for CMIP7 aims to use 95 vertical layers.

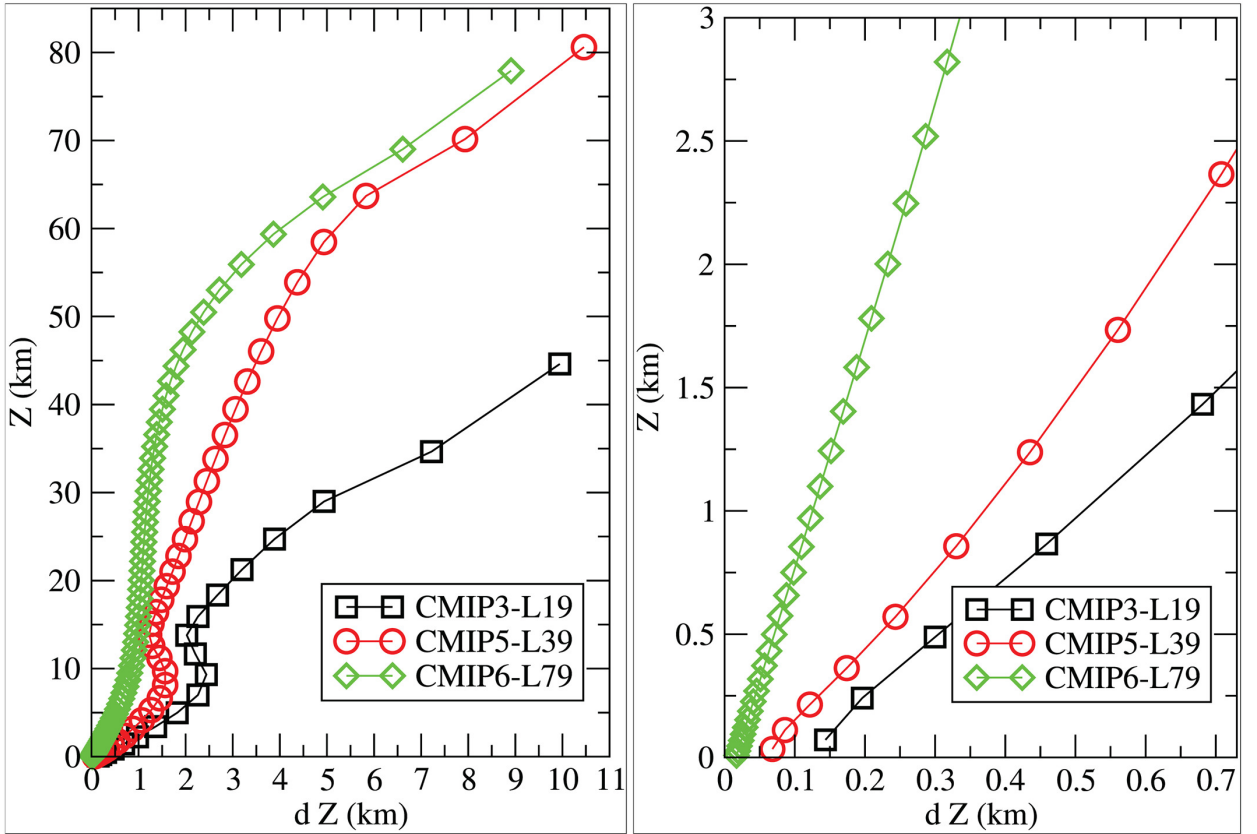


Figure V.3: LMDZ vertical grids used for the reference configurations of the successive phases of the CMIP exercise with the IPSL model showing the altitude in y axis as a function of the layer thickness in x. A focus on the first 3 km is shown in the right panel. (Taken from Hourdin et al., 2020)

3.2 Dynamics of LMDZ

The fundamental component of an atmospheric model is its dynamical core which solves the primitive equations of motion of the atmosphere. The finite difference discretization of the primitive equations are made to ensure numerical conservation of the air mass, the enstrophy (the square of the wind rotational, see Equation V.3) and the angular momentum. The choice was made to conserve enstrophy rather than explicitly conserving energy, though significant effort is made to limit the rate of energy leak in the model.

$$\varepsilon(u) = \int_{\Omega} |\nabla u|^2 dx \quad (\text{V.3})$$

The dynamical core is **hydrostatic**: it is assumed that the horizontal scales are larger than the scale height, this assumption is valid up to around 10 to 20 km. The advection of liquid water and vapour as well as tracers is carried out by a Van Leer monotonic second order finite volume scheme (Van Leer, 1977). It is however important to note that in LMDZ, water is considered weightless and therefore does not influence surface pressure. The impact of this on atmospheric tracer transport is partly counter-balanced by the nudging of the winds, which will be described below in subsection 3.4. A horizontal dissipation term is added to the equation of the dynamics to represent the transfer of enstrophy at the cut-off scale. In addition to this horizontal dissipation, a so-called “sponge layer” is used to damp the vertically-propagated gravity waves from the model top layer. In practice it consists in forcing the wind and/or temperature of the topmost vertical layers (usually 4) to relax towards zero or towards the mean zonal values. The relaxation time of this function can be parametrized.

3.3 Physics of LMDZ

While the dynamics of LMDZ have essentially stayed the same since its creation, the physics on the other hand have greatly evolved since its inception. Major reworks of the physics lead to a new version numbering of LMDZ. We will present here the main characteristics of the physics brought by either LMDZ4 (Hourdin et al., 2006), LMDZ5A (Hourdin et al., 2013a), LMDZ5B (Hourdin et al., 2013b) or LMDZ6A (Hourdin et al., 2020), the last version corresponding to the most up to date physics used in this work.

In LMDZ the physics is called at a less frequent time-stepping than the dynamics in order to save computational resources while still maintaining numerical stability. In the latest LMDZ6A physics, it is called every 15 minutes regardless of the horizontal resolution, to be compared to every 30 minutes for LMDZ5A. Even in the physics, some processes are run at a lower frequency, the radiation model being called every 1.5 hours and the deep convection every second physical time step.

The LMDZ5B version (Hourdin et al., 2013b) using the so-called “New Physics” also used in LMDZ6A introduced a new object-oriented approach to the parametrization of the convection. Instead of using a unified parametrization for the convection, they are instead configured independently and their coupling itself is parametrized after that. The vertical motion modeled by the sub-grid processes of convection and turbulence are divided into three components:

- Small scale turbulence, dominant in the surface layer
- Boundary layer convection
- Deep convection associated with cumulonimbus

In the boundary layer, the first two components are dominant with the turbulence modeled with the Yamada diffusive scheme (Yamada, 1983) and the dry and cloudy shallow convection modeled by a mass flux scheme. The mass flux approach represents mean ascending thermal plumes in the grid cell. The diffusive scheme was revisited and improved in LMDZ6A.

The deep convection uses a version of the mass flux scheme of Emanuel (Emanuel, 1991). It is coupled to the shallow convection and triggered if the lifting energy provided by the thermal plumes is sufficient.

The radiative code is inherited from the one of the European Centre for Medium-Range Weather Forecasts (ECMWF) weather forecast model, and in LMDZ6A the switch was made to a rapid radiative transfer model (RRTM) with 16 spectral bands (Mlawer et al., 1997).

Tuning of these numerous physical parameters is principally done as part of the CMIP work with the goal of reproducing observed climatology.

3.4 Nudging of atmospheric variables

In most simulations we ran with the GCM throughout this work, we used the *nudging* option (also called guided mode) for the wind fields. This option allows some meteorological variables (winds, temperature, pressure...) from the model to be nudged towards pre-determined fields, either from a reanalysis or a previous simulation for example. In our case we only used this option to nudge wind fields to the ERA5 reanalysis. Since wind fields have such a strong impact on the atmospheric transport of tracers, using reanalysis fields instead of only the modeling from LMDZ is necessary to properly compare the CO₂ concentrations to observations, and for generating mass fluxes that are useful to the inverse system.

At each time step of the dynamics the atmospheric variable, in this case for example the zonal winds u , is nudged towards the target zonal winds u_t at a rate depending on the relaxation coefficient α that is itself defined by the time step of the dynamics Δt and the relaxation time τ . These equations describing the new value of the nudged winds u_n can be written:

$$u_n = \alpha u + (1 - \alpha)u_t \quad (\text{V.4})$$

$$\alpha = 1 - \frac{\Delta t}{\Delta t + \tau} \quad (\text{V.5})$$

From this equation we can see that when the relaxation time τ gets smaller, the nudging gets stronger and the closer the final wind field will be to the reanalysis field. In practice in all of our GCM simulations using nudging on the winds, we chose a relaxation time of 3 hours.

As mentioned previously in [subsection 3.2](#), water is considered weightless in LMDZ and does not affect the calculation of the surface pressure which in turn impacts the atmospheric transport. However since the meteorological variables from ERA5 take water mass into account, the nudging of the winds to the ERA5 fields mostly counterbalances that approximation from LMDZ in practice.

3.5 Parallelization of LMDZ

LMDZ can be run in parallel on many processors following a hybrid MPI/OpenMP approach as briefly described back in [section 2](#) of [chapter III](#) in [Part I](#). The parallelization strategy is different in the dynamics and the physics:

Dynamics There are many interactions and data exchange between neighboring cells, requiring a robust spatial distribution of the processes illustrated in [Figure V.4](#) on page [Figure V.4](#). The first level of parallelization is done through an MPI tiling by bands of latitude, each MPI process is responsible for at least 2 bands of latitude. A limitation of this approach however is that there is a maximum number of MPI processes for a given horizontal resolution. Different latitude bands also have different computational loads, particularly for bands where the polar filter is activated. The second level of parallelization is on the vertical levels with OpenMP. Each column is divided into chunks of a few vertical levels, with each chunk assigned to a singular thread.

Physics Since the physics of LMDZ is modeled within individual columns of the atmosphere, there is no interaction between neighboring columns of the atmosphere, therefore the columns can be distributed over all the available cores. The total number of columns is first distributed among the MPI processes, then the columns assigned to each MPI domain are further subdivided into OpenMP domains that will be assigned to individual threads. The division and number of columns in each MPI and OpenMP domain can be configured as long as the total number of columns of atmosphere (k_{glo}) is equal to the sum of the number of columns per MPI domain (k_{MPI}), which is itself equal to the sum of the number of columns per OMP thread (k_{OMP}):

$$k_{\text{glo}} = \sum k_{\text{MPI}}$$

$$k_{\text{MPI}} = \sum k_{\text{OMP}}$$

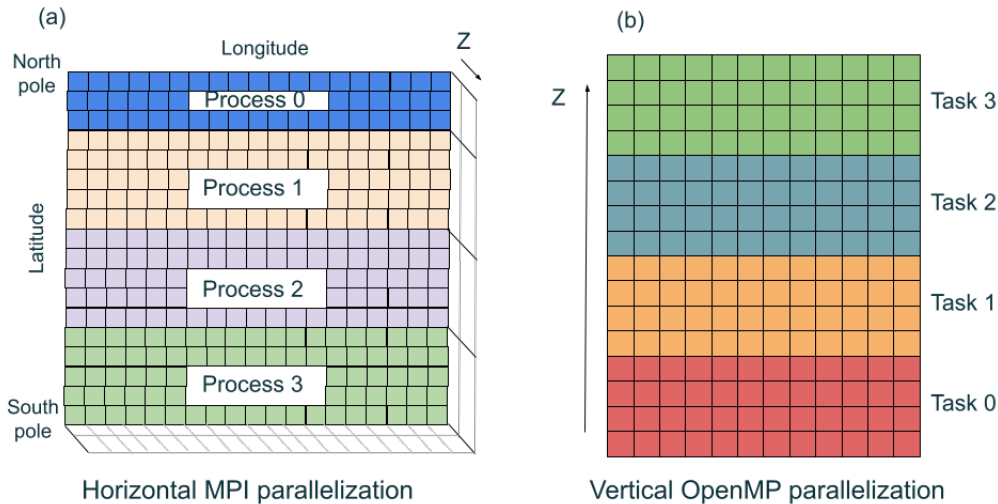


Figure V.4: Schematic view of the hybrid MPI/OpenMP parallelization of the dynamics of LMDZ. Panel (a) shows the horizontal MPI parallelization per latitude band and panel (b) shows the vertical parallelization.

4 Icosahedral dynamical core: DYNAMICO

In 2009, the IPSL started to develop a new dynamical core that could be used to replace the one of LMDZ. The main objective of this new dynamical core was for it to use a quasi-uniform grid to avoid the computational bottleneck created by the polar singularities of a regular latitude-longitude grid as discussed back in [section 3 of chapter III \(Part I\)](#). This made it particularly appealing for future high-resolution modeling running on massively parallel supercomputers. This work resulted in the first hydrostatic version of DYNAMICO in 2015, presented in [Dubos et al., 2015](#). DYNAMICO was then quickly extended to solve non-hydrostatic equations and described in the DCMIP2016 comparison exercise ([Ullrich et al., 2017](#)).

We will now describe the specifics of DYNAMICO.

4.1 Description of DYNAMICO

DYNAMICO aims to have exact discrete conservation properties similar to the original dynamical core of LMDZ, but ensuring this for all equation sets, even the non-hydrostatic ones, requires formulating a general approach valid for all of them. The answer to this problem chosen for DYNAMICO was to use a Hamiltonian formulation of the equations of motion. This approach was not new and already used in the ICON-IAP dynamical core ([Gassmann, 2013](#)), but was limited to fully compressible equations in Eulerian coordinates. This approach was extended in [Tort and Dubos, 2014](#) and [Dubos and Tort, 2014](#) to include the equations of compressible hydrostatic flows and non-Eulerian vertical coordinates.

The vertical coordinates of DYNAMICO are non-Eulerian (the geopotential depends on time) and mass-based rather than pressure-based. In our use case of shallow-atmosphere hydrostatic equations it is equivalent to the hybrid $\sigma - p$ coordinate of LMDZ, but this is not the case in general.

The kinetics responsible for the transport equations of mass, scalars and entropy do not use any information from the momentum equations. The vertical advection follows a Van Leer scheme, identical to that of LMDZ. A positive-definite finite-volume scheme is used for the horizontal advection.

The enstrophy-conserving finite difference scheme for the horizontal dynamics used in LMDZ was generalized to non-orthogonal dual meshes (Thuburn, Cotter, and Dubos, 2014).

We now present the horizontal grid of DYNAMICO.

4.2 Icosahedral grid

The icosahedral grid also called hexagonal grid or geodesic grid used by DYNAMICO is based on a tessellation of the sphere. A triangular mesh is generated first from a spherical icosahedron divided into 20 spherical triangles, then further divided an arbitrary number of times (nbp) which will define the final resolution of the grid. See Figure V.5. We call this the primal mesh.

Then, a mesh is constructed with a Voronoi diagram by joining the points of the triangles, such that they become the center of the hexagonal grid. This guarantees the orthogonality of the two meshes, necessary for the numerical scheme, and creates a grid where all the cells are hexagonal, except for 12 pentagons. We call this hexagonal mesh the dual mesh.

The mass will be associated with the center of the hexagons volume i , and other quantities can be associated to the edges v joining three hexagons, or the edges e joining the centers of two hexagons (see Figure V.6 on page 48).

The generation of this dual mesh is optimized with the Jigsaw mesh generator library, which homogenizes the surface of the cells over the sphere (Engwirda, 2018). This results in a horizontal quasi-uniform icosahedral C-grid.




Figure V.5: Icosahedral grid construction: (a) regular icosahedron, (b) after division of each triangular face into four subtriangles, (c) after decomposition into 42 subtriangles, and (d) the dual pentagonal-hexagonal grid of (c). Staniforth and Thuburn, 2012

4.2.1 Horizontal resolution

The number of total grid points n for a resolution nbp is equal to the number of vertices of the dual mesh:

$$n = 10nbp^2 + 2$$

In practice, only a few resolutions have been extensively tested and studied, ranging from $nbp40$ to $nbp320$. In this work we have only used two horizontal resolutions of DYNAMICO:

- $nbp40$: 16002 cells in total with hexagons of side 110 km. It is the closest to the reference resolution 2.500° in longitude and 1.250° in latitude of LMDZ, at least at the Equator, used for CMIP6.
- $nbp80$: 64002 cells in total with hexagons of side 55 km. It is the closest to the reference resolution 1.250° in longitude and 0.625° in latitude of LMDZ.




Figure V.6: Representation of the mesh of DYNAMICO, with the location of the prognostic and diagnostic variables relative to the mesh. (Taken from Dubos et al., 2015)

4.3 Parallelization of DYNAMICO

DYNAMICO is also parallelized with both MPI and OpenMP. This is done by first dividing the icosahedral grid into at least 10 tiles formed by rhombi. Then each of these tiles can be further subdivided into two directions i and j to create parallelograms composed of $iim \times jjm$ hexagons, see Figure V.7. The optimal performance is obtained when each sub-tile is dedicated to an individual MPI process. If there are more sub-tiles than MPI processes then they can still take care of several at a time. Data transfer between these sub-tiles is done by asynchronous MPI calls. The vertical calculations of the dynamics are then distributed on OpenMP threads.

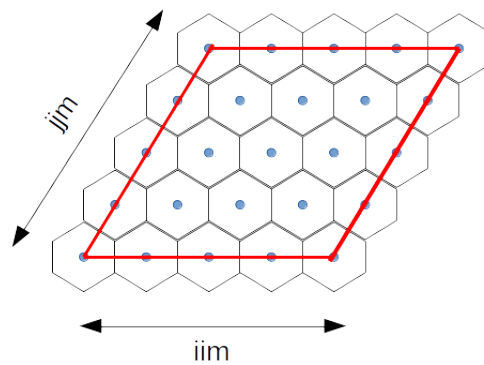


Figure V.7: Subdivision of the primary mesh along the two directions i and j .

The maximum number of MPI processes usable in parallel must be smaller than $10 \times iim \times jjm$. This gives a much higher maximum threshold of MPI processes than when parallelizing the original dynamical core of LMDZ. For example with the `nbp40` resolution and a 4×4 split, we can use up to 160 MPI processes, and for the `nbp80` resolution with a 8×8 split, a maximum of 640 MPI processes. This is much higher than the maximum limit created by the 2 latitude band per MPI process of the original

LMDZ dynamical core which for equivalent resolutions corresponds to 71 and 128 MPI processes.

At the beginning of this PhD, DYNAMICO was already coupled to LMDZ and could freely replace the original dynamical core of LMDZ while still using the same physics described above in [subsection 3.3](#). However, it was not yet properly coupled with the other models for a complete master GCM configuration capable of generating mass fluxes for an offline version. This configuration had not been validated for long-term CO₂ atmospheric transport either and was missing some necessary features. DYNAMICO is able to run on GPU when not coupled with the other models, but this was not used in this thesis. The work carried out to prepare this configuration is presented in [chapter VII](#). The validation of the DYNAMICO-LMDZ coupled configuration of the master GCM for CO₂ atmospheric transport is presented in [chapter VIII](#).

5 Chemistry model: INCA

The atmospheric chemistry model used in the master coupled configurations of the GCM is the Interaction between Chemistry and Aerosol (INCA) model, presented in Hauglustaine et al., 2004 and Folberth et al., 2006. An updated version of the model is also presented in Szopa et al., 2013.

The model takes into account emissions and sink both from natural sources and anthropogenic activities and models tropospheric chemistry through up to 85 chemical species and 264 chemical reactions. It also models dry and wet depositions, aerosols and stratospheric ozone. The concentration fields are updated every 15 minutes when called by LMDZ.

Atmospheric tracers, passive or reactive, can also be modeled by INCA and coupled with LMDZ so that their concentration is stored in INCA while being transported by the relevant atmospheric routines in LMDZ. This is necessary when studying chemically reactive tracers such as CH₄, but passive tracers like CO₂ would behave identically if decoupled from INCA and only used in LMDZ. Nonetheless, even if the subject of this study is to improve the atmospheric transport of CO₂ in our models, validating a configuration with INCA could be very useful to other members of the team that would want to use it. The INCA model is also maintained by Anne Cozic, the co-supervisor of this PhD, who provides an expertise of the model that would be harder to access if only using LMDZ here.

6 Land model: ORCHIDEE

The land model used in the master coupled configurations of the GCM is the Organizing Carbon and Hydrology in Dynamic Ecosystems (ORCHIDEE), presented for the first time in detail in Krinner et al., 2005. Since then, major updates have been made to improve the ORCHIDEE model and follow the evolution of LMDZ. The latest version is described in Cheruy et al., 2020.

The land model simulates moisture in the soil as well as energy exchanges between the soil and atmosphere. These processes are computed at the same time-step as the physics of LMDZ. It can also simulate the evolution of vegetation dynamics across time and space. However, in our case, we use pre-computed yearly land cover maps instead of simulating the vegetation dynamics.

The hydrology is modeled by discretizing Darcy's law describing the flow of water through porous soil over 11 vertical layers (Rosnay et al., 2002). This vertical grid is non-uniform and the grid is smaller near the surface than at the maximum depth of 10 m. A major role of the land model is to control the albedo of the surface as well as the near-surface temperature (i.e., 2 m). The land model also interacts with the atmosphere model by prescribing the roughness heights used in the treatment of the surface layer, depending on the vegetation maps.

Since the influence of the land model on CO₂ atmospheric transport was not the subject of our study, the parameters of ORCHIDEE were kept identical between all configurations, and no coupling of the specific CO₂ components were used.




Figure V.8: Overview of the XIOS architecture. Each model process communicates to a XIOS server using asynchronous MPI messages. (Taken from Yepes-Arbós et al., 2022).

7 Input/Output module: XIOS

All the inputs and outputs of the models described previously are handled by a module called the XML Input/Output Server (XIOS, see <https://forge.ipsl.jussieu.fr/ioserver/wiki>). The performance of XIOS for Earth system models has been evaluated in Yepes-Arbós et al., 2022.

XIOS is an asynchronous parallel I/O server running on dedicated MPI processes. The configuration and management of all inputs and outputs is made easier with the use of XML configuration files. The post-processing and re-gridding of the input or output files can be done directly in-line and configured in the XML files.

The schematic structure of XIOS is presented in Figure V.8 (page 50). Each process of the model runs an individual XIOS client and asynchronously communicates with the XIOS servers via MPI messages when ready. The servers are run on independent nodes and can either communicate between them to aggregate the data and write a single file to storage, or each write an individual file which can be aggregated later. The performance of the XIOS clients and servers can be easily evaluated thanks to a profiler tool. The number of XIOS servers can then be adjusted so that a minimum amount of time is spent waiting.

Virtually all climate models and GCM possess some way of writing outputs in parallel, and their efficiency is becoming crucial now that higher resolutions, more frequent outputs and a growing number of diagnostics are becoming more common. The specific advantage of asynchronous servers is that the writing to disk can be done concurrently to the model running without halting the time steps.

At the beginning of this PhD, XIOS was already fully implemented into all models described above, but not in the offline LMDZ transport model used for atmospheric inversions.

Besides accelerating the input/output process of our models, XIOS was also used throughout this work to interpolate fields, particularly by re-gridding from a regular latitude-longitude grid to the icosahedral grid of DYNAMICO and back. This process can either be done automatically at the runtime by the master GCM, or as a standalone operation by XIOS. This was necessary in particular for creating the files used for the initial state of the ICO configuration, but also for outputting some variables from that configuration on the regular latitude-longitude grid instead of its native icosahedral grid.

Interpolating values between different meshes while ensuring the conservation of physical quantities such as fluxes is a complex task. To do so, XIOS uses a local, explicit, second order, conservative

interpolation algorithm described in Kritsikis et al., 2017. A *supermesh* is created as a common refinement of both the source and target mesh: every cell from either of those meshes is the union of cells from the supermesh. The area of those cells from the supermesh give the weight used for the interpolation, and the gradients, if present, are estimated from the barycenters of the three meshes.

The construction of this supermesh is computationally expensive and computing all the intersections of the source and target mesh should in theory have a quadratic algorithmic complexity $O(N^2)$ for meshes with $O(N)$ cells. In practice, most cells do not intersect, and the supermesh construction can be carried out by a fast search algorithm resulting in a $O(N \log N)$ complexity. This algorithm does not need prior information on the connectivity of the source or target mesh and even recreates that connectivity as part of its initialization. The local interpolation error is bounded by the local second order derivative of the interpolated function and by $O(h^2)$, with h the square root of the cell surface of the target or source mesh, whichever is larger.

Atmospheric inverse modeling: theoretical concept and description of the CAMS/LSCE inverse system

Contents

1	Concept of atmospheric inverse modeling	54
1.1	Definition of the variables and parameters	54
1.2	Error types	55
1.2.1	Prior error	55
1.2.2	Observation error	55
2	Solving the inverse problem	56
2.1	Bayesian formulation	56
2.2	The Gaussian assumption	56
2.3	Cost function	57
2.4	Adjoint operator	57
2.5	Methods for solving the inverse problem	58
2.5.1	Analytical method	58
2.5.2	Ensemble method	58
2.5.3	Variational method	59
3	The CAMS/LSCE inverse system	59

WE PRESENTED THE CONCEPT OF A GCM as well as the specific one used in our study in [chapter V](#). As we explained in [chapter II](#) of [Part I](#), our final goal is to use a top-down approach to estimate the surface fluxes of carbon and their spatial and temporal variations.

To do so, we have to find numerical solutions to an atmospheric inverse problem. The theoretical basis of this problem is presented in [section 1](#) of this chapter. Then, different methods to solve it are explained in [section 2](#), the variational method in particular that we use in this work to solve the inverse problem is presented in detail in [section 2.5.3](#). In [section 3](#), CAMS/LSCE, the inverse system used in this study is described, and we explain how it is linked with the master GCM described in the previous chapter.

1 Concept of atmospheric inverse modeling

The previous chapter presented the master GCM used to resolve a forward (or direct) problem. In our case, it consists of calculating the atmospheric concentrations of CO₂ at a given time as a function of the concentrations at the previous time-step and of specific CO₂ surface fluxes.

Solving an inverse problem, on the other hand, consists in looking at a final state (in our case concentrations of CO₂), and finding the causes (surface fluxes of CO₂) that were responsible for this state.

However, inverse problems do not generally have a unique solution and need to be constrained to provide a range of realistic ones. This is especially relevant for atmospheric transport of greenhouse gases: there are many sources, and tracers are then transported and diluted through complex processes in the atmosphere. The answer to this issue is to constrain the problem by using prior information, on the fluxes we want to estimate and on the greenhouse gas concentration we are transporting.

In our work we have exclusively looked at global inversions of CO₂, which simplifies the problem, since given its long lifetime in the atmosphere compared to our simulations length, we assume that it does not react with other species in the atmosphere.

1.1 Definition of the variables and parameters

In our inverse problem, we assume a deterministic behavior of the atmosphere: a true initial state of CO₂ concentrations, sources and sinks as well as a specific meteorology results in a single final state of concentrations after a given time. The physical space of this atmospheric state and observations need to be discretized numerically and can be expressed by this equation:

$$\mathbf{y}^o = \mathcal{H}(\mathbf{x}^t) \quad (\text{VI.1})$$

Here, \mathbf{y}^o is the observation vector containing all the observations we wish to use. It is a part of the observation space \mathcal{R}^p , with p the number of observations available. \mathbf{x}^t is the vector of the true state and is part of the state space \mathcal{R}^n , with n being the dimension of the problem. They are linked by the observation operator \mathcal{H} which maps the state space \mathcal{R}^n to the observation space \mathcal{R}^p . This equation corresponds to an idealized case where no observation errors exist. This is not the case in our application of the inverse problem and we will describe the impact of these errors in the next subsections.

By solving the atmospheric inverse problem, we are trying to approach the true state \mathbf{x}^t as closely as possible. The solutions corresponding to the posterior state vector \mathbf{x}^a , obtained after optimization, being closer to the true state than the prior state vector \mathbf{x}^b by assimilating information from observations.

1.2 Error types

In reality, both the initial state of the atmosphere and the observed state cannot be perfectly modeled. Taking into account these errors and unknowns, affects the simplified Equation VI.1 (page 54). This error has multiple sources which we will describe in the following subsections.

1.2.1 Prior error

Firstly, the error corresponding to the uncertainty of the prior vector \mathbf{x}^b compared to the reality:

$$\epsilon^b = \mathbf{x}^t - \mathbf{x}^b \quad (\text{VI.2})$$

We associate the prior error covariance matrix \mathbf{B} to this error, with \mathbb{E} the expected value of a random variable:

$$\mathbf{B} = \mathbb{E}[(\epsilon^b)(\epsilon^b)^T] \quad (\text{VI.3})$$

In our case, this is the error associated with the prior fluxes of CO_2 and with the initial state of the CO_2 field. Evaluating and quantifying this prior error covariance matrix is a complex topic and different strategies exist. In the following, we have assigned the diagonal elements of the matrix based on estimates of the uncertainties of CO_2 fluxes from the land model ORCHIDEE (Chevallier et al., 2006) and ad hoc values for the ocean. For the non-diagonal elements we take into account a spatial exponential decrease from the diagonal:

$$e(i,j) = \sigma_i \cdot \sigma_j \cdot \exp\left(-\frac{D(x_i, x_j)}{L}\right) \quad (\text{VI.4})$$

With D the distance between the center of two cells x_i and x_j , σ_i and σ_j the standard deviation at each of those cells and L the correlation length chosen as 500 km over land and 1000 km over the ocean. We also take into account a temporal correlation of four weeks in a similar way.

1.2.2 Observation error

The error associated to the observation space is twofold: the measurements error and the modeling error.

Every measurement or observation of a physical state, be they from in-situ sampling of air or satellite observations of CO_2 for example, are subject to errors. This error can be systematic, which often happens for satellite observations, or be random. A random error in this case means that for two successive measurements of a state that is supposed to be identical, the instrument does not result in the same value measured. When they are significant, statistics of the random errors are often estimated by the data producers for each observation individually from Bayesian principles. A large part of the systematic errors from satellite observations are often already empirically removed by the data producers (Keely et al., 2023), but the remaining part can still impact the posterior state of the fluxes after the inversion.

The modeling error is defined as the differences between an idealized model associating the state space to the observation space and a numerical model doing the same. This modeling error has multiple factors, the first one being representativeness error. The state and observation spaces have to be discretized, turning what is a physical space with infinite dimension into a finite amount of cells where values are averaged in each one. This gives a minimum scale under which there is no variability, creating the representativeness error. This error is directly linked to the spatial resolution of the model and adds substantial uncertainty to the final flux estimates (Tolk et al., 2008).

The modeling error is also due to the advection scheme and the physical parameterization of atmospheric processes. The modeling of convection and advection for example are necessarily different from the true process in the atmosphere. In the end, the total observation error ϵ^o , the sum of the measurement and modeling errors, can be defined as:

$$\epsilon^o = \mathbf{y}^o - \mathcal{H}(\mathbf{x}^t) \quad (\text{VI.5})$$

Its statistics are usually empirically defined, for instance from proxies like the synoptic-scale variability for in situ measurements (Chevallier et al., 2010) or neglected in the case of satellite retrievals.

In the same way as for the prior error covariance matrix, we then define \mathbf{R} the observation error covariance matrix which takes into account the observation errors and the estimated errors from the transport model:

$$\mathbf{R} = \mathbb{E}[(\epsilon^\circ)(\epsilon^\circ)^\top] \quad (\text{VI.6})$$

In practice, we assume it to be diagonal because the observation configuration evolves in time, making it challenging to formulate a computationally efficient correlation model. We even average some of the observations at a preprocessing stage in order to remove correlated errors: for a few local hours in the case of in situ data (Chevallier et al., 2010) and for a few seconds in the case of dense satellite retrievals (Baker et al., 2022).

2 Solving the inverse problem

2.1 Bayesian formulation

The usual method for solving an inverse problem uses a Bayesian framework. This describes how information obtained from the observations can be combined with prior information about the final state to estimate a new probability density function (pdf) of the final state.

The classical formulation of Bayesian inference is:

$$p(\mathbf{x}|\mathbf{y}^\circ) = \frac{p(\mathbf{x})p(\mathbf{y}^\circ|\mathbf{x})}{p(\mathbf{y}^\circ)} \quad (\text{VI.7})$$

also presented in [Figure VI.1](#). With the different components in our atmospheric inverse problem:

- $p(\mathbf{x}|\mathbf{y}^\circ)$ is the probability density function of the posterior. It corresponds to the probability of the state \mathbf{x} , after assimilating the observations.
- $p(\mathbf{x})$ is the probability density function of the prior. It is the initial estimate of \mathbf{x} before any new information is taken from the observations.
- $p(\mathbf{y}^\circ|\mathbf{x})$ is the probability density function of the observations \mathbf{y}° , given a specific state \mathbf{x} .
- $p(\mathbf{y}^\circ)$ is the probability density function of the observations independent of \mathbf{x} and therefore constant.

The best estimate of the solution for our inverse problem given our information from the observations is the one that maximizes $p(\mathbf{x}|\mathbf{y}^\circ)$.

2.2 The Gaussian assumption

A common assumption when setting up an atmospheric inverse problem is to consider the errors described above as Gaussian, or put another way, that they follow a normal distribution. The main argument for this simplification comes from the central limit theorem, which holds that the sum of independent and identically distributed random variables (in practice, the outcome of many error sources) tends to converge to a Gaussian distribution. Given the number of different causes of errors described in [subsection 1.2](#) which are then aggregated into a single random variable, this hypothesis is generally accepted.

A careful look at this hypothesis however shows its limits. A regular normal distribution can include negative values that are not physical for a tracer concentration or for strict emissions. This can be mitigated through various methods such as a truncated Gaussian (Stohl et al., 2009) or using a log-normal distribution (Brioude et al., 2011). But the initial assumption that the error distribution should be considered Gaussian in the first place is even more contested. The prior error distribution of CO_2

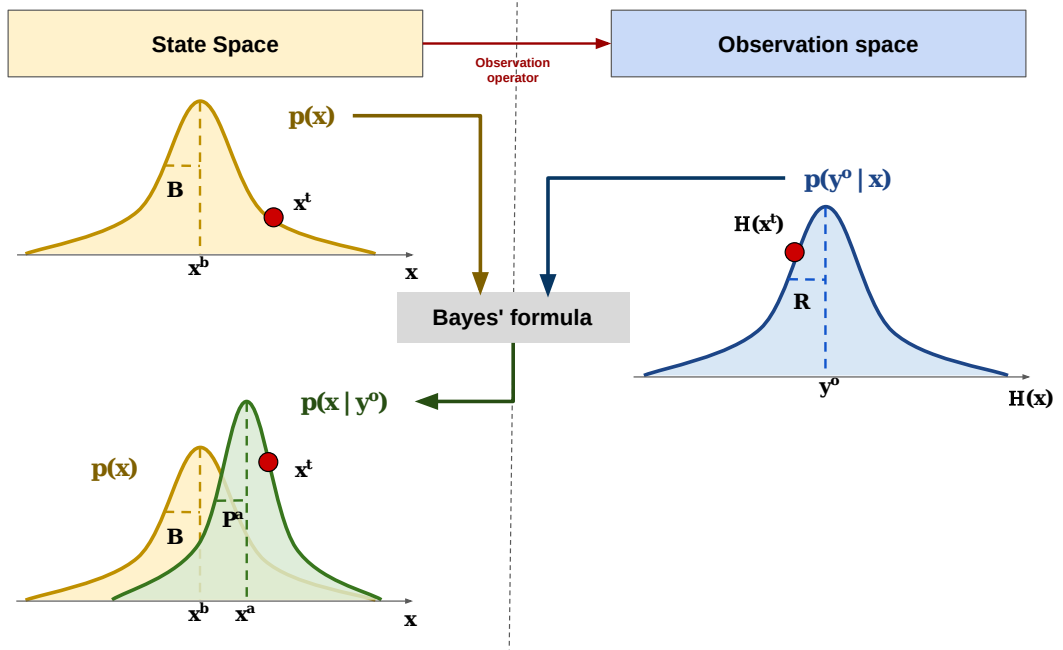


Figure VI.1: Schematic representation of the use of Bayes' formula to solve the inverse problem. With the Gaussian assumption, the pdfs $p(y^o|x)$ and $p(x)$ are fully described by their averages (respectively y^o and x^b) and covariance matrices (respectively R and B); so can $p(x|y^o)$, its average being x^a and its covariance matrix being P^a . From Thanwerdas, 2021.

fluxes for example can vastly differ from a normal distribution (Chevallier et al., 2006). Despite these limitations, the Gaussian assumption greatly reduces the computational cost of numerically solving the inverse problem and is the easiest way to obtain a closed-form solution. It will be used in the rest of this work. We also assume that the errors do not have any systematic bias.

2.3 Cost function

When we apply the Gaussian assumption to the previous Bayesian formulation of the atmospheric inverse problem and assume that errors are unbiased, we obtain:

$$p(x|y^o) \propto \exp \left\{ -\frac{1}{2}(x - x^b)^T B^{-1}(x - x^b) - \frac{1}{2}(\mathcal{H}(x) - y^o)^T R^{-1}(\mathcal{H}(x) - y^o) \right\} \quad (\text{VI.8})$$

$$p(x|y^o) \propto \exp \{-J(x)\} \quad (\text{VI.9})$$

With $J(x)$ being the **cost function**:

$$J(x) = \frac{1}{2}(x - x^b)^T B^{-1}(x - x^b) + \frac{1}{2}(\mathcal{H}(x) - y^o)^T R^{-1}(\mathcal{H}(x) - y^o) \quad (\text{VI.10})$$

Finding the maximum probability of $p(x|y^o)$ is therefore equivalent to minimizing the cost function $J(x)$. In the variational method described in [subsection 2.5.3](#) we aim for the convergence to zero of the gradient of this cost function:

$$\nabla J(x) = B^{-1}(x - x^b) + \mathcal{H}^* R^{-1}(\mathcal{H}(x) - y^o) \quad (\text{VI.11})$$

We introduce here \mathcal{H}^* , the adjoint operator.

2.4 Adjoint operator

The adjoint operator \mathcal{H}^* is derived from the direct observation operator. The direct observation operator is the numerical transport model that links the fluxes in the state space to the atmospheric concentrations (observation space). If this operator is linear then the resulting cost function is

quadratic which greatly simplifies the process of minimizing its gradient, but in practice the numerical transport model introduces non-linearities.

From this non-linear operator, we can create the tangent-linear operator by individually linearizing each operation of the model with a first-order Taylor expansion. The adjoint operator is represented by the transpose of the Jacobian matrix \mathbf{H} of \mathcal{H} corresponding to this tangent-linear operator. This is possible since for a linear matrix, the adjoint is equal to the transpose.

$$\mathbf{H}^* = \mathbf{H}^T \quad (\text{VI.12})$$

2.5 Methods for solving the inverse problem

There are three main methods for solving the inverse problem: the analytical method, the ensemble method and the variational method. The inverse system studied in this manuscript follows the variational method, but in the following sections we will describe them all.

2.5.1 Analytical method

If the observation operator is linear then there is a direct way of obtaining an optimal solution, with a state \mathbf{x}^a and posterior error covariance \mathbf{A} :

$$\mathbf{x}^a = \mathbf{x}^b - \mathbf{B}\mathbf{H}^T(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})^{-1}(\mathbf{H}\mathbf{x}^b - \mathbf{y}^o) \quad (\text{VI.13})$$

$$\mathbf{A} = \mathbf{B} - \mathbf{B}\mathbf{H}^T(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})^{-1}\mathbf{H}\mathbf{B} \quad (\text{VI.14})$$

Which can be rewritten as:

$$\mathbf{x}^a = \mathbf{x}^b - (\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H} + \mathbf{B}^{-1})^{-1}\mathbf{H}^T\mathbf{R}^{-1}(\mathbf{H}\mathbf{x}^b - \mathbf{y}^o) \quad (\text{VI.15})$$

$$\mathbf{A} = (\mathbf{B}^{-1} + \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H})^{-1} \quad (\text{VI.16})$$

This method provides an easy way of calculating \mathbf{x}^a and its associated uncertainty \mathbf{A} if we have previously calculated \mathbf{H} . Directly calculating this matrix carries a significant computational cost when the dimension of the inverse problem is large, which is the case with atmospheric inversions at the global scale (Bousquet et al., 2000). Moreover this method also requires directly inverting matrices of the size of \mathbf{B} or \mathbf{R} depending on which formulation of the equations is chosen. These matrices are respectively of size $n \times n$, the dimensions of the inverse problem and of size $p \times p$ the square of the number of observations. Even if assumptions can be made that the covariance error matrices are diagonal (which comes with its own drawbacks), directly inverting matrices of such size is prohibitively expensive and most often not the preferred method.

2.5.2 Ensemble method

The ensemble methods are a subset of the sequential Monte Carlo methods, also called particle filters. The particles here refer to state vectors, part of an ensemble of a given dimension m , that vary in a manner representative of the state and uncertainty of the system. The observations are assimilated sequentially, with a running computation window, ensuring that only a small number of observations are assimilated at a time thereby reducing the size p of the observation space. This assimilating window is then moved forward step by step until the end of the inversion. For this to be valid we assume that the errors of the observations are independent from each others since they propagate from one time window to another.

One of the most popular ensemble method is the Ensemble Kalman Filter (Chatterjee et al., 2012, Kong et al., 2022), which is a type of particle filter that assumes a Gaussian distribution of all the probability distributions, which greatly simplifies the inverse problem. Rather than trying to directly calculate the prior error covariance matrix \mathbf{B} , the full products $\mathbf{B}\mathbf{H}^T$ and $\mathbf{H}\mathbf{B}\mathbf{H}^T$ are instead approximated

as:

$$\mathcal{H}\mathbf{B}\mathbf{H}^T \simeq \frac{1}{m-1}(\mathcal{H}(\mathbf{x}_1), \mathcal{H}(\mathbf{x}_2), \dots, \mathcal{H}(\mathbf{x}_m)) \cdot (\mathcal{H}(\mathbf{x}_1), \mathcal{H}(\mathbf{x}_2), \dots, \mathcal{H}(\mathbf{x}_m))^T \quad (\text{VI.17})$$

$$\mathbf{B}\mathbf{H}^T \simeq \frac{1}{m-1}(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m) \cdot (\mathcal{H}(\mathbf{x}_1), \mathcal{H}(\mathbf{x}_2), \dots, \mathcal{H}(\mathbf{x}_m))^T \quad (\text{VI.18})$$

The approximated values of these products can then be used to analytically solve the inverse problem as seen in the previous section avoiding the most computationally expensive matrices inversions. This method is highly parallelizable and well suited to run on large supercomputers, allowing comparatively quick results.

2.5.3 Variational method

The variational method is the preferred method in this thesis. We rewrite the previously defined cost function in equation (VI.10) with a linear hypothesis:

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}^b)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^b) + \frac{1}{2}(\mathbf{H}\mathbf{x} - \mathbf{y}^o)^T \mathbf{R}^{-1}(\mathbf{H}\mathbf{x} - \mathbf{y}^o) \quad (\text{VI.19})$$

As explained in subsection 2.3, we want to solve the inverse problem by minimizing the cost function in an iterative manner. This is equivalent to reducing the norm of the gradient of the cost function $\nabla J(\mathbf{x})$ with:

$$\nabla J(\mathbf{x}) = \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^b) + \mathbf{H}^T \mathbf{R}^{-1}(\mathbf{H}\mathbf{x} - \mathbf{y}^o) \quad (\text{VI.20})$$

This method requires inverting the prior error covariance matrix \mathbf{B} , but this computationally expensive operation is possible because the matrix is either diagonal if neither temporal correlations nor spatial correlations are taken into account, or is sparse if they are included. The rest of the matrix products necessary to calculate $\nabla J(\mathbf{x})$ would however still be prohibitively expensive if done directly. To solve this, we use the adjoint model described in subsection 2.4.

Instead of explicitly calculating each element of \mathbf{H} , the adjoint model uses the chain rule to decompose \mathbf{H}^T within $\mathbf{H}^T \mathbf{R}^{-1}(\mathbf{H}\mathbf{x} - \mathbf{y}^o)$, which is used to calculate the gradient of the cost function $\nabla J(\mathbf{x})$. The decomposition is done line by line of the forward model \mathcal{H} , so that only elementary (hence small) Jacobian matrices are processed at a time.

The function is minimized iteratively according to quasi-Newtonian algorithms, such as the MIQN3 software (Gilbert and Lemaréchal, 1989) or the Lanczos version of conjugate gradient algorithm (CON-GRAD, Fisher, 1998). The iterative process is then repeated until either the norm of the gradient $\nabla J(\mathbf{x})$ has been reduced by a fixed amount, or when a fixed number of iterations have been carried out. This gets us a value of the posterior state vector \mathbf{x}^a optimized given the information from the observations \mathbf{y} .

3 The CAMS/LSCE inverse system

The CAMS/LSCE inverse system has two main components:

- An offline version of an atmospheric transport model based on LMDZ described in chapter V. In addition to this direct offline transport model, its tangent-linear and adjoint are also coded in Fortran and integrated in the inverse system in order to solve the inverse problem with a variational method. The offline version of the transport model only solves tracer transport equations and is driven by pre-computed air mass fluxes from the master GCM.
- The PYVAR module, coded in Python, which handles the inversion process itself. This module parametrizes the different components of the inversion: the error covariance matrices, initial

states, prior fluxes and assimilated observations. It then performs the optimization according to a 4D variational method, in our case following the CONGRAD algorithm.

The first version of this inverse system was created in 2005 (Chevallier et al., 2005), and has been maintained and improved ever since, following the developments of the master GCM and its atmospheric model LMDZ (Chevallier, Bréon, and Rayner, 2007, Chevallier et al., 2010, Chevallier, 2013). It can be used in conjunction with either in-situ measurements of CO₂ concentrations, or with satellite observations of total CO₂ atmospheric columns. This inverse system was also modified to be used for atmospheric inversions of N₂O (Thompson et al., 2014), CH₄, CO and H₂ (Berchet et al., 2021, Pison et al., 2009).

At the beginning of this thesis, the CAMS/LSCE inverse system ran on a regular grid with $96 \times 96 \times 39$ cells, equivalent to 3.75° in longitude by 1.875° in latitude. In Part III of this work, we present two successive increases in resolution of the operational CAMS/LSCE inverse system:

- The first resolution increase is presented in chapter IX resulting in a grid with $144 \times 143 \times 79$ cells equivalent to 2.50° in longitude by 1.27° in latitude with 79 vertical levels. The effect of the increase in resolution on the transport model had been studied previously by Remaud et al., 2018, but its implementation in the inverse system had been postponed until this work due to the high computational cost it entails.
- The second resolution increase is presented in chapter X resulting in a grid with $256 \times 256 \times 79$ cells equivalent to 1.41° in longitude by 0.70° in latitude over the same 79 vertical levels. The impact of this increase in resolution on the global atmospheric CO₂ inversions is directly studied in the chapter.

The next part of this manuscript will present the scientific findings and the developments carried out in this work. In chapter VII, we start by describing the preliminary developments necessary for the following of the studies.

Part III

Scientific findings

Preliminary developments

Contents

1	Initial state of the models	64
1.1	Master GCM	64
1.2	Inverse system	64
1.3	Technical challenges	65
1.3.1	Computational cost	65
1.3.2	File size	66
2	Developments	66
2.1	Nudging	66
2.2	Mass fluxes	67
2.3	Tracers	68
2.4	Using XIOS in the inverse system	68
3	Dealing with the icosahedral grid	69
3.1	NetCDF files	69
3.2	Comparison to a regular grid	69
3.3	Drawing maps	70

THIS CHAPTER PRESENTS THE DEVELOPMENTS and work that were necessary to make the new ICO coupled configuration using DYNAMICO presented in [section 2](#) of [chapter V](#) functional. This work was needed for the study in [chapter VIII](#), but did not make it to the final publication. The initial state of the GCM and inverse system is described in [section 1](#), as well as some of the technical challenges that were faced in the course of this work and how they were solved. Then in [section 2](#), some of the features added to the GCM are described.

1 Initial state of the models

1.1 Master GCM

At the beginning of this thesis, the master GCM in its configuration coupling the atmospheric model LMDZ, land model ORCHIDEE, and chemistry model INCA was used to generate mass fluxes for the inversion system at a low resolution of $96 \times 96 \times 39$ (3.75° by 1.875°). This resolution was already being slowly phased out for a lot of applications and we wanted to exclusively use the higher $144 \times 143 \times 79$ resolution (1.40° by 0.70°) to validate the coupled configuration using DYNAMICO and for the inversions of CO_2 .

The version of the GCM at the $96 \times 96 \times 39$ resolution was using versions of the models that were no longer evolving. In this thesis, we wanted to use the latest developments in each of the GCM components, and we also wanted to be able to switch easily from one version of the dynamics to the other. To this end, we have created a specific model configuration that groups together both versions of the dynamics and can thus be used to run the REG and ICO configurations described previously in [section 2](#) of [chapter V](#). We could then choose for each simulation which model and dynamics to use while only maintaining one code base.

At the beginning of this thesis, DYNAMICO was already coupled with the LMDZ atmospheric model thanks to a specific module linking the dynamics of DYNAMICO to the physics of LMDZ. However, this coupling did not yet work when integrated into the master GCM. It was also missing some essential functionalities for our work and future integration in our inverse system, notably proper integration of tracers and nudging of atmospheric variables.

A coupled configuration of the GCM with so many models and parameters, able to run on two different dynamical cores and grids required extensive testing and much trial and error to find a good setup. These will be discussed in [section 2](#).

1.2 Inverse system

As explained above, the CAMS/LSCE inverse system originally ran on a $96 \times 96 \times 39$ grid. The advantage of an inverse system based on a dedicated offline version of a larger atmospheric model is the relatively small size of the model code which can be limited to only the essential atmospheric transport routines and optimized for the inversion process. However, in our case the inverse system had been developed over almost two decades and suffered from a huge technical debt on the GCM side since it was not up to date with recent LMDZ technological developments. It was not synchronized with the LMDZ atmospheric model and could not naturally follow its future developments. Even if the results in tracer transport of LMDZ coupled with DYNAMICO would prove to be successful, the inverse model was at the time not ready to make the switch in the same manner without a major rewrite as explained in [chapter IV \(Part I\)](#).

1.3 Technical challenges

1.3.1 Computational cost

Running large-scale atmospheric inversions of CO₂ already required a rather large amount of computing resources at the original resolution. It quickly became clear when extrapolating the time required to run these inversions at the new required resolution, which increased both horizontally and vertically, that this would not be sustainable. Inversions are run month by month, over large periods of time, often decades, and take many iterations. At the 96 × 96 × 39 resolution, run in parallel on 15 CPUs, each month took roughly 5 minutes real-time in total. This includes both the tangent-linear and the adjoint of the transport model. Our global inversions are usually done over 40 iterations, which means that in total, a 10-year-long inversion at this resolution takes around 400 hours in real time or a bit more than 2 weeks. Increasing the resolution to 144 × 143 × 79 would increase the size of the grid by a 4.5 factor. At the same speed, a decade-long global inversion would take more than two months. In practice, inversions of more than three years can be parallelized further by running multiple years concurrently, bringing the wall clock time down to that of a 1.5-year inversion. This is still too long compared to the time constraints required by CAMS and too computationally expensive in terms of CPU hours, and the goal of sub 1° resolution is strictly impossible with this performance.

The inversions carried out in this thesis were done on *Obelix*, the computing cluster of the LSCE. This cluster is composed of 46 nodes: 12 of them with 24 CPUs and 36 of them with 20 CPUs for a total of 1008 CPUs. It is now also equipped with 4 GPUs. These inversions require pre-computed mass fluxes generated by the master GCM presented in [chapter V \(Part II\)](#), and therefore need an additional computing step. This initial simulation is much more computationally expensive since it runs on a fully coupled configuration, not just an atmospheric transport model.

The simulations of the master GCM carried out for this work were all done on the Joliot-Curie (also called Irene) supercomputer, part of the CEA's Très Grand Centre de Calcul (TGCC), this is also where the products for the CAMS project are carried out. This supercomputer possesses a theoretical peak processing capacity of around 23 petaflops and is ranked as the 82nd most powerful supercomputer in the world as of June 2024 (<https://www.top500.org/system/179412/>). The Irene supercomputer is divided into multiple partitions, each with different hardware characteristics. The main partition used in this work for simulations of the master GCM is *Skylake*. This partition has 1 653 nodes, with 48 cores per node for a total of 79 344 cores. The RAM is limited to 180 GB per node which is often too low for post-processing purpose where you need to store large amounts of data at the same time in the memory. For this, we use instead the *XLarge* partition which only has 560 cores but each where each node has a RAM of 3TB. Skylake is the oldest partition of the Irene supercomputer, and the GCM was already used extensively on it which means that the compilation of the code running on the Intel Skylake CPUs was robust and already optimized.

At the end of 2019, a year before the beginning of this thesis, a new partition was added to Irene: *AMD Rome*. This partition has 2 286 nodes, with 128 cores per node for a total of 292 608 cores. We initially wanted to use this partition for the majority of our simulations. It had the advantage of being relatively new and therefore had few users, rendering the traditionally long queuing wait times of supercomputers almost non-existent. Our team was also able to easily get a large number of CPU hours on it compared to busier partitions since much of the available capacity was unused. However the actual use of this partition turned out to be very problematic. The CPUs of this partition are from AMD instead of Intel, and while this is not inherently a problem nor are they of inferior quality, it means that compiling the GCM's Fortran code to run on them requires a lot of work. This work was graciously done by people at the IPSL but at the time still necessitated more testing. Added to that was the instabilities of the partition itself, and the very frequent maintenances, sometimes lasting more than a week. Simulations would often get stuck indefinitely for no apparent reason or be excessively slow. Because of these problems, we chose to exclusively use the Skylake partition from June 2021 onwards.

Resolution	96×96 ×39	144×143 ×79	nbp40	256×256 ×79	nbp80
Atmospheric model	35 GB	160 GB	121 GB	-	-
Chemistry model	265 GB	1.2 TB	1 TB	-	-
Air mass fluxes	250 GB	1.1 TB	820 GB	3.3 TB	3.3 TB

Table VII.1: File size per year of the outputs of the different models (LMDZ/Dynamico for the atmospheric model, INCA for the chemistry) of the GCM. The outputs of the models are hourly, the mass fluxes have a 3 hour time-step.

Simulations on the GCM were usually run on more than 300 cores, with each month simulated taking around 30 to 40 minutes depending on the configuration and the output frequency. To this must be added the time spent queuing until each computing job is accepted, which can be significant when using hundreds of cores at the same time. This high computational cost meant that each test took a long time and errors requiring a re-run of a simulation were large setbacks. The total amount of CPU hours used for this work on Skylake in 2023 for example was around 500 000 hours, comprising simulations done for research purpose and generating mass fluxes used for the production of CO₂ inversions for CAMS. With future resolution increase, these constraints of CPU hours consumed and long time-to-solution will only become more significant.

1.3.2 File size

Another issue to deal with when running simulations of the master GCM is the large size of the output files. In [Table VII.1](#) we present the average file size for yearly outputs at different resolutions. These are the typical sizes for the outputs of the simulations in [chapter VIII](#), and of the tests carried out to prepare them.

The size of these files means that post-processing also takes a long time and must be done in parallel on computing clusters. To do so, we have mainly used Python, and in particular the Xarray and Dask libraries. This provides an efficient way of reading NetCDF files of any dimensions, selecting the data and then operating over it very easily. When paired with Dask, the arrays are divided into small chunks and operations can be queued as series of tasks mapped over these chunks. The data is only loaded into memory when the computation is needed and is then threaded over multiple processors. This parallelization necessitates no extra manual work which makes it very easy to integrate into post-processing scripts. This method was necessary since the post-processing nodes have a limited memory size, as explained in the previous section, and the full files usually cannot be loaded into memory. It also reduces the time spent on the post-processing, which would be too high if done sequentially. This is particularly the case when sampling the CO₂ concentration fields to compare them to observations.

The large size of the air mass fluxes files leads to the need of compressing them for their use as inputs in our inverse system, which was done and presented in [chapter IX](#).

2 Developments

2.1 Nudging

An important functionality of the master GCM is the *nudging*: the ability to guide specific variables towards pre-defined values, usually atmospheric fields from a meteorological reanalysis, described back in [section 3](#) in [chapter V](#). For atmospheric transport of tracers, in our case CO₂, nudging the wind fields to a reanalysis greatly improves the accuracy of the tracer transport since horizontal winds are one of the largest source of error in the observed CO₂ variability (Agustí-Panareda et al., 2019). In particular, our global atmospheric CO₂ inversions made for CAMS use air mass fluxes from a simulation

of the master GCM where winds are nudged to the ERA5 reanalysis. This is the case for the inversions in [chapter X](#) for example.

The regular latitude-longitude configuration of the GCM already had this nudging functionality for variables such as the wind speeds, temperature or surface pressure. In the icosahedral configuration using DYNAMICO however, it did not exist yet and had to be developed in time for this work. This guided mode is also usually available at either the global scale or only restricted to an area defined by latitude and longitude boundaries to drive a Limited Area Model (LAM).

2.2 Mass fluxes

The generation of the air mass fluxes by the master GCM was previously present in the regular latitude-longitude configuration for the old resolution. But, as explained in [section 1](#), this was done in a now obsolete version-controlled branch of the development tree of the master GCM and not integrated into the main version of LMDZ. We therefore had to redo this work since our REG configuration was recreated from the main version, both the configuration of this *offline output* mode and the output of these mass flux files had to be rewritten and adapted. Since we wanted to eventually also be able to generate mass fluxes from the ICO configuration, we had to integrate the parametrization of the flags governing the activation of this mode into another module instead of directly into LMDZ as was the case previously. This module called *ICOSA-LMDZ* handles the coupling of the dynamics from DYNAMICO and the physics of LMDZ in the ICO configuration. The parametrization in this module configures the activation and output frequency of the mass fluxes and is now correctly piloted by definition files like the rest of the GCM parameters instead of only being defined in the code.

The mass fluxes are divided into two parts: the variables derived from the dynamics and the ones from the physics.

Physics

Since the physics is identical for the two configurations of the GCM, the implementation of this output only had to be done once in LMDZ. The variables are first accumulated over the time steps of the physics for the duration chosen as the output frequency (3 hours as a default), then divided by this duration to give an average value. The output variables depend on the desired convection scheme, in our case the Emanuel scheme. The thermal plumes in the deep convection are described by multiple variables governing for example the updrafts or saturated downdrafts, the entrainment or the adiabatic ascent. The variables used to describe the thermal plumes in the deprecated Tiedtke scheme (not described here) were also kept in an effort of backward compatibility. The variables describing the physics are a lot heavier than the ones describing the dynamics, most of this is due to a 2D array of exchange coefficients in each column that potentially links all vertical levels together. It is therefore saved as a global 5D variable. In practice we ignore the layers very high in the atmosphere which have a negligible mixing from this process (cutoff at 50 layers). This still leaves 50 instances of a 3-hourly 4D variable at the model resolution. However, in practice each layer only mixes with few other layers around it in a time step: a thermal plume originating from the surface will not reach the stratosphere in that time. This means that most of the values of this variable are zeros and it can be assimilated to a sparse matrix. The outputted files can therefore be compressed a lot and without loss, a process which is presented in [chapter IX](#). Without this step, using these monthly input files weighing 100 GB or 300 GB (see [Table VII.1](#) on page 66) in the inverse system would be impossible with realistic runtimes.

Dynamics of LMDZ

The variables describing the dynamics of the air mass fluxes are necessarily different between the two configurations since they use different dynamical cores. When using the dynamics of LMDZ we use the mass flux of the advection in the zonal and meridional directions as well as the air mass, geopotential and potential temperature in each cell. These variables are outputted at the same frequency as the ones from the physics. These variables are altogether a small portion of the size of the mass fluxes.

Dynamics of DYNAMICO

The creation and development of an *offline* version of DYNAMICO was carried out in parallel to this work, as explained back in [chapter IV](#) of [Part I](#). One part necessary to make this version work that I took charge of, was the generation of the mass fluxes to be read by the future offline version. Developing this version was a constant back-and-forth between trying which variables could be outputted or not, identifying the variables that were essential and those that could be derived and recalculated from simpler ones.

A major difference of variables on the icosahedral grid is that some of them are defined on the cell edges instead of the cell centers. This can be the case for example with wind fields, or horizontal tracer flux. These variables can not be outputted easily, and moreover would be three times larger than equivalent variables at the cell center. However each of these variables can also be re-centered and divided into two standard components: an eastward and northward one. Once read by the offline version of the transport model, they can be converted back onto the cell edges and used normally.

Initially, the mass flux output consisted of 4 variables: the wind fields and the horizontal tracer flux at the cells' edges, the vertical tracer flux, and the air mass at the centers of the cells. But for the final version of the offline mode of DYNAMICO made operational at the very end of this thesis and presented in [Part IV](#), we settled on a much simpler set of variables. The wind fields are outputted, divided into two components at the cell centers along with the surface pressure at each cell center. These two input variables alone are enough to run the transport model (associated with the physics of LMDZ described above). In our case calculations are generally cheaper computationally than reading data from disk making our use of a minimal amount of variables ideal.

2.3 Tracers

As presented in [chapter V \(Part II\)](#), the master GCM used configurations coupling an atmospheric model, a land model and INCA, an atmospheric chemistry model. While the use of a chemistry model was not strictly necessary to study the atmospheric transport of CO₂ we could use it to track the CO₂ tracers concentration and make sure that our results would be generalizable to reactive tracers, eventually paving the way for inversions of CH₄ at new higher resolutions for example.

The modeling of tracers has to be synchronized between INCA and the atmospheric model so that tracers are moved by the transport model after having their concentrations updated in each cell according to their respective local sources and sinks. For the REG configuration of the master GCM this was already operational from the start, but for the ICO configuration it was not the case: tracers were stored differently in DYNAMICO compared to LMDZ and INCA and could therefore not be coupled.

In LMDZ and INCA, tracers and isotopes are identified by name, type and phase. They can also have so-called children, linking species that descend from another one. This was different in DYNAMICO: tracers were all stored in a variable q of size $nqtot$, the total number of tracers, and could only be accessed by their index number. This meant for example that tracer selection by name using a parser was not implemented yet, and that the position of each tracer in the memory had to be known in advance. The definition of tracer types in DYNAMICO was rewritten to match the one used in LMDZ and INCA. This way, the storage of tracers in the GCM is unified and can be controlled from any of the models interchangeably. These tracers are now also defined from a file *tracer.def* identical to the one used for LMDZ. The quality of the representation of the tracer transport with the ICO configuration could then be studied, which will be the main topic of [chapter VIII](#).

2.4 Using XIOS in the inverse system

We presented the input/output server XIOS back in [section 7](#) of [chapter V](#), and it was already integrated into the master GCM at the beginning of this work. While it necessitated some trial-and-error

to optimize its parametrization for use with our configurations, the biggest question regarding XIOS was whether or not to use it in the CAMS/LSCCE inverse system.

The main advantage of XIOS is that it allows asynchronous reading and writing of files as well as integrating post-processing directly in the runtime. While the time taken to write the outputs is not really an issue in our inverse system, the non-computing bottleneck for now is reading the files from disk, and transferring that data. XIOS could be useful in that regard but re-writing the input/output system to integrate it would have been costly. One of the reason this is so difficult is the fact that XIOS has been developed with an assumption of increasing time steps, while the adjoint works backward in time. This small but important difference prevented us from simply replacing the input routines with a ready-made version of XIOS.

3 Dealing with the icosahedral grid

Since DYNAMICO runs on an icosahedral grid, the usual tools used for post-treatment of climate data are not always suited for the task. These programs or Python libraries were often only made with regular latitude-longitude grids in mind and have to be either replaced or adapted to work on our new grid.

3.1 NetCDF files

At the beginning of this thesis, only a few tools existed that were specifically suited for the study of climate data on an icosahedral grid, and I instead chose to stay with the Python environment which had more technical support available. When variables are outputted into a NetCDF file on the icosahedral grid by DYNAMICO, the default horizontal dimension is represented by an unordered list of coordinates corresponding to each cell. The horizontal coordinates can be associated to their corresponding longitude and latitude which are also outputted to locate in space all the hexagonal cells.

We briefly discussed the use of the Xarray python library in [subsection 1.3.2](#), for most of our work it became our tool of choice for opening and manipulating NetCDF files whether on the regular latitude-longitude or icosahedral grid. This library is able to easily compute pre-determined operations on labeled multi-dimensional arrays on the regular latitude-longitude grid but is flexible enough to manually adapt those operations to the icosahedral grid.

3.2 Comparison to a regular grid

The first and most obvious way of comparing simulations on two different grids is to simply re-grid one of them so they are identical. For this goal, XIOS is a tool that can not only accelerate the input/output process, as described back in [section 7 of chapter V](#), but it can also directly integrate post-processing and re-gridding of variables during this process. This means that the outputs of simulations using the ICO configuration of the GCM can be automatically re-gridded to a regular latitude-longitude grid of a similar resolution.

This approach, though it is the easiest and most straightforward one, does not allow us to see the direct local impact of a grid change on CO₂ tracer transport. This is the reason why besides for marginal use at the beginning of this thesis to check the validity and stability of our simulations, we almost exclusively outputted fields on the native icosahedral grid.

When using these different grids, it is impossible to directly create maps of the difference between two simulations, which is often a very common way of comparing them, to find areas of interest for example. It is of course always possible to visually assess two side-by-side maps on different grids, which is often enough to spot if any large-scale trend differs between the two grids.

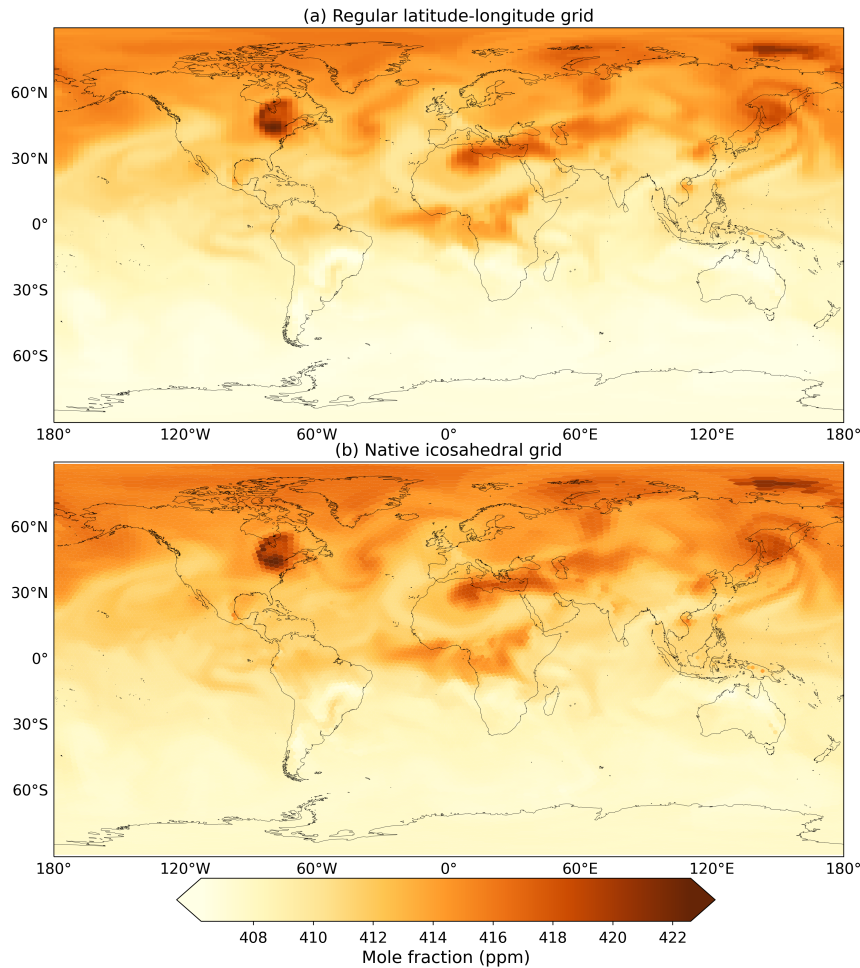


Figure VII.1: Instantaneous field of CO₂ mole fractions in ppm on a re-gridded latitude-longitude grid (a) and the native icosahedral grid (b) at the 25th vertical level of the model. The re-gridding process was done automatically at the output step of the model by XIOS.

For example, [Figure VII.1](#) shows the automatic re-gridding of a CO₂ molar fraction field by XIOS from the native icosahedral grid at the nbp40 resolution to the regular 144 × 143 resolution. It is immediately apparent that these two fields are essentially identical, which is consistent with a re-gridding operation, though this comparison alone is not sufficient to assess the validity of this operation at a small scale.

Manual zonal cuts and averages, however, can still be computed on the icosahedral grid and then directly differentiated to identical ones from the regular grid to obtain quantifiable data.

3.3 Drawing maps

Visualizing and drawing maps of fields on the icosahedral grid was not a straightforward process either. At the time of writing, there exist many tools to easily draw such maps from NetCDF files, for example, [Psyplot](#) (Sommer, 2017) for a standalone software or [Geovista](https://pypi.org/project/geovista/) (<https://pypi.org/project/geovista/>) when staying in the Python environment. These were not available for most of the duration of this thesis and we instead had to use workarounds with the more standard [Cartopy](#) Python library.

To do so, we used the *Collections* class of [Matplotlib](#), to draw the large amounts of polygons recreated from the list of longitude and latitude boundaries of the cells. This is how panel (b) of [Figure VII.1](#) on page 70 was created, but this process is extremely slow compared to newer methods specifically tailored for this purpose and takes several minutes per drawing.

Direct modeling with DYNAMICO

Contents

1	Motivation	73
2	Outline of the study	73
2.1	Setup of the study	73
2.2	Summary of the results	74
2.3	Conclusion of the study	75
3	Introduction	76
4	Presentation of the model and experiments	78
4.1	Configuration description	78
4.1.1	General description of the two models	79
4.1.2	ICO - New icosahedral grid configuration	81
4.2	Description of the simulations	82
4.3	Observational data	83
4.4	Evaluation methodology	84
4.4.1	Surface stations	84
4.4.2	Annual gradient between stations	85
4.4.3	Seasonal cycle	85
4.4.4	Synoptic variability	85
4.4.5	ERA5	85
4.4.6	Aircraft measurements	85
4.4.7	AirCore measurements	85
5	Results and discussions	86
5.1	Mass conservation	86
5.2	Computational efficiency	86
5.2.1	Computational setup	86
5.2.2	Computational gains	88
5.3	Vertical temperature profiles	88
5.4	Seasonal analysis	91
5.4.1	Annual gradient	91
5.4.2	Seasonal cycle	94
5.4.3	Synoptic variability	95

5.5	Vertical profiles of CO ₂ mole fractions	96
5.5.1	Troposphere	96
5.5.2	Low stratosphere	96
6	Conclusion	99
7	Code and data availability	99
8	References	100

THIS CHAPTER CONSISTS OF A STUDY looking at the impact on atmospheric transport of CO₂ when using DYNAMICO as a dynamical core in our master general circulation model. A direct simulation over 40 years was run for two configurations of the GCM, one using the regular latitude-longitude grid of LMDZ for reference (REG), and one on the new icosahedral grid using DYNAMICO (ICO).

The motivation for this study is presented in [section 1](#). A short summary and outline of the study and article is then presented in [section 2](#). The complete article that was created from this study is then presented in the final sections of this chapter.

This article was initially meant to be published in the Geoscientific Model Development Journal (GMD) but was rejected after the review process because the study was focused on the evaluation of the ICO configuration for CO₂ transport instead of the model development of the DYNAMICO-LMDZ coupling itself. It is planned to resubmit it to a more appropriate journal.

1 Motivation

In [chapter VII](#), we established and made functional a new configuration of the master GCM using DYNAMICO as its dynamical core. A crucial step to achieve if we want to make this new configuration usable by the wider scientific community is to validate it against a reference configuration.

With this study we aim to answer several questions:

- Does this new configuration using DYNAMICO perform as well or better at CO₂ atmospheric transport than the reference REG configuration at the global scale?
- Does the absence of grid singularity at the poles in the ICO configuration lead to local improvements in CO₂ atmospheric transport?
- Does the new configuration perform or scale better computationally than the reference?

2 Outline of the study

To validate and evaluate our new configuration of the master GCM we run several direct simulations, either using the reference LMDZ dynamical core on a regular latitude-longitude grid or using DYNAMICO on its icosahedral grid. We then compare the output of these simulations to independent measurements of CO₂ concentrations and see how they perform. We also directly compare the performance of the configurations against one another.

2.1 Setup of the study

The main simulations to evaluate the quality of atmospheric CO₂ transport were run over the 1979–2020 period. The two different configurations were run with identical initial and boundary conditions and had hourly outputs. The large-scale atmospheric circulation of these simulations was nudged to the horizontal winds of the ERA5 reanalysis. CO₂ surface fluxes were prescribed every 3 hours and taken from an atmospheric inversion of the Copernicus Atmosphere Monitoring Service (CAMS, version 20r2).

We evaluated the atmospheric CO₂ transport by comparing the outputs of CO₂ mole fraction to independent measurements at surface stations around the globe. In particular, we extracted growth rate, average seasonal cycle, and synoptical variability from the time-series at each station and analyzed them. We also compared the vertical profiles of temperature and CO₂ mole fraction from the models to data from ERA5 and measurements from aircrafts and AirCore.

We also run short-term simulations in order to evaluate the computational gains of the ICO configuration in a more optimized setup. These runs use a variety of CPU numbers and MPI/OMP parallel setups (see Table VIII.1, page 87). We also evaluated the difference in computing speed when the configurations run at a higher resolution, around double that of our main simulations.

2.2 Summary of the results

The annual gradients of CO₂ mole fraction growth at surface stations relative to the SPO reference station have a bias of less than 0.01 ppm per year for both configurations compared to measurements. The bias is slightly larger when restricted to high latitudes, around 0.1 ppm per year, but there is again no significant difference between the two configurations. Furthermore, there is no difference between them when grouping the bias by latitude bands.

The average seasonal cycle modeled at surface stations shows a strong correlation to measurements for both configurations, with only 3 stations not exceeding a correlation coefficient of 0.8. The amplitude of this seasonal cycle is also well captured.

The modeling of the synoptic variability only has a mean correlation coefficient of 0.54 for both configurations. The ICO configuration has a lower but non-significant normalized standard deviation than the regular configuration.

The vertical temperature profile is different for each configuration with both of them differing from ERA5 profiles by 2 to 5 K. The difference in temperature at the stratopause between the two configurations is highest during winters at high latitudes when driven mainly by gravity waves.

Both configurations show very similar CO₂ vertical profiles compared to aircraft measurements up to 15 km, with a general negative bias of around 1 ppm. Above 15 km the vertical profiles differ, with the DYNAMICO configuration having an increased negative bias compared to the reference. This is also confirmed by comparing the model to AirCore measurements.

Both configurations lose some amount of mass over time, but the value is low enough to have a negligible impact in our study. In the main simulation, running on 384 CPUs and 336 CPUs for the REG and ICO configuration respectively, the ICO configuration is faster by around 20% which is the same order of magnitude as the reduction in the number of cells from the change of grid. When optimizing the number of CPUs for each configuration in their ideal state the number of CPUs hours consumed per month simulated by the ICO configuration is around 50% lower than the regular configuration.

The model was then tested at a higher resolution: 1.4° in longitude and 0.7° in latitude for the regular configuration, and hexagons of side 55 km for the ICO configuration. At this resolution, the regular configuration can use up to 1024 CPUs but is not faster than at a lower scale. The ICO configuration can use up to 2560 CPUs which can bring the monthly time-to-solution down to 323 seconds, compared to the 830 seconds of the regular configuration.

2.3 Conclusion of the study

When comparing the two configurations of the GCM with the winds nudged to the ERA5 reanalysis we find that:

- The new configuration using DYNAMICO was just as effective as the reference configuration of the master GCM for atmospheric CO₂ transport. Both configurations accurately capture the seasonal variations of CO₂ at the surface stations. They also perform almost identically in terms of spatial gradient of annual growth of CO₂ compared to measurements, meaning that they are both efficient at long-term transport of tracers at the global scale.
- Both configurations perform similarly regarding the modeling of the synoptic variability. However, the relatively coarse horizontal resolution of the main simulation limits the upper bound of the quality of this metric.
- The configurations show some difference in vertical transport around the stratosphere but are very similar at lower altitudes. This difference is dependent on the season and latitude, and is largely influenced by the modelization of gravitational waves. The ICO configuration needs to be further tuned in this regard. But at around 15 km both configurations differ from aircraft and AirCore measurements.
- The two configurations did not show any significant difference in CO₂ atmospheric transport at high latitudes near the poles despite the big difference in spatial resolution between them in these areas. This can be explained in large parts by the effect of the polar filter in the regular configuration. The absence of a polar singularity in the icosahedral grid therefore did not lead to any major change in terms of atmospheric transport. However this does not mean that the grid does not provide other advantages.
- The ICO configuration has the major advantage of requiring a smaller number of cells for an equivalent spatial horizontal resolution at the Equator. This already provides a significant speed-up compared to the regular configuration that only gets bigger at higher resolution. Moreover, the DYNAMICO configuration can scale much better than the regular configuration. On the regular grid, the maximum number of MPI processes is limited by the resolution since each one needs at least two latitude bands, this is not the case with DYNAMICO. Since more CPUs can be used with this new configuration the time-to-solution can be drastically reduced and scaling at higher resolutions is more efficient.

Simulating the variations of carbon dioxide in the global atmosphere on the hexagonal grid of DYNAMICO coupled with the LMDZ6 model

Zoé Lloret¹, Frédéric Chevallier¹, Anne Cozic¹, Marine Remaud^{1*}, Yann Meurdesoif¹

¹Laboratoire des Sciences du Climat et de l'Environnement, LSCE/IPSL, CEA-CNRS-UVSQ, Université Paris-Saclay, Gif-sur-Yvette, France

* currently at Faculty of Science, A-LIFE, Vrije Universiteit Amsterdam, 1081 HV Amsterdam, the Netherlands

Correspondence to: Zoé Lloret (zoe.lloret@lsce.ipsl.fr)

Abstract. Efforts to monitor the emissions and absorptions of atmospheric carbon dioxide (CO₂) over the globe and to understand their varying regional patterns with greater accuracy have intensified in recent years. This study evaluates the performance of a new model coupling, ICO, built around the Laboratoire de Météorologie Dynamique atmospheric general circulation model (LMDZ) for simulating CO₂ transport. ICO utilizes the new icosahedral hydrostatic dynamical core called DYNAMICO running on an unstructured grid, which enables potential improvements in spatial resolution at the Equator while removing artificial distortions and numerical filters at the poles. Comparisons with a reference configuration using a structured latitude-longitude grid reveal that ICO well captures seasonal variations in CO₂ concentrations at surface stations. While not significantly improving the simulation of complex seasonal patterns, ICO maintains comparable accuracy. Both configurations exhibit similar vertical CO₂ concentration profiles and display a consistent bias in the lower stratosphere relative to observational data. ICO demonstrates advantages in computational efficiency and storage, thanks to its reduced cell count per level and a homogeneous grid structure. It holds promise for future developments, including with the LMDZ offline model and associated inversion system, which contribute to the Copernicus Atmosphere Monitoring Service. Overall, the ICO configuration showcases the efficacy of utilizing an unstructured grid for the physics, and the capability of DYNAMICO in accurately simulating CO₂ transport. This study emphasizes the importance of advanced modeling approaches and innovative grids in enhancing our understanding of the global carbon cycle and refining climate models.

1 Introduction

The key role of carbon dioxide (CO₂) in climate change has motivated increasing efforts in recent decades to monitor its variations in the global atmosphere. Sources and sinks of this trace gas are found primarily on the Earth's surface. They induce the highest CO₂ gradients in the boundary layer, for example around anthropogenic emission hotspots, while their direct influences gradually mix over time at all altitudes to contribute to the overall CO₂ background. The distribution of CO₂ in the atmosphere therefore spans a wide range of spatial and temporal scales, mainly combining influences from surface sources, surface sinks and meteorology. This complexity is sampled by growing high-quality observation networks on the ground, in the atmosphere (aircraft, balloons, drones) and in space (e.g., Ciais et al., 2014; Crisp et al., 2018). It is also simulated, more or less well, by Atmospheric General Circulation Models (GCMs) and dedicated tracer transport models (e.g., Remaud et al., 2018; Basu et al., 2018; Agustí-Panareda et al., 2022). Many uncertainties in the model input data (boundary conditions, meteorology) and the model equations (advection schemes, subgrid parameterizations) still limit these simulations. However, there is a strong incentive towards higher spatial resolutions in order to benefit from an increased realism for orography, coastlines, and known emission or absorption hot-spots, and to reduce any artificial smoothing of the 3D fields (Agustí-Panareda et al., 2019). However,

this wish is tempered by the need to carry out long simulations of this long-lived tracer, typically several years, that may be massively repeated in the case of inverse modeling. Increasing the resolution without affecting the time-to-solution leads to revisiting the numerical efficiency of models in order to gain computing time margins. Porting codes on Graphical Processing Units (GPUs) may largely contribute to this effort (in particular for the simpler codes of the offline models, as shown by Chevallier et al., 2023), but not enough to close the gap between model resolution and, for exemple, the kilometer-resolving resolution of the current space-borne observations. In particular, models running on a regular longitude-latitude grid face scaling limitations due to advection at the poles requiring significant data communication to solve the problem of resolution clustering. This data exchange can create a computing bottleneck on supercomputers using large amounts of processors (Staniforth & Thuburn, 2012). Moreover, the efficiency of porting existing GCMs to GPUs depends on the structure of their code, which may have to be redesigned. In this paper, we are addressing these issues for the simulation of CO₂ transport using an unstructured quasi-uniform grid made of non-quadrilateral grid cells. Such a solution is gaining popularity in recent years, either for use in Earth system models, transporting tracers or directly for atmospheric inversion (Niwa et al., 2017; Giorgetta et al., 2018; Sakaguchi et al., 2020; Zheng et al., 2021).

We build on the dynamical core DYNAMICO (Dubos et al., 2015), which has been integrated for the first time here into LMDZ, LMDZ being the GCM of the Laboratoire de Météorologie Dynamique (Hourdin et al., 2020). LMDZ has been used as the atmospheric component of the Institut Pierre-Simon-Laplace (IPSL) Earth system model (Sepulchre et al., 2020) and for the Climate Model Intercomparison Project (CMIP) with its traditional regular longitude-latitude grid. For comparisons with real observations, e.g., for inverse modeling, it is nudged to horizontal wind fields obtained from a numerical weather forecast reanalysis.

This paper evaluates the ability of this new configuration of the LMDZ GCM using the DYNAMICO dynamical core to transport a long-lived tracer like CO₂. DYNAMICO has been extensively compared to other dynamical cores (Ullrich et al., 2017) in a stand-alone fashion, but has not been evaluated yet when integrated into a GCM for use with tracer transport. Coupling DYNAMICO to the LMDZ GCM also represents the first step towards the use of DYNAMICO for inverse modeling of CO₂. We compare it to the previous reference configuration running on a regular longitude-latitude grid and to various observations of CO₂ mole fractions over a 40-year period. We also compare their computational performance at a higher resolution to judge the future scalability of this new configuration.

This study is but one stepping stone in the overall strategy of the development of the LMDZ GCM to increase spatial resolution while leveraging the advancements in high-performance computing, emphasizing manycore systems and hardware accelerators. DYNAMICO was created to meet these needs (Dubos et al., 2015) and implementing its coupling with the rest of the GCM as well as validating its performance against a reference configuration for tracer transport are pivotal steps in this development. Analyzing the computational performance of this new configuration and its scalability will also help inform decisions on future porting of these models and coupled configurations to GPUs.

Section 2 describes the two configurations of our GCM, the developments done to create our new configuration, the experiments we ran to compare them, and the method for our study. Section 3 presents a performance comparison of our configurations and the results of the direct comparison between our models and the observations. Section 4 concludes the study.

2 Presentation of the model and experiments

2.1 Configuration description

In this study, we compare two configurations of the LMDZ GCM that each couple different individual models. The full configurations are schematized in Figure 1. The individual models are presented in section 2.1.1 and the newly developed configuration is presented in detail in section 2.1.2.

Our reference general circulation model configuration (Fig. 1) consists of the coupling between the LMDZ model of Hourdin et al. (2013, 2020) itself, an aerosol and reactive chemistry model called INteractions between Chemistry and Aerosols (INCA, Hauglustaine, et al., 2004) and the Organizing Carbon and Hydrology in Dynamic Ecosystems land surface model (ORCHIDEE, Krinner et al., 2005). ORCHIDEE simulates the water and energy exchanges between the soil and the atmosphere, but yearly land cover maps were used here instead of simulating vegetation dynamics. In the following, we will refer to this LMDZ-INCA-ORCHIDEE coupled model on the regular latitude-longitude grid as REG for simplicity.

We have developed a novel configuration that couples the ORCHIDEE land surface model, the INCA chemistry model, and the physics module of the LMDZ model. The previous dynamical core in LMDZ has been replaced by a new one, known as DYNAMICO (Dubos et al., 2015), which operates on a quasi-uniform icosahedral C-grid for its horizontal mesh (Fig. 1). In the following, we will refer to this DYNAMICO-LMDZ-INCA-ORCHIDEE coupled model as ICO for simplicity.

Both configurations use XIOS, a tool that allows asynchronous and parallel input and output of files (<https://forge.ipsl.jussieu.fr/ioserver/wiki>).

Each of these pre-existing models can be operated either independently using precomputed files as input, or using information from other models to which they are coupled. We will describe them and their role in our configurations in the following subsection.

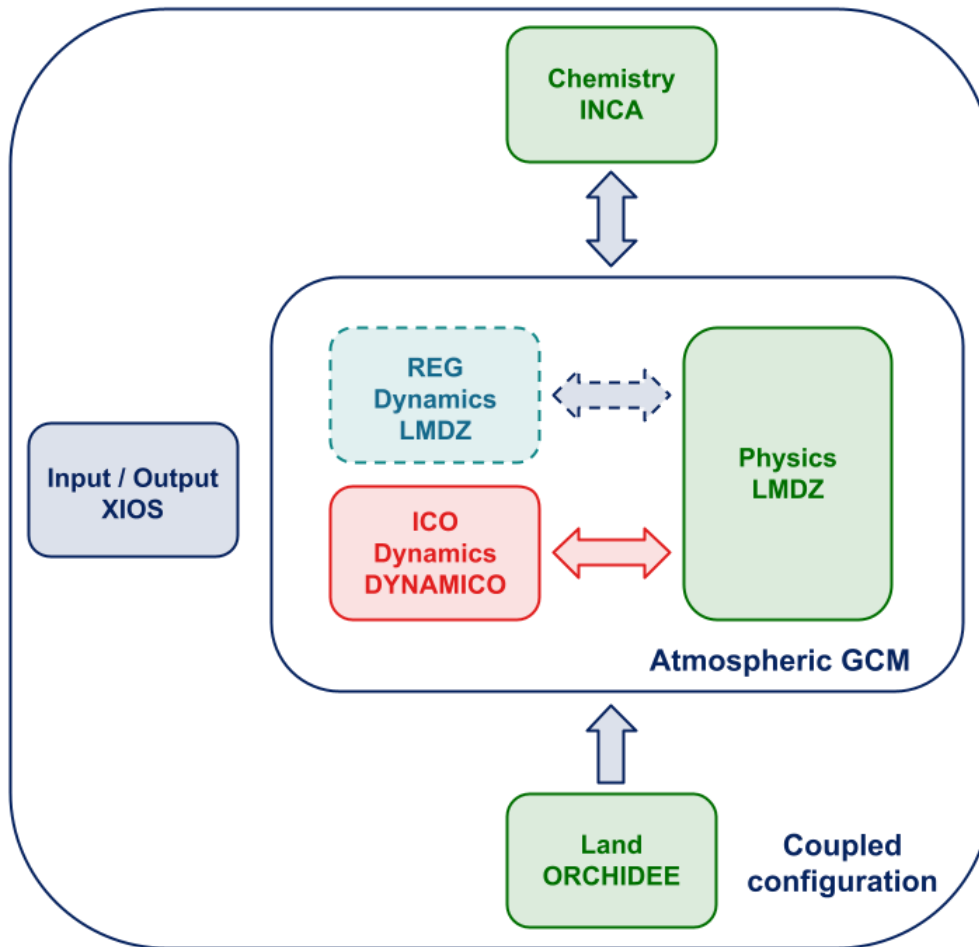


Figure 1: The structure of the two coupled configurations REG and ICO. Both use the same models for the physics, land surface and chemistry but they each use a different dynamical core. For the ICO configuration, the new coupler between DYNAMICO and LMDZ as well as XIOS ensure a seamless transition with the new icosahedral grid for all the models.

2.1.1 General description of the two models

The two configurations use the same individual physics, land and chemistry models but use different dynamical cores.

In our configurations, tracers, such as CO_2 , are outputted by INCA, and their transport is calculated and synced with the LMDZ GCM physics time step. Chemical processes are also calculated at this same frequency by computing differential equations to update the atmospheric mole fraction fields of each cell. Using tracers from INCA instead of only having them in LMDZ allows interaction between chemical reactions and the tracer transport process, which is crucial for some tracers such as CH_4 , although it has no impact on CO_2 . In our study, these chemical processes are not applied to the CO_2 tracer.

The latest version of LMDZ physics is described in Hourdin et al. (2020). Most notably for tracer transport, dry and cloudy shallow convection is separated from deep convection. Shallow convection is unified and combines the Mellor and Yamada (1974) diffusive approach for small-scale turbulence with a thermal plume model (Rio & Hourdin, 2008) for the boundary layer. Deep convection uses a modified version of the mass-flux formulation of Emanuel (1991) (Grandpeix et al., 2004, Rochetin et al., 2014). Longwave radiation is modeled using the Rapid Radiation Transfer Model (RRTM; Mlawer et al., 1997), and shortwave radiation uses a 6-band code derived from Fouquart and Bonnel (1980).

The dynamical core of LMDZ discretizes the sphere of the primitive equations of meteorology and of transport equations (Hourdin et al., 2006, 2013). Water and other tracers are advected with a scheme from Van Leer (1997), and angular momentum is conserved numerically. This full configuration was previously evaluated for CO₂ transport by Remaud et al. (2018).

The dynamical core of LMDZ is parallelized in latitude using distributed memory with the Message Passing Interface standard (MPI) and in the vertical with shared memory using the Open Multi-Processing interface (OpenMP). A longitudinal filter near the poles avoids the use of very small time steps, but limits the efficiency of any parallelism along the longitudes.

The parallelization of the physical parameterizations within LMDZ follows a different approach. It utilizes a combination of MPI and OpenMP processes with shared memory by splitting a single vector that runs through the entire horizontal grid into independent domains. This is possible due to the fundamental 1-D nature of the LMDZ physical parameterizations that only compute vertical transfers. The performance of the model is optimized by using domain decomposition parallelism on the horizontal layer with MPI and shared memory parallelism with OpenMP.

DYNAMICO is a dynamical core that solves the hydrostatic and shallow-atmosphere non-hydrostatic Euler equations (Ullrich et al., 2017). In this study, the hydrostatic mode was used. The mesh is based on a tessellation of the sphere into triangles, which when joined, creates the primal hexagonal-pentagonal mesh. A quasi-uniform grid avoids any singularity at the poles, thereby improving the load balancing on parallel computers. By construction, this grid has a coarser spatial resolution than a regular longitude-latitude grid in the high latitudes, even when accounting for the longitudinal filter (Herrington et al., 2022). The cells have similar areas across the globe, from the equator to the poles, unlike in the regular longitude-latitude grid where cell size gets systematically smaller when approaching the poles. Figure 2 provides an example of a visual representation of the icosahedral C-grid.

The kinematics handle the transport of mass, potential temperatures, and tracers using the mass fluxes computed by the dynamics. The vertical transport uses a slope-limited Van Leer's scheme (Van Leer, 1977) and does not differ from the REG configuration. The fully discrete finite volume horizontal advection scheme is described in Dubey et al. (2015). It uses a flux-corrected transport approach to stay positive-definite rather than slope limiters.

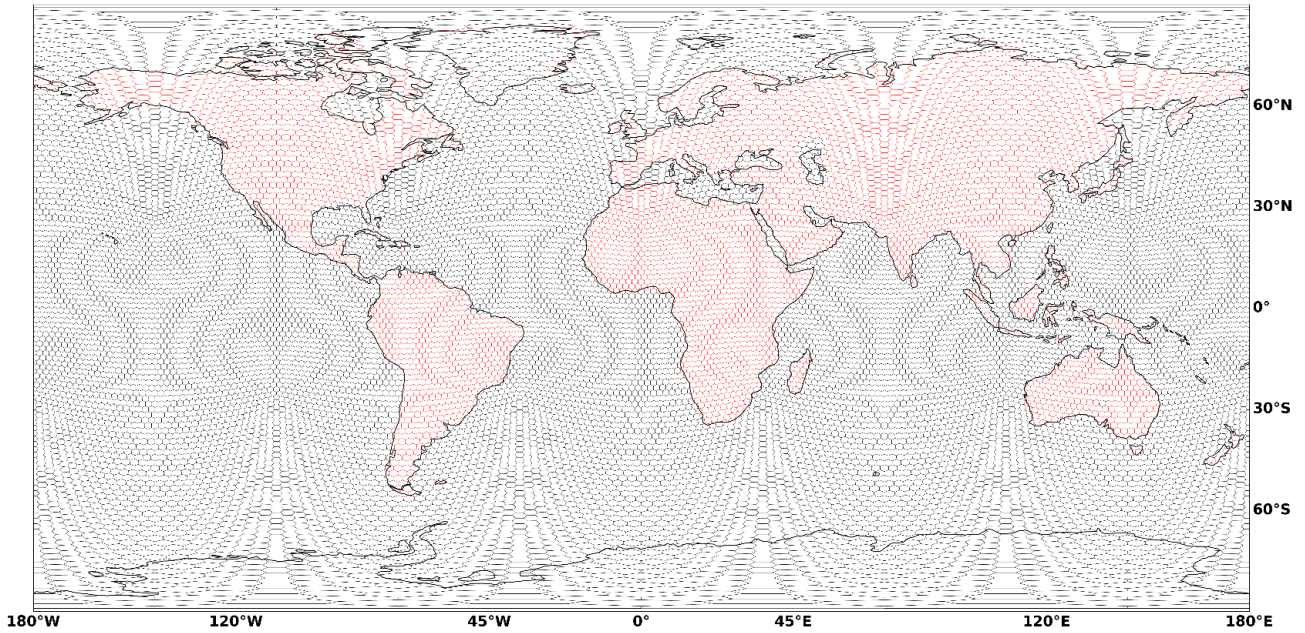


Figure 2: Icosahedral grid of the ICO configuration with with a horizontal grid of 16002 cells (see section 2.2 for a description of the resolution).

2.1.2 ICO - New icosahedral grid configuration

DYNAMICO, the dynamical core itself of this new configuration, was already developed and presented in Dubos et al. (2015) but not coupled with the physics of LMDZ at the time, nor used as part of a larger coupled configuration with a chemistry or land model. To that goal, a specific module handling the coupling of the dynamical core to the physics of LMDZ was needed.

This coupler specifically interfaces the dynamics with the already existing physics of LMDZ, with an asynchronous time step. The time step of the dynamics is inextricably linked to the resolution of the model. For the main resolution used in this study and described in section 2.2, the dynamics time step is 7.5 minutes, half that of the physics and identical to the time step in the REG configuration. This time step has to be reduced accordingly when increasing the resolution in Section 3.2 to satisfy convergence conditions. Another important part of this coupler is interpolating the interfaced variables since some variables in the dynamics are computed at the edges or vertices, whereas the physics uses centered variables.

The global mesh used in the ICO configuration is partitioned into 10 quadrilateral clusters of similar size to map the sphere. Each quadrilateral is paved with the same number of hexagonal cells, depending on the chosen resolution. It can be subdivided along two directions, i and j , thus generating sub-tiles in the form of parallelograms composed of $iim \times jjm$ hexagons (Fig. 3).

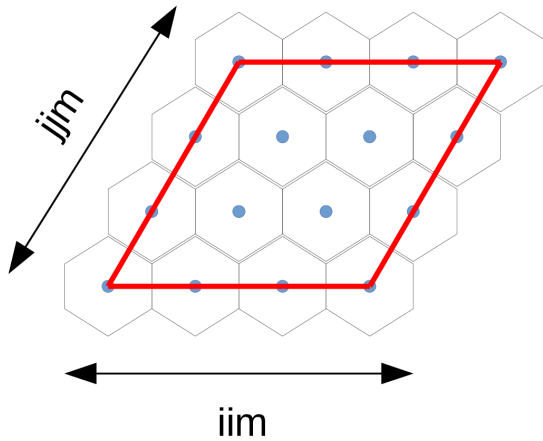


Figure 3: Subdivision of the primary mesh quadrilateral along the two directions i and j .

MPI parallelism is achieved by distributing the sub-tiles thus created to each MPI process, the optimum performance being achieved for sub-tiles of identical size and with only one sub-tile per MPI process. When solving the various numerical schemes of the dynamic core, the data required at the domain boundaries is transferred by asynchronous MPI calls from one sub-tile to another. OpenMP parallelism operates through shared memory, distributing computational iterations on vertical levels over threads created within an MPI process, in a similar fashion to the REG configuration.

Tracer storage was handled very differently between DYNAMICO and LMDZ and had to be uniformized. Initially, DYNAMICO tracers were only identified by a unique number and handling them necessitated always knowing which tracer corresponded to which index number in the tracer variable defined in a unique tracer definition file. LMDZ on the other hand was able to dynamically access tracers identified by chemical species or isotope name. Definition and initialization of tracers is now unified through a single file, and a parser was created so that tracers in DYNAMICO can be linked to these same chemical species names.

Another important feature that had to be developed was the ability to nudge the atmospheric GCM to variables such as winds, temperature or surface pressure. This feature was already available in the REG configuration but had to be re-created for the ICO configuration. A new guided mode was added to the dynamical core DYNAMICO, and special care had to be taken to keep the ability to nudge the variables either at a global scale or restricted to certain areas only of the icosahedral grid. Such restrictions are often parametrized in latitude-longitude coordinates and must be properly interpolated to the icosahedral grid while handling the conditions at the edge of the nudged area. This development is also essential for future use of the DYNAMICO - LMDZ coupling as a limited-area model.

2.2 Description of the simulations

For each configuration (REG and ICO), we have run a simulation from 1979 to 2020. The first year is used for spin-up and is not analyzed. In both configurations, the large-scale atmospheric circulation was nudged to the 6-hourly horizontal winds from the

ERA5 reanalysis (Hersbach et al., 2020) with a relaxation time of 3 hours. The nudging drives the large-scale atmospheric circulation of the model. Initial atmospheric CO₂ mole fractions values were set using the Copernicus Atmosphere Monitoring Service (CAMS) atmospheric inversion, version 20r2 (Chevallier et al., 2005; <https://atmosphere.copernicus.eu/>, access 31 May 2023). This same product prescribed the CO₂ surface fluxes every 3 hours. Prior fluxes used for this product were GCP-GridFED version 2021.2 for anthropogenic emissions (Jones et al., 2021), GFED 4.1 inventories for biomass burning, ocean fluxes from Chau et al. (2022) and climatological biosphere-atmosphere fluxes from an ORCHIDEE simulation, version 4.6.9.5.

These surface fluxes carry some imprint from the REG model with a regular grid, since the CAMS database used an older REG model version at coarser spatial resolution. Still, after Remaud et al. (2018) who tested a distinct set of surface fluxes for their model evaluation within a similar framework, we consider that this imprint hardly affects our conclusions.

The boundary conditions used for the two simulations were identical (aerosol, oxidants and ozone mole fraction, solar forcing, land use maps). However, for the simulation running on ICO, the boundary conditions were either interpolated or recreated onto the new grid ahead of time to fit the unstructured grid or interpolated during execution. The initial total mass of CO₂ in the atmosphere had a difference of only 0.01% between the two simulations because of these operations.

We had an hourly model output for all variables. This high frequency output was chosen in order to well assess the differences in synoptic variability of tracer transport between our two model configurations.

We used two different horizontal resolutions in this study, the lower resolution one was used in our main simulations to compare the difference in CO₂ tracer transport between our two configurations. We chose this resolution because it was an already established and studied resolution for our reference REG configuration, allowing us to study only the influence of the new dynamical core and grid. We also performed computational performance tests at a high resolution to test the scaling of our two configurations and make better informed decisions about their future development.

For our main simulation, we ran REG on a horizontal grid of 144 points in longitude \times 143 grid points in latitude, which corresponds to a resolution of 2.5° in longitude and 1.27° in latitude, equivalent to 278 km by 140 km at the equator. We use 79 vertical layers going up to 80 km in altitude, with around 25 layers dedicated to the first 2 km. The complete grid configuration is described in more detail in Hourdin et al. (2020).

We compared this configuration with ICO running on a horizontal grid of 16002 cells, and the same 79 vertical layers. This gives an horizontal resolution at the equator of around 2.5° in longitude and 1.25° in latitude, each cell has the area of an hexagon of side 110 km in order to have similar resolution at the equator to the equivalent longitude-latitude grid from REG. With this setup, ICO has 22% less cells than REG.

For the performance tests of the configurations running at a higher resolution, REG has an horizontal grid of 256 points in longitude \times 256 grid points in latitude, which corresponds to a resolution of 1.4° in longitude and 0.7° in latitude, equivalent to 157 km by 78 km at the equator. In this test, ICO was run on an horizontal grid of 64002 cells, with hexagons of side 55 km. Both configurations still have 79 vertical layers and at this higher resolution ICO only has only 2% less cells in total than REG but with a higher spatial horizontal resolution.

2.3 Observational data

To compare our simulated tracer mole fractions to observations, we sampled the mole fraction fields at the nearest cell center, model level and timestamp for each data point. We used the high-quality measurements of the CO₂ GLOBALVIEWplus v8.0_2022-08-27 ObsPack database (Schuldt et al., 2022, Miles et al., 2017, Miles et al., 2018, ICOS RI, et al., 2023, Lan et al., 2023). For AirCore, we used the dataset from NOAA Version 20230831 (Baier et al., 2021).

In this dataset, observations were calibrated according to the WMO CO₂ X2019 scale (Hall et al., 2021). Like for inverse modeling with LMDZ (Chevallier et al., 2010), only afternoon non flagged data from 12:00 to 16:00 local time were selected for continuous in-situ surface stations under 1000 m above sea level (a.s.l.), and only night time data from 00:00 to 4:00 local time were kept for in-situ stations above 1000 m a.s.l. This selection accounts for the usual failure of transport models to well represent the accumulation of tracers at low altitude during the night as well as the inability to model the phenomenon in mountain stations where air masses are advected during daytime through updrafts on the sun-exposed slopes (Geels et al., 2007). By selecting data differently between high-altitude and low-altitude surface stations, we ensure that the cases described here are avoided. All flask data, and all upper-air data (aircraft data and AirCore measurements) were kept.

We divided the observations into three groups: surface in situ and flask data, aircraft observations and observations from AirCore flights. We used the aircraft measurements and AirCore data to obtain vertical profiles of CO₂ mole fraction. For surface data, 106 stations have been selected from the Obspack dataset out of the original 222 stations. Surface stations with less than 5000 measurement points that passed the initial data selection described above over the entire duration of the study were excluded from the analysis. For aircrafts, we have selected 33 sites and campaigns out of a possible 51, only keeping those with more than 2000 measurement points. For the AirCore data, we kept all observations. The full list of sites and datasets used is presented in Table 1 and Table 2 as a supplement.

The uncertainty of the reference CO₂ mole fraction measurements used here is on the order of 0.1 ppm (see, e.g., Crotwell et al., 2020, for the systematic errors and Hazan et al., 2016, for the standard deviation). It is negligible compared to the model uncertainty due to transport error which is on the order of 1 ppm under 3000 m (Lauvaux et al., 2009) and is not further discussed in the following. Collection altitude determination error from AirCore measurements can be high and depends on the altitude, and is on the order of 250 m below 20 km and up to 1 km above that altitude (Wagenhäuser et al., 2021).

2.4 Evaluation methodology

2.4.1 Surface stations

For surface stations with continuous measurements, we used a curve-fitting method using a non-linear least squares method on both the model and observations CO₂ mole fractions time series to extract the annual mean, the seasonal cycles and the synoptic variations. A smoothed function consisting of a second-order polynomial and eight harmonics was used to fit the time series over the 1980-2020 period (equation 1). We excluded stations where the fit of the measurements had a coefficient of determination under 0.75, for a total of 85 stations. The polynomials were used to calculate the annual trend and growth rate, while the harmonics were used to get the seasonal cycle.

$$f(t) = p_1 \cdot t + p_2 \cdot t^2 + \sum_{k=3}^{10} p_k \cdot \sin(2\pi kt) \quad (1)$$

$$r(t) = x(t) - f(t) \quad (2)$$

Equation 1 corresponds to the fitted function of CO₂ mole fractions f with t being the time. The different parameters p_k are coefficients fitted after optimization and unique to each model and station. The parameters p_1 and p_2 correspond to the general trend of the growth rate of atmospheric CO₂, and the parameters of the sinusoidal function to the seasonal cycle. The residual, corresponding to the synoptic variations, r , are obtained from the difference between the raw measured or modeled CO₂ mole fraction x and the fitted smooth curve f (Equation 2).

To evaluate the two model configurations performance between each other and compared to observations we use metrics which we will describe in the following subsections.

2.4.2 Annual gradient between stations

We use the measurements from South Pole station (SPO), which is far from any major CO₂ source or sink, to validate the simulated background growth rate of CO₂ mole fraction. Then, we study the cross-site gradients by calculating the yearly growth rate at each site relative to SPO. To do so, we average the annual growth rate of the CO₂ mole fraction over the 1980-2020 period for each site and subtract the value at SPO. Comparing the observed and modeled values of this variable informs us on both the growth rate of the CO₂ mole fraction at each site, and on mole fraction gradients of our transport model which are key for use in an inverse system. To study the interannual variation of these growth rates, we calculate their standard deviation for both measurements and models. We normalize the average model's standard deviation by dividing it by the measurement standard deviation. This gives us information on how well the model captures the magnitude and direction of these variations.

We compute the yearly growth rate for each year of the 1980-2020 period using the smooth curved fit described above, before averaging it. To evaluate this variable, we then look at the mean bias and the root-mean-square error (RMSE) of the CO₂ mole fraction gradient for each station relative to SPO.

2.4.3 Seasonal cycle

We evaluate the capacity of our model to represent the CO₂ seasonal cycle by comparing the phase and amplitude of the harmonics of their smoothed fitted curve to the one of the measurements at each station. At each measurement site we calculated the Pearson correlation coefficient between measurements and model time series to evaluate the phase of the seasonal cycle. And we evaluated the amplitude of the seasonal cycle by looking at the ratio between peak-to-peak amplitudes of the harmonics. We normalized this variable by dividing the values of the model's seasonal cycle peak-to-peak amplitude at each station by the ones from the observations.

2.4.4 Synoptic variability

To evaluate our model ability to represent the phase of the CO₂ synoptic variability we again used the Pearson correlation coefficient between the residual from the smoothed fitted curve of the model and the measurements. The amplitude of the synoptic variations at each station were evaluated by the normalized standard deviation.

2.4.5 ERA5

To compare the simulated temperature with the ERA5 reanalysis, we divided the output into seasons and then into bins of 30° latitudes. For each bin, we averaged the data for each model level for each season. We then did an identical operation on the ERA5 reanalysis data before comparing the two.

2.4.6 Aircraft measurements

The aircraft measurements have been binned into 1 km altitude bins, and then averaged for each hour and over each bin for each site or campaign. Then the data was averaged over all sites and campaigns. This process was done for each season and for the whole year.

2.4.7 AirCore measurements

For measurements from AirCore, we binned and averaged the data into 50 altitude bins, from the ground to the maximum altitude of the data (27 km) to get an average vertical profile of CO₂ mole fractions.

3 Results and discussions

3.1 Mass conservation

Conservation of mass is closely examined for the simulation of long-lived tracers as it directly supports the simulation of the tracer's global growth rate. In inverse systems, it makes it possible to infer surface fluxes far from observations, far in space as well as in time. In practice, numerical approximations may make the model lose or gain tracer mass (Houweling et al., 2010).

In this section, we evaluate mass conservation in our models by calculating the total mass of atmospheric CO₂ at the beginning and at the end of the simulations.

To do that, we calculate Δm , the observed increase in the total amount of atmospheric CO₂ over a certain period of time. In Equation (3), superscript e corresponds to the end time step of a given period studied, and i is the initial time step. We multiply the CO₂ mass fraction w with the dry air mass m_{air} in each cell and sum it over the whole grid (N cells).

We then separately calculate the total emitted mass of CO₂ over a period of time, $m_{(\text{CO}_2)}^{\text{emi}}$ in Equation 4, by multiplying the surface fluxes q with the area of each cell A^n and summing it over time.

The difference between these two values, m_{loss} in equation 5, is the total mass of atmospheric CO₂ lost or gained by our model over a certain period of time.

For REG, the difference is equal to -0.13 % of the CO₂ mass emitted over the 1979 - 2020 period. For ICO it is -0.28 % for this same period.

Therefore, while our models do not exactly conserve mass, they lose only around 0.014 GtC integrated over 10 years for REG, and 0.027 GtC for ICO.

$$\Delta m_{(\text{CO}_2)} = m_{(\text{CO}_2)}^e - m_{(\text{CO}_2)}^i = \sum_{n=1}^N m_{\text{air}}^{n,e} \times w_{(\text{CO}_2)}^{n,e} - \sum_{n=1}^N m_{\text{air}}^{n,i} \times w_{(\text{CO}_2)}^{n,i} \quad (3)$$

$$m_{(\text{CO}_2)}^{\text{emi}} = \sum_{t=1, n=1}^{T, N} q_{(\text{CO}_2)}^{t, n} \times A^n \quad (4)$$

$$m_{\text{loss}} = m_{(\text{CO}_2)}^{\text{emi}} - \Delta m_{(\text{CO}_2)} \quad (5)$$

The total amount of CO₂ in each model also depends on the prescribed surface CO₂ fluxes described in section 2.2 ($m_{(\text{CO}_2)}^{\text{emi}}$) which are interpolated on the 2 different grids and therefore, not strictly identical either for each configuration. However, the average difference in yearly emitted CO₂ between the two model configurations is 0.0006 % only.

We verified that the routines in the LMDZ physics and in the dynamical cores of both configurations perfectly conserve mass. Therefore, the small mass difference comes from discrepancies between the time-integrated values of emissions and mole fractions of our tracer, but we did not investigate it further given its negligible impact in our study.

3.2 Computational efficiency

3.2.1 Computational setup

Simulations were run on the Skylake partition of Joliot Curie, a BullSequana X1000 supercomputer operated since 2017 by Très Grand Centre de Calcul (TGCC, Bruyère-le-Châtel, France). This partition is composed of 1656 nodes, each of which has an Intel Skylake 8168 dual-processor. We used the Intel Fortran compiler version 20.0.0.

For our main simulations (called “Production run” in Table 1), REG used 47 MPI processes and 8 OpenMP threads for a total of 384 Central Processing Unit (CPU) cores, while ICO employed 80 MPI processes and 4 OpenMP threads for a total of 336 CPU cores (Table 1). This choice was made as a compromise between fast time-to-solution for the simulations, and small number of nodes for lower queue time. This “Production run” is made with an hourly output for 26 physical and tracer variables to have a precise understanding of the CO₂ tracer transport dynamics. This output frequency significantly increases the execution time and would generally be lower for most routine use of this configuration.

To better compare the configurations scaled up and in their ideal state, speed tests were run with different numbers of CPUs and with additional CPUs used for XIOS servers. XIOS is a tool used for reading the input files in parallel and we chose 8 servers to ensure that this operation does not become a computational bottleneck for our models. Only monitoring files tracking the progress of the simulations were output, no physical variables were saved in order to avoid comparing the time it takes to write the files on disk. To avoid variability due to individual node performance, the tests were performed multiple times over several days, and outlier months caused by node performance issues were removed. In addition, we used timers in the code to evaluate what percentage of the time is spent in the routines of the physics versus the dynamics.

The first experiment (called REG/ICO-Speed test in Table 1) consisted in using identical numbers of CPUs for both configurations, running at the same resolution as our main simulations: 71 MPI processes and 8 OpenMP threads. Another experiment (ICO-Optimal scale) was run using 160 MPI processes and 4 OpenMP threads, more suited for the parallelization scheme of the ICO configuration. To test the scaling potential of the new ICO configuration at a higher horizontal resolution (HR) compared to REG, similar speed tests were run for each configuration at high resolution at two different scales. The first test was run with a small amount of CPU cores (called Low-scale in Table 1) using the same optimal number of MPI processes and OpenMP threads as the lower resolution one. A second test was run with more CPUs (called High-scale in Table 1), optimized for this higher resolution. For REG, this meant using 128 MPI processes and 8 OpenMP threads. For ICO, using 640 MPI processes and 4 OpenMP threads.

Simulation	Total CPU cores	MPI processes	OpenMP threads	Average monthly CPU (hours)	Monthly time to solution (seconds)	Time spent in dynamics (%)	Time spent in physics (%)
REG-Production	384	47	8	277	2594		
ICO-Production	336	80	4	209	2238		

REG-Speed test	576	71	8	132	823	20	77
ICO-Speed test	576	71	8	106	662		
ICO-Optimal scale	640	160	4	63	357	15	75
REG-HR-Low-scale	576	71	8	132	829	42	54
REG-HR-High-scale	1024	128	8	237	832	33	60
ICO-HR-Low-scale	640	160	4	125	702	30	64
ICO-HR-High-scale	2560	640	4	230	323	35	55

Table 1: Computational setup and results of the simulations. The total number of CPU cores used in each simulation may be higher than the product of MPI processes by OpenMP threads since entire nodes were reserved for better memory efficiency even though not all of their CPUs were needed.

3.2.2 Computational gains

For our production runs, on average, over the whole simulation, REG achieved a wall-clock-time of 2594 seconds and consumed 277 CPU hours per month simulated, while ICO executed in 2238 seconds and consumed 209 CPU hours per month simulated (REG-Production and ICO-Production in Table 1): in this case, ICO consumes 25% less CPU hours compared to REG, with 22% less cells. Most of the gains for this setup with heavy output levels could therefore be attributed to the reduced grid size.

The average monthly time to completion for the first speed test for REG was 823 seconds (132 CPU hours, REG-Speed test in Table 1), and for ICO 662 seconds (106 total CPU hours, ICO-Speed test in Table 1). This shows that for identical computational setups at this reference resolution, ICO is on average 20% faster than REG. This is again only of the order of the reduced grid size, showing that in this setup, other differences such as the absence of a longitudinal filter did not significantly improve the computational speed.

The ICO gain was increased by optimizing the distribution of MPI processes and OpenMP threads to better fit the parallelization scheme of ICO (ICO-Optimal scale in Table 1). In this optimized setup, ICO consumes 52% less CPU hours than the REG configuration, a gain in computing resources much greater than the one only due to the reduced grid size. This highlights the importance in optimizing the computing resources to each configuration for better performance.

At the high resolution (REG-HR-Low-scale, REG-HR-High-scale, ICO-HR-Low-scale and ICO-HR-High-scale in Table 1), ICO only has 2% less total cells than REG but is still faster. More importantly, it scales much better than the REG configuration, with the monthly time-to-solution going as low as 323 seconds when using 2560 CPUs. This is not the case for the REG configuration which plateaus at around 830 seconds per month irrespective of the increase in number of used CPUs. Further scaling is impossible for the REG configuration since it needs at least two latitude bands per MPI process, which means that 128 is the maximum number of MPI processes possible for this resolution.

This shows that while performance gains were modest in our main simulation, and only of the same order as the reduction in grid size, the performance gain increases when resolution increases or computing resources scale up. This gives a strong incentive to use an icosahedral grid compared to a regular latitude-longitude grid for future high-resolution studies using the LMDZ GCM.

3.3 Vertical temperature profiles

To get a first idea of the differences between REG and ICO simulations, we consider atmospheric temperature and compare it with ERA5 values. Note that our models are nudged toward ERA5 horizontal winds (Section 2.2), but do not use the ERA5 temperature fields. Figure 4 shows the vertical profiles of the average temperature over the year 2000 for different zonal cuts in 60° latitude increments. We can already see that REG and ICO differ on several aspects for different altitudes. The tropopause height, as identified by the change in the vertical temperature gradient, is the same in both configurations, but its temperature varies between 2.5 K to 5 K for each configuration outside of the tropics. At the stratopause, a difference of up to 10 K for the yearly temperature average in high latitudes is observed between the simulations from REG and ICO (Not shown on the figure).




Figure 4: Vertical profile of zonal temperatures averaged over the year 2000 for the two model configurations and the ERA5 reanalysis, with REG in blue, ICO in red and ERA5 in yellow.

Looking at the temporal change of the temperature rather than yearly averages reveals a different pattern. We can see on Fig. 5 that the large temperature difference at the stratopause between our configurations is only present during winters for high latitudes. During summers, both configurations have much more similar temperatures in these latitudes, and all year around in the tropics. This is explained by the fact that during summers, the polar stratopause is mainly driven by ozone, whereas in winter it is driven by gravity-waves (Hitchman et al., 1989). The difference in parametrization and tuning of gravity waves in DYNAMICO used in our new configuration ICO compared to the previously used and much-tested REG version likely explains the observed differences in temperature of the stratopause. This large difference in temperature in the stratosphere also affects temperature lower in the troposphere, as has been shown for the stratospheric dynamics of the LMDZ GCM in Lott et al. (2005). Future versions of the ICO configuration will contain a better parametrization of gravity waves as well as the introduction of a so-called “sponge layer” (Shepherd et al., 1996) to nudge high atmospheric winds towards zonal averages, which was already present in the REG configuration but not in ICO yet.

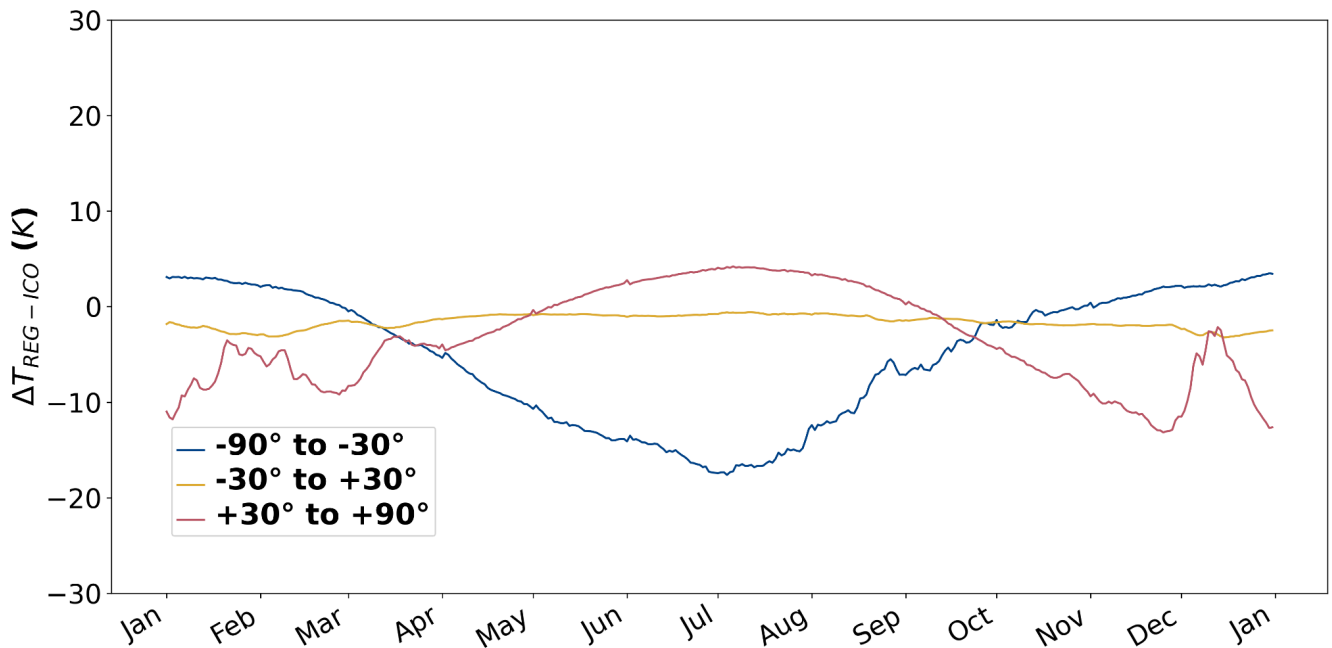


Figure 5: Time series of the average difference in zonal temperature at the stratopause (53 km) between the two model configurations REG and ICO for the year 2000, divided in 3 latitude zones of 60°.

We now turn to CO₂ mole fractions to see how the different models affect tracer transport.

3.4 Seasonal analysis

3.4.1 Annual gradient

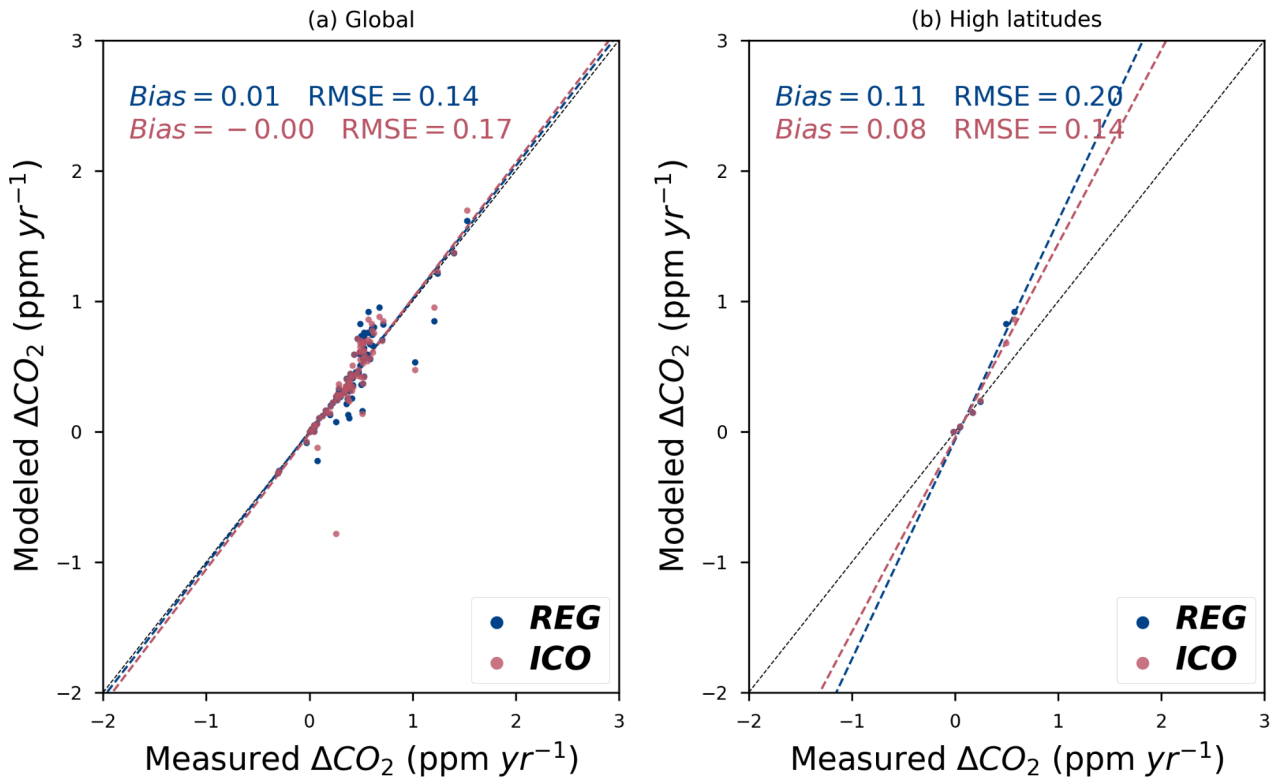


Figure 6: Annual gradients of CO₂ mole fraction compared to SPO averaged over the 1980-2020 period for every station (a) or only stations at high latitudes (> 70°N/S) (b). Blue circles are the model outputs for the REG configuration, and red circles for the ICO configuration. The dotted lines correspond to the linear fitted lines of the corresponding colored configurations, and the black dotted lines correspond to the 1:1 relation.

Figure 6 shows the annual gradients of surface stations compared to SPO averaged over the 1980-2020 period, and for the two model configurations the differences between the modeled and observed values of this gradient. We find an average yearly growth rate of CO₂ mole fraction at SPO of 1.79 ppm per year from observations, and of 1.74 ppm per year for both the REG and ICO configurations. This difference of 0.05 ppm between our models and observations shows that the background growth rate of CO₂ mole fraction is well modeled and within the small uncertainty range of the observations.

When looking at all surface stations (a), the ICO configuration exhibits a bias that is not significantly different from the REG configuration, and an almost identical spread as seen by the root mean square error (RMSE). This is confirmed by the results of a two sample t-test of the annual gradients of all stations for the two configurations ($t=0.25$ and $p=0.80$). Both configurations show a bias of less than 0.1 ppm per year compared to observations. The two model configurations therefore successfully model the annual gradients between surface stations over the globe.




Figure 7: Normalized standard deviation of the annual gradient for both configurations for each station. Blue circles are the model outputs for the REG configuration, and red circles for the ICO configuration. The dotted line corresponds to the ideal normalized standard deviation of 1. The stations are ordered on the abscissa by increasing latitude from -90° to $+90^\circ$.

The average normalized standard deviation of the interannual variation in the annual gradient for both configurations is 1.01, therefore both configurations show a good agreement in magnitude of these variations for the majority of stations (Fig. 7).

Since the biggest change regarding the grid and resolution takes place at the poles, we also checked the statistics and linear fit of these gradients restricted to stations at high latitudes (higher than 70°N and lower than 70°S). ICO performs just as well as REG for these stations in terms of both general bias and spread (Fig. 6 (b)). Even though the spatial resolution is much coarser for the ICO configuration at these latitudes, it has not significantly affected the simulation of long term trends of CO_2 mole fractions. We can also verify this by looking at the bias in the annual gradient of CO_2 mole fraction per station (Fig. 8) according to latitude. We see that there is no difference in between the two configurations related to latitude. This shows that forced resolution clustering at the poles of the regular latitude-longitude grid is not necessary for properly resolving tracer transport.




Figure 8: Bias per station (a) and average bias per 10° latitude band (b) of the annual gradients of CO₂ mole fraction compared to SPO averaged over the period 1980-2020, with blue circles for REG and red circles for ICO, the gray line is the difference between the two. The stations in (a) are ordered on the abscissa by increasing latitude from -90° to +90°.

3.4.2 Seasonal cycle

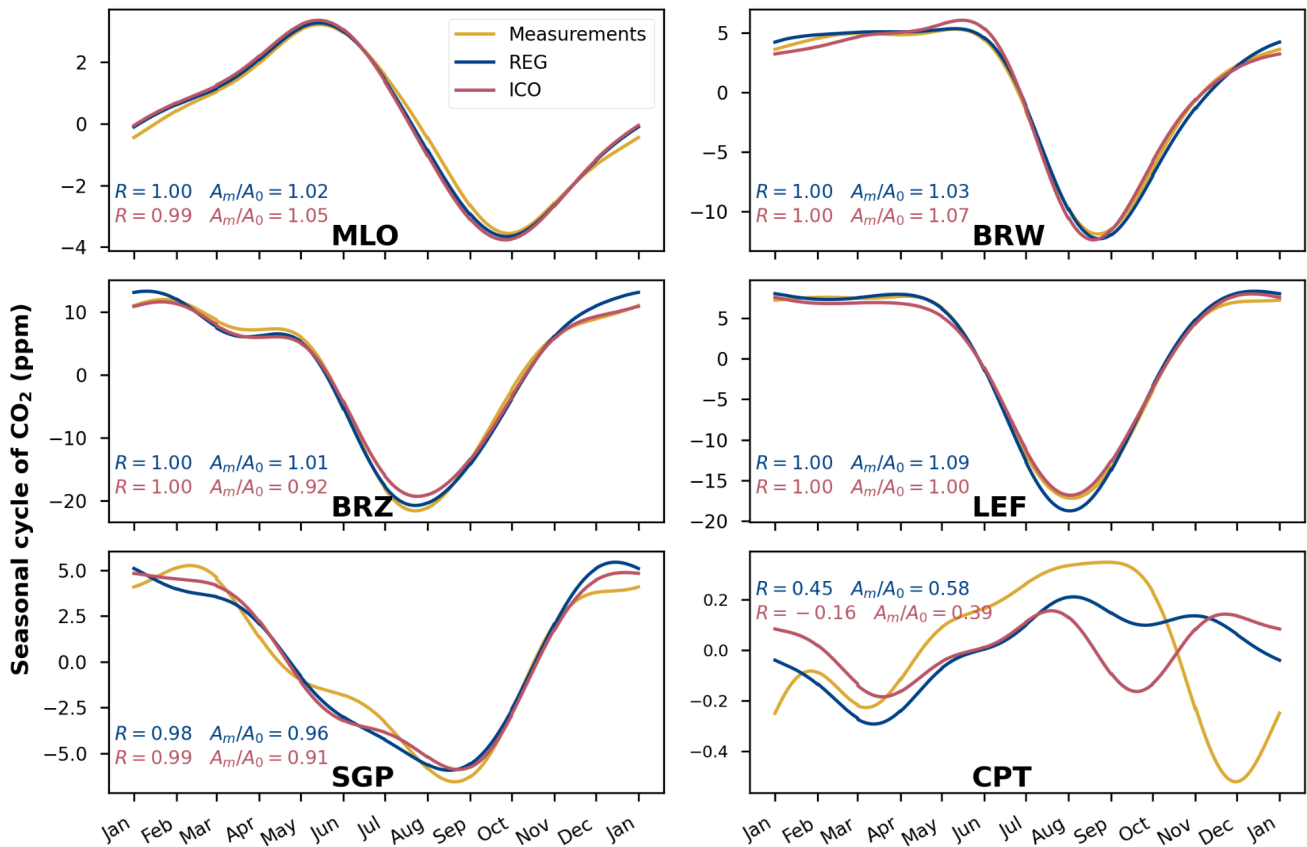


Figure 9: Seasonal cycle of the CO₂ mole fraction averaged over the period 1980–2020, in ppm, at a selection of surface stations for measurements, REG configuration and ICO configuration in yellow, blue and red respectively. The correlation coefficient and the peak-to-peak amplitude between the two model's output and measurements are displayed for each station. The selected stations were chosen to exemplify diverse behaviors: where both configurations successfully capture the seasonal cycle, neither configuration does so, or only one out of the two model configurations achieves it.

The seasonal cycles at most surface stations are well captured by both configurations, with regards to both phase and amplitude, as illustrated in Fig. 9. Some stations exhibiting more complex and higher frequency patterns of CO₂ mole fractions variation throughout seasons have a lower correlation coefficient. This pattern is observed for both configurations. However, almost all stations that are adequately modeled by the REG configuration with regards to seasonal cycles (correlation coefficient higher than 0.8) are equally well represented in the ICO configuration, as shown in Fig. 10 (a). Out of the 85 stations analyzed, only 3 stations did not exceed a correlation of 0.8 with either configuration. Only one station, CPT, performs significantly worse for the ICO configuration than for the REG configuration. The amplitude of the seasonal cycle is also well captured for almost all stations, as shown in Fig. 10 (b).




Figure 10: Pearson correlation coefficient (a) and normalized standard deviation (b) of the seasonal cycle for all surface stations studied averaged over the period 1980-2020, with blue circles for REG and red circles for ICO, the gray line is the difference between the two. The stations are ordered on the abscissa by increasing latitude from -90° to $+90^\circ$.

3.4.2 Synoptic variability




Figure 11: Pearson correlation (a) and normalized standard deviation (b) of the daily average residue between our modeled and measured CO_2 mole fractions at the surface stations described in section 2.5 for the period 1980-2020. The model output from the REG

and ICO configurations are in blue and red respectively. The stations are ordered on the abscissa by increasing latitude from -90° to $+90^\circ$.

To study the synoptic variability modeled by our two model configurations, we look at the correlation and the normalized standard deviation (NSD) of the daily averaged residue of our seasonal analysis for each surface station (Fig. 11). This gives us information on the accuracy of our simulation for higher frequency than the seasonal cycles. Both configurations have correlation coefficients over 0.63 for 25% of all stations and a mean value of 0.54 for REG and ICO. The ICO configuration has a lower mean NSD of 1.09 compared to the one of REG of 1.20. However, the results of a two sample t-test show that this difference is not significant ($t=1.38$ and $p=0.17$). And stations that offer a good correlation also tend to exhibit a better spread of the synoptic variability characterized by the NSD. These results are in line with what can be expected of a simulation at these resolutions as shown in Agustí-Panareda et al. (2019).

3.5 Vertical profiles of CO₂ mole fractions

3.5.1 Troposphere

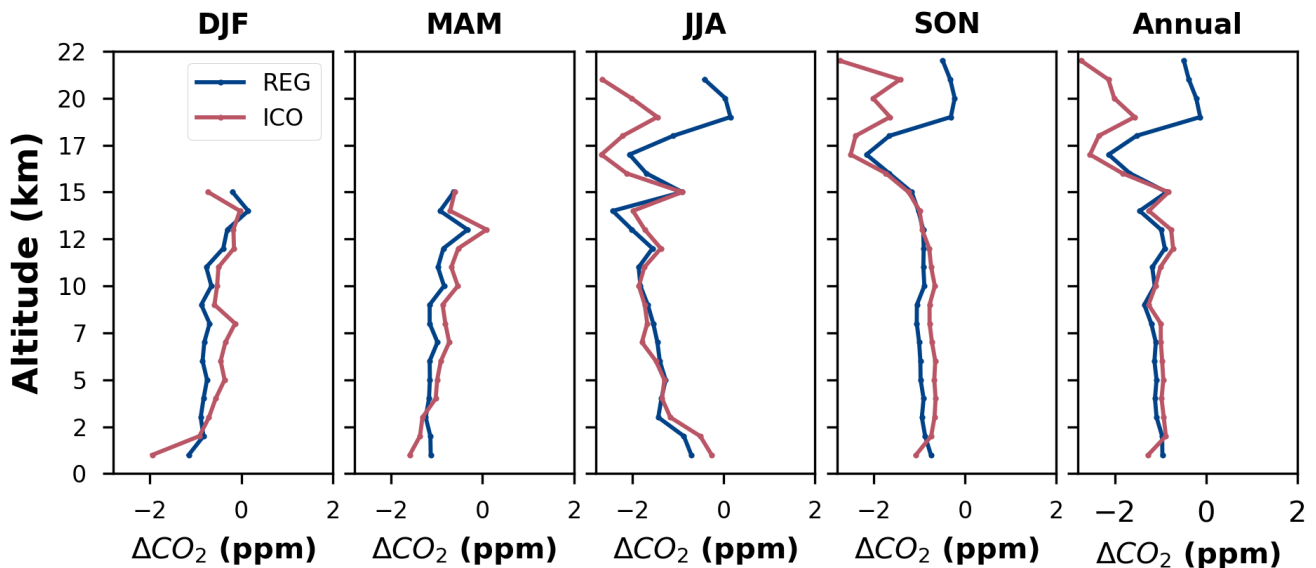


Figure 12: Seasonal and annual means of the difference in CO₂ vertical profile between the two model configurations of the model and aircraft measurements. The data has been binned into 1 km altitude bins for each season of the 1980-2020 period, then averaged per hour, and finally averaged across all aircraft sites and campaigns. The blue line represents the difference between REG and the measurements, while the red line represents the difference between ICO and the measurements.

In the troposphere, we studied CO₂ vertical profiles using various aircraft measurements described in section 2.4.2. Figure 12 shows the differences between the simulated and observed values for our two model configurations, REG and ICO. Only a small number of aircraft campaigns reach high altitudes above 15 km and not all seasons are covered. Both configurations show very similar vertical profiles up to 15 km altitude, before diverging above. Both configurations show a general negative bias compared to measurements. The vertical profiles are almost identical for all altitudes, but the extent of the differences between model's output and measurements differ at high altitudes. REG has much greater variations in CO₂ mole fractions while ICO has an increased negative bias at high altitudes. This is similar to the results in the next section 3.5.2.

3.5.2 Low stratosphere

We utilized data from AirCore flights to compare the CO₂ mole fractions of our model with observed data and obtain vertical profiles extending to the low stratosphere, in order to investigate the potential effects of the change in dynamics on vertical mixing within a column. However, since these measurements were only conducted in latitudes higher than 30°N and lower than 30°S, information about vertical tracer transport in the tropics was not obtained. As shown in Fig. 13, both model configurations exhibit an excess of CO₂ mole fractions around the 12.5 km range. Above that altitude, the vertical profiles differ in amplitude between the configurations but both show a similar decrease in CO₂ mole fraction. While the ICO configuration employs the same van Leer vertical transport scheme, differences in the vertical temperature profiles (see Fig. 4 and Fig. 5) discussed in section 3.4 could explain the disparity in the amplitude of the vertical profile at the stratosphere. The previously discussed conclusions are independently verified by the aircraft measurements that do not suffer from the Aircore altitude determination error and show similar differences in CO₂ mole fractions at 20 km (Fig. 12).

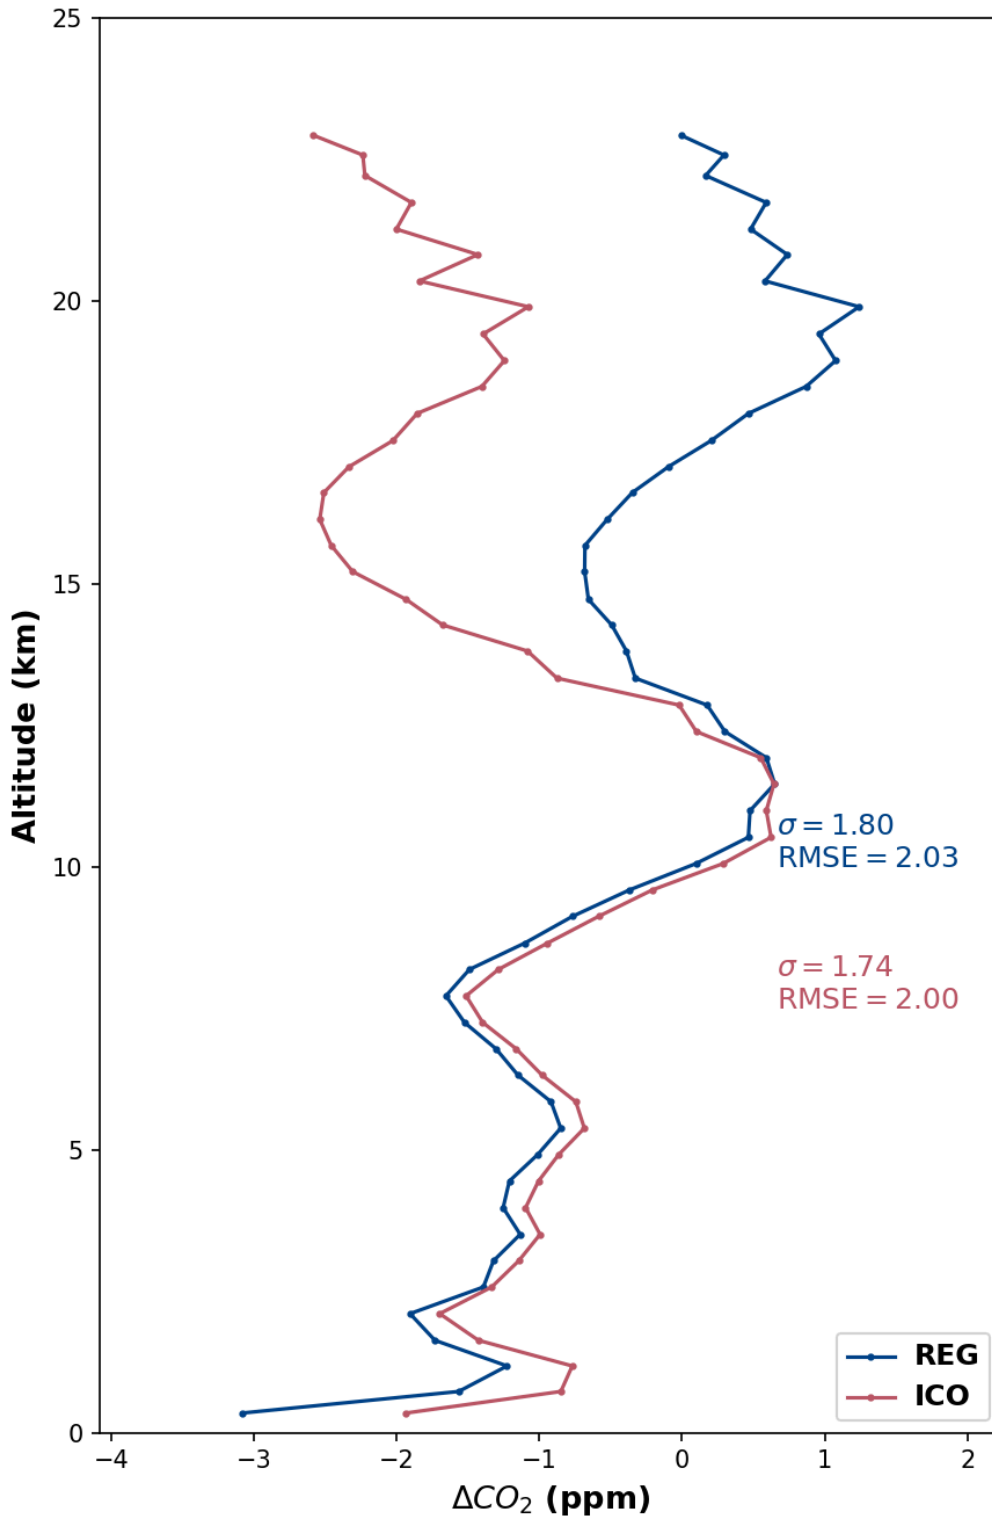


Figure 13: Difference in CO_2 mole fraction vertical profile between the two model configurations of the model and AirCore measurements. The blue line represents the difference between the REG model output and measurements, while the red line indicates the difference between the ICO model output and measurements. The fitted lines were generated by averaging the data over 50 altitude bins.

4 Conclusion

As demonstrated in the previous section's results, the configuration ICO based on the new dynamical core DYNAMICO, using an unstructured grid is just as effective as the reference configuration that used a structured latitude-longitude grid for modeling atmospheric CO₂ transport when the dynamics was nudged to horizontal winds of an ERA5 reanalysis. Both configurations accurately capture the seasonal variations in CO₂ mole fractions at most surface stations, and while the ICO configuration did not better capture more complex seasonal patterns, it did not worsen it either. The annual gradient is almost identical between both configurations, and they both perform well regardless of latitude. Regarding synoptic variability both configurations also perform almost identically. Nevertheless, both configurations provide an inadequate modeling of synoptic variability, as the local high-frequency emissions are poorly constrained and the horizontal resolution is still too coarse.

Additionally, both configurations offer comparable vertical CO₂ mole fraction profiles and exhibit the same bias in the lower stratosphere relative to observational data. In contrast, gravity waves in the new ICO configuration may need some tuning in order to improve the quality of temperature at the tropopause and in the stratosphere (Lott et al., 2005). Their impact on atmospheric transport of CO₂ at lower altitudes has not been specifically evaluated but seems to be minimal given the small differences shown between the two model configurations. Tuning the LMDZ - DYNAMICO coupling in general is still an ongoing process.

The new ICO configuration offers new opportunities in terms of development. Its use of fewer cells per level for a comparable resolution at the equator results in faster computation times of around 20% in our main simulation and easier-to-store outputs thanks to their smaller size on disk. Unlike regular latitude-longitude grids, ICO does not require a polar filter, whereas these filters generally parallelize badly, on both CPU and GPU. ICO provides a more homogeneous grid at the poles compared to a regular latitude-longitude grid where cell size in kilometer gets smaller when approaching the singularities of the poles. The ICO configuration also allows for much better scaling of computing resources when used at higher resolution, increasing even further the computational speed of the simulations relative to the structured latitude-longitude grid.

While running, REG and ICO can archive specific meteorological variables like air mass fluxes which can then be read by an offline version of the model dedicated to tracer transport. This economical transport model forms the basis of the inversion system of Chevallier et al. (2005) to generate the CO₂ and N₂O inversion products of the Copernicus Atmosphere Monitoring Service of the European Commission (CAMS service, (<https://atmosphere.copernicus.eu/ghg-services>)). Our next task is the implementation of DYNAMICO in this offline model in order to prepare future resolution increases, while limiting the induced increase in computational cost.

Code and data availability

The source code for the REG and ICO configurations is freely available online via the following address: https://forge.ipsl.jussieu.fr/igcmg/browser/CONFIG/publications/ICOLMDZORINCA_CO2_Transport_GMD_2023

under the CeCILL v2 Free Software License (<http://www.cecill.info/index.en.html>, last access: 11 September 2023, CECILL, 2020). The exact version of the model used to produce the results used in this paper is archived on Zenodo, as are input data and scripts to run the model and produce the plots for all the simulations presented in this paper (Lloret et al., 2023).

Author contributions

ZL designed, ran and analyzed the simulations and prepared the paper under direction and advice of FC. AC contributed to the development and preparation of the simulations. MR was also involved in the development of the simulations. YM supported the development of the model. All authors contributed to the text.

Disclaimer

The authors declare that they have no conflict of interest.

Acknowledgements

This work was granted access to the HPC resources of TGCC under the allocation A0130102201 made by GENCI. It was funded by the Copernicus Atmosphere Monitoring Service, implemented by ECMWF on behalf of the European Commission (Grant: CAMS2 55). We thank Olivier Boucher for fruitful discussions on vertical temperature profiles and the LMDZ model, and Sakina Takache for discussions on DYNAMICO. We also thank Julien Derouillat for his help in the implementation of XIOS for our application. We also gratefully acknowledge the many people who contributed atmospheric observations (aircraft, AirCore and surface air sample measurements). Measurements at Jungfraujoch were supported by ICOS Switzerland, which is funded by the Swiss National Science Foundation, in-house contributions and the State Secretariat for Education, Research and Innovation. We thank Kathryn McKain for her comments and suggestions on the best way to describe aircraft and AirCore data. We also thank Bianca Baier for her help with the AirCore data.

References

- Agustí-Panareda, A., Diamantakis, M., Massart, S., Chevallier, F., Muñoz-Sabater, J., Barré, J., Curcoll, R., Engelen, R., Langerock, B., Law, R. M., Loh, Z., Morguí, J. A., Parrington, M., Peuch, V.-H., Ramonet, M., Roehl, C., Vermeulen, A. T., Warneke, T., and Wunch, D.: Modelling CO₂ weather – why horizontal resolution matters, *Atmospheric Chemistry and Physics*, 19, 7347–7376, <https://doi.org/10.5194/acp-19-7347-2019>, 2019.
- Agustí-Panareda, A., McNorton, J., Balsamo, G., Baier, B. C., Bousserez, N., Boussetta, S., Brunner, D., Chevallier, F., Choulga, M., Diamantakis, M., Engelen, R., Flemming, J., Granier, C., Guevara, M., Denier van der Gon, H., Elguindi, N., Haussaire, J.-M., Jung, M., Janssens-Maenhout, G., Kivi, R., Massart, S., Papale, D., Parrington, M., Razinger, M., Sweeney, C., Vermeulen, A., and Walther, S.: Global nature run data with realistic high-resolution carbon weather for the year of the Paris Agreement, *Sci Data*, 9, 160, <https://doi.org/10.1038/s41597-022-01228-2>, 2022.
- Baier, B., Sweeney, C., Newberger, T., Higgs, J., Wolter, S., & NOAA Global Monitoring Laboratory. NOAA AirCore atmospheric sampling system profiles (Version 20230831) [Data set]. NOAA GML. <https://doi.org/10.15138/6AV0-MY81>, 2021.
- Basu, S., Baker, D. F., Chevallier, F., Patra, P. K., Liu, J., and Miller, J. B.: The impact of transport model differences on CO₂ surface flux estimates from OCO-2 retrievals of column average CO₂, *Atmospheric Chemistry and Physics*, 18, 7189–7215, <https://doi.org/10.5194/acp-18-7189-2018>, 2018.
- Cadule, P., Friedlingstein, P., Bopp, L., Sitch, S., Jones, C. D., Ciais, P., Piao, S. L., and Peylin, P.: Benchmarking coupled climate-carbon models against long-term atmospheric CO₂ measurements, *Global Biogeochemical Cycles*, 24, <https://doi.org/10.1029/2009GB003556>, 2010.
- CeCILL: CeCILL and Free Software, <http://www.cecill.info/index.en.html> (last access: 3 July 2023), 2020.
- Chau, T. T. T., Gehlen, M., and Chevallier, F.: A seamless ensemble-based reconstruction of surface ocean pCO₂ and air–sea CO₂ fluxes over the global coastal and open oceans, *Biogeosciences*, 19, 1087–1109, <https://doi.org/10.5194/bg-19-1087-2022>, 2022.

- Chevallier, F., Fisher, M., Peylin, P., Serrar, S., Bousquet, P., Bréon, F.-M., Chédin, A., and Ciais, P.: Inferring CO₂ sources and sinks from satellite observations: Method and application to TOVS data, *Journal of Geophysical Research: Atmospheres*, 110, <https://doi.org/10.1029/2005JD006390>, 2005.
- Chevallier, F., Ciais, P., Conway, T. J., Aalto, T., Anderson, B. E., Bousquet, P., Brunke, E. G., Ciattaglia, L., Esaki, Y., Fröhlich, M., Gomez, A., Gomez-Pelaez, A. J., Haszpra, L., Krummel, P. B., Langenfelds, R. L., Leuenberger, M., Machida, T., Maignan, F., Matsueda, H., Morguá, J. A., Mukai, H., Nakazawa, T., Peylin, P., Ramonet, M., Rivier, L., Sawa, Y., Schmidt, M., Steele, L. P., Vay, S. A., Vermeulen, A. T., Wofsy, S., and Worthy, D.: CO₂ surface fluxes at grid point scale estimated from a global 21 year reanalysis of atmospheric measurements, *Journal of Geophysical Research: Atmospheres*, 115, <https://doi.org/10.1029/2010JD013887>, 2010.
- Chevallier, F., Lloret, Z., Cozic, A., Takache, S., and Remaud, M.: Toward High-Resolution Global Atmospheric Inverse Modeling Using Graphics Accelerators, *Geophysical Research Letters*, 50, e2022GL102135, <https://doi.org/10.1029/2022GL102135>, 2023.
- Ciais, P., Dolman, A. J., Bombelli, A., Duren, R., Peregón, A., Rayner, P. J., Miller, C., Gobron, N., Kinderman, G., Marland, G., Gruber, N., Chevallier, F., Andres, R. J., Balsamo, G., Bopp, L., Bréon, F.-M., Broquet, G., Dargaville, R., Battin, T. J., Borges, A., Bovensmann, H., Buchwitz, M., Butler, J., Canadell, J. G., Cook, R. B., DeFries, R., Engelen, R., Gurney, K. R., Heinze, C., Heimann, M., Held, A., Henry, M., Law, B., Luysaert, S., Miller, J., Moriyama, T., Moulin, C., Myneni, R. B., Nussli, C., Obersteiner, M., Ojima, D., Pan, Y., Paris, J.-D., Piao, S. L., Poulter, B., Plummer, S., Quegan, S., Raymond, P., Reichstein, M., Rivier, L., Sabine, C., Schimel, D., Tarasova, O., Valentini, R., Wang, R., van der Werf, G., Wickland, D., Williams, M., and Zehner, C.: Current systematic carbon-cycle observations and the need for implementing a policy-relevant carbon observing system, *Biogeosciences*, 11, 3547–3602, <https://doi.org/10.5194/bg-11-3547-2014>, 2014.
- Crotwell, A., Lee, H., Steinbacher, M., Eds.: “20th WMO/IAEA Meeting on Carbon Dioxide, Other Greenhouse Gases and Related Measurement Techniques (GGMT-2019),” GAW REPORT No. 255 (World Meteorological Organization, 2020); https://library.wmo.int/doc_num.php?explnum_id=10353, 2020.
- Dubey, S., Dubos, T., Hourdin, F., and Upadhyaya, H. C.: On the inter-comparison of two tracer transport schemes on icosahedral grids, *Applied Mathematical Modelling*, 39, 4828–4847, <https://doi.org/10.1016/j.apm.2015.04.015>, 2015.
- Dubos, T. and Tort, M.: Equations of Atmospheric Motion in Non-Eulerian Vertical Coordinates: Vector-Invariant Form and Quasi-Hamiltonian Formulation, *Monthly Weather Review*, 142, 3860–3880, <https://doi.org/10.1175/MWR-D-14-00069.1>, 2014.
- Dubos, T., Dubey, S., Tort, M., Mittal, R., Meurdesoif, Y., and Hourdin, F.: DYNAMICO-1.0, an icosahedral hydrostatic dynamical core designed for consistency and versatility, *Geosci. Model Dev.*, 8, 3131–3150, <https://doi.org/10.5194/gmd-8-3131-2015>, 2015.
- Ducoudré, N. I., Laval, K., and Perrier, A.: SECHIBA, a New Set of Parameterizations of the Hydrologic Exchanges at the Land-Atmosphere Interface within the LMD Atmospheric General Circulation Model, *J. Climate*, 6, 248–273, [https://doi.org/10.1175/1520-0442\(1993\)006<0248:SANSOP>2.0.CO;2](https://doi.org/10.1175/1520-0442(1993)006<0248:SANSOP>2.0.CO;2), 1993.
- Emanuel, K. A.: A Scheme for Representing Cumulus Convection in Large-Scale Models, *Journal of the Atmospheric Sciences*, 48, 2313–2329, [https://doi.org/10.1175/1520-0469\(1991\)048<2313:ASFRCC>2.0.CO;2](https://doi.org/10.1175/1520-0469(1991)048<2313:ASFRCC>2.0.CO;2), 1991.
- Evangelio, N., Balkanski, Y., Cozic, A., and Møller, A. P.: Simulations of the transport and deposition of ¹³⁷Cs over Europe after the Chernobyl Nuclear Power Plant accident: influence of varying emission-altitude and model horizontal and vertical resolution, *Atmos. Chem. Phys.*, 13, 7183–7198, <https://doi.org/10.5194/acp-13-7183-2013>, 2013.

- Folberth, G. A., Hauglustaine, D. A., Lathière, J., and Brocheton, F.: Interactive chemistry in the Laboratoire de Météorologie Dynamique general circulation model: model description and impact analysis of biogenic hydrocarbons on tropospheric chemistry, *Atmospheric Chemistry and Physics*, 6, 2273–2319, <https://doi.org/10.5194/acp-6-2273-2006>, 2006.
- Fouquart, Y. and Bonnel, B.: Computations of solar heating of the earth's atmosphere: A new parameterization, *Beitr. Phys. Atmos.*, 53, 35–62, 1980.
- Geels, C., Gloor, M., Ciais, P., Bousquet, P., Peylin, P., Vermeulen, A. T., Dargaville, R., Aalto, T., Brandt, J., Christensen, J. H., Frohn, L. M., Haszpra, L., Karstens, U., Rödenbeck, C., Ramonet, M., Carboni, G., and Santaguida, R.: Comparing atmospheric transport models for future regional inversions over Europe – Part 1: mapping the atmospheric CO₂ signals, *Atmospheric Chemistry and Physics*, 7, 3461–3479, <https://doi.org/10.5194/acp-7-3461-2007>, 2007.
- Giorgetta, M. A., Brokopf, R., Crueger, T., Esch, M., Fiedler, S., Helmert, J., Hohenegger, C., Kornblueh, L., Köhler, M., Manzini, E., Mauritsen, T., Nam, C., Raddatz, T., Rast, S., Reinert, D., Sakradzija, M., Schmidt, H., Schneek, R., Schnur, R., Silvers, L., Wan, H., Zängl, G., and Stevens, B.: ICON-A, the Atmosphere Component of the ICON Earth System Model: I. Model Description, *Journal of Advances in Modeling Earth Systems*, 10, 1613–1637, <https://doi.org/10.1029/2017MS001242>, 2018.
- Grandpeix, J.-Y., Phillips, V., and Tailleux, R.: Improved mixing representation in Emanuel's convection scheme, *Q. J. R. Meteorol. Soc.*, 130, 3207–3222, <https://doi.org/10.1256/qj.03.144>, 2004.
- Hall, B. D., Crotwell, A. M., Kitzis, D. R., Mefford, T., Miller, B. R., Schibig, M. F., and Tans, P. P.: Revision of the World Meteorological Organization Global Atmosphere Watch (WMO/GAW) CO₂ calibration scale, *Atmospheric Measurement Techniques*, 14, 3015–3032, <https://doi.org/10.5194/amt-14-3015-2021>, 2021.
- Hazan, L., Tarniewicz, J., Ramonet, M., Laurent, O., and Abbaris, A.: Automatic processing of atmospheric CO₂ and CH₄ mole fractions at the ICOS Atmosphere Thematic Centre, *Atmospheric Measurement Techniques*, 9, 4719–4736, <https://doi.org/10.5194/amt-9-4719-2016>, 2016.
- Herrington, A. R., Lauritzen, P. H., Lofverstrom, M., Lipscomb, W. H., Gettelman, A., and Taylor, M. A.: Impact of grids and dynamical cores in CESM2.2 on the surface mass balance of the Greenland Ice Sheet, *Journal of Advances in Modeling Earth Systems*, 14, e2022MS003192, <https://doi.org/10.1029/2022MS003192>, 2022.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thépaut, J.-N.: The ERA5 global reanalysis, *Quarterly Journal of the Royal Meteorological Society*, 146, 1999–2049, <https://doi.org/10.1002/qj.3803>, 2020.
- Hourdin, F., Musat, I., Bony, S., Braconnot, P., Codron, F., Dufresne, J.-L., Fairhead, L., Filiberti, M.-A., Friedlingstein, P., Grandpeix, J.-Y., Krinner, G., LeVan, P., Li, Z.-X., and Lott, F.: The LMDZ4 general circulation model: climate performance and sensitivity to parametrized physics with emphasis on tropical convection, *Clim Dyn*, 27, 787–813, <https://doi.org/10.1007/s00382-006-0158-0>, 2006.
- Hourdin, F., Grandpeix, J.-Y., Rio, C., Bony, S., Jam, A., Cheruy, F., Rochetin, N., Fairhead, L., Idelkadi, A., Musat, I., Dufresne, J.-L., Lahellec, A., Lefebvre, M.-P., and Roehrig, R.: LMDZ5B: the atmospheric component of the IPSL climate model with revisited parameterizations for clouds and convection, *Clim Dyn*, 40, 2193–2222, <https://doi.org/10.1007/s00382-012-1343-y>, 2013.
- Hourdin, F., Rio, C., Grandpeix, J.-Y., Madeleine, J.-B., Cheruy, F., Rochetin, N., Jam, A., Musat, I., Idelkadi, A., Fairhead, L., Foujols, M.-A., Mellul, L., Traore, A.-K., Dufresne, J.-L., Boucher, O., Lefebvre, M.-P., Millour, E., Vignon, E., Jouhaud, J.,

Diallo, F. B., Lott, F., Gastineau, G., Caubel, A., Meurdesoif, Y., and Ghattas, J.: LMDZ6A: The Atmospheric Component of the IPSL Climate Model With Improved and Better Tuned Physics, *Journal of Advances in Modeling Earth Systems*, 12, e2019MS001892, <https://doi.org/10.1029/2019MS001892>, 2020.

Houweling, S., Aben, I., Breon, F.-M., Chevallier, F., Deutscher, N., Engelen, R., Gerbig, C., Griffith, D., Hungershofer, K., Macatangay, R., Marshall, J., Notholt, J., Peters, W., and Serrar, S.: The importance of transport model uncertainties for the estimation of CO₂ sources and sinks using satellite measurements, *Atmos. Chem. Phys.*, 10, 9981–9992, <https://doi.org/10.5194/acp-10-9981-2010>, 2010.

ICOS RI, Bergamaschi, P., Colomb, A., De Mazière, M., Emmenegger, L., Kubistin, D., Lehner, I., Lehtinen, K., Lund Myhre, C., Marek, M., Platt, S.M., Plaß-Dülmer, C., Schmidt, M., Apadula, F., Arnold, S., Blanc, P.-E., Brunner, D., Chen, H., Chmura, L., Conil, S., Couret, C., Cristofanelli, P., Delmotte, M., Forster, G., Frumau, A., Gheusi, F., Hammer, S., Haszpra, L., Heliasz, M., Henne, S., Hoheisel, A., Kneuer, T., Laurila, T., Leskinen, A., Leuenberger, M., Levin, I., Lindauer, M., Lopez, M., Lunder, C., Mammarella, I., Manca, G., Manning, A., Marklund, P., Martin, D., Meinhardt, F., Müller-Williams, J., Necki, J., O'Doherty, S., Ottosson-Löfvenius, M., Philippon, C., Piacentino, S., Pitt, J., Ramonet, M., Rivas-Soriano, P., Scheeren, B., Schumacher, M., Sha, M.K., Spain, G., Steinbacher, M., Sørensen, L.L., Vermeulen, A., Vítková, G., Xueref-Remy, I., di Sarra, A., Conen, F., Kazan, V., Roulet, Y.-A., Biermann, T., Heltai, D., Hensen, A., Hermansen, O., Komínková, K., Laurent, O., Levula, J., Pichon, J.-M., Smith, P., Stanley, K., Trisolino, P., ICOS Carbon Portal, ICOS Atmosphere Thematic Centre, ICOS Flask And Calibration Laboratory, ICOS Central Radiocarbon Laboratory, 2023. European Obspack compilation of atmospheric carbon dioxide data from ICOS and non-ICOS European stations for the period 1972-2023; obspack_co2_466_GLOBALVIEWplus_v8.0_2023-04-26. <https://doi.org/10.18160/CEC4-CAGK>

Jones, M. W., Andrew, R. M., Peters, G. P., Janssens-Maenhout, G., De-Gol, A. J., Ciais, P., Patra, P. K., Chevallier, F., and Le Quéré, C.: Gridded fossil CO₂ emissions and related O₂ combustion consistent with national inventories 1959-2018, *Scientific Data*, 8, <https://doi.org/10.1038/s41597-020-00779-6>, 2021.

Kenneth N. Schuldt, John Mund, Ingrid T. Luijckx, Tuula Aalto, James B. Abshire, Ken Aikin, Arlyn Andrews, Shuji Aoki, Francesco Apadula, Bianca Baier, Peter Bakwin, Jakub Bartyzel, Gilles Bentz, Peter Bergamaschi, Andreas Beyersdorf, Tobias Biermann, Sebastien C. Biraud, Harald Boenisch, David Bowling, Gordon Brailsford, Willi A. Brand, Huilin Chen, Gao Chen, Lukasz Chmura, Shane Clark, Sites Climadat, Aurelie Colomb, Roisin Commene, Sébastien Conil, Cedric Couret, Adam Cox, Paolo Cristofanelli, Emilio Cuevas, Roger Curcoll, Bruce Daube, Kenneth Davis, Martine De Mazière, Stephan De Wekker, Julian Della Coletta, Marc Delmotte, Joshua P. DiGangi, Ed Dlugokencky, James W. Elkins, Lukas Emmenegger, Shuangxi Fang, Marc L. Fischer, Grant Forster, Arnoud Frumau, Michal Galkowski, Luciana V. Gatti, Torsten Gehrlein, Christoph Gerbig, Francois Gheusi, Emanuel Gloor, Vanessa Gomez-Trueba, Daisuke Goto, Tim Griffis, Samuel Hammer, Chad Hanson, László Haszpra, Juha Hatakka, Martin Heimann, Michal Heliasz, Daniela Heltai, Arjan Hensen, Ove Hermanssen, Eric Hintsa, Antje Hoheisel, Jutta Holst, Viktor Ivakhov, Dan Jaffe, Armin Jordan, Warren Joubert, Anna Karion, Stephan R. Kawa, Victor Kazan, Ralph Keeling, Petri Keronen, Joil Kim, Tobias Kneuer, Pasi Kolari, Katerina Kominkova, Eric Kort, Elena Kozlova, Paul Krummel, Dagmar Kubistin, Casper Labuschagne, David H. Lam, Xin Lan, Ray Langenfelds, Olivier Laurent, Tuomas Laurila, Thomas Lauvaux, Jost Lavric, Bev Law, Olivia S. Lee, John Lee, Irene Lehner, Kari Lehtinen, Reimo Leppert, Ari Leskinen, Markus Leuenberger, Ingeborg Levin, Janne Levula, John Lin, Matthias Lindauer, Zoe Loh, Morgan Lopez, Chris R. Lunder, Toshinobu Machida, Ivan Mammarella, Giovanni Manca, Alistair Manning, Andrew Manning, Michal V. Marek, Melissa Y. Martin, Giordane A. Martins, Hidekazu Matsueda, Kathryn McKain, Harro Meijer, Frank Meinhardt, Lynne Merchant, N. Mihalopoulos, Natasha Miles, John B. Miller, Charles E. Miller, Logan Mitchell, Stephen Montzka, Fred Moore, Heiko Moossen, Eric Morgan, Josep-Anton Morgui, Shinji Morimoto, Bill Munger, David Munro, Cathrine L. Myhre, Meelis Mölder,

- Jennifer Müller-Williams, Jaroslaw Necki, Sally Newman, Sylvia Nichol, Yosuke Niwa, Simon O'Doherty, Florian Obersteiner, Bill Paplawsky, Jeff Peischl, Olli Peltola, Salvatore Piacentino, Jean M. Pichon, Steve Piper, Joseph Pitt, Christian Plass-Duelmer, Stephen M. Platt, Steve Prinzivalli, Michel Ramonet, Ramon Ramos, Enrique Reyes-Sanchez, Scott Richardson, Haris Riris, Pedro P. Rivas, Michael Rothe, Thomas Ryerson, Kazuyuki Saito, Maryann Sargent, Motoki Sasakawa, Bert Scheeren, Martina Schmidt, Tanja Schuck, Marcus Schumacher, Thomas Seifert, Mahesh K. Sha, Paul Shepson, Michael Shook, Christopher D. Sloop, Paul Smith, Martin Steinbacher, Britton Stephens, Colm Sweeney, Lise L. Sørensen, Pieter Tans, Kirk Thoning, Helder Timas, Margaret Torn, Pamela Trisolino, Jocelyn Turnbull, Kjetil Tørseth, Alex Vermeulen, Brian Viner, Gabriela Vitkova, Stephen Walker, Andrew Watson, Ray Weiss, Steve Wofsy, Justin Worsley, Doug Worthy, Dickon Young, Sönke Zaehle, Andreas Zahn, Mirosław Zimnoch, Rodrigo A. de Souza, Alcide G. di Sarra, Danielle van Dinther, Pim van den Bulk; Multi-laboratory compilation of atmospheric carbon dioxide data for the period 1957-2021; obspack_co2_1_GLOBALVIEWplus_v8.0_2022-08-27; NOAA Earth System Research Laboratory, Global Monitoring Laboratory. <http://doi.org/10.25925/20220808>, 2022
- Krinner, G., Viovy, N., de Noblet-Ducoudré, N., Ogée, J., Polcher, J., Friedlingstein, P., Ciais, P., Sitch, S., and Prentice, I. C.: A dynamic global vegetation model for studies of the coupled atmosphere-biosphere system, *Global Biogeochemical Cycles*, 19, <https://doi.org/10.1029/2003GB002199>, 2005.
- Lan, X., J.W. Mund, A.M. Crotwell, M.J. Crotwell, E. Moglia, M. Madronich, D. Neff and K.W. Thoning (2023), Atmospheric Carbon Dioxide Dry Air Mole Fractions from the NOAA GML Carbon Cycle Cooperative Global Air Sampling Network, 1968-2022, Version: 2023-08-28, <https://doi.org/10.15138/wkgj-f215>
- Lauvaux, T., Pannekoucke, O., Sarrat, C., Chevallier, F., Ciais, P., Noilhan, J., and Rayner, P. J.: Structure of the transport uncertainty in mesoscale inversions of CO₂ sources and sinks using ensemble model simulations, *Biogeosciences*, 6, 1089–1102, <https://doi.org/10.5194/bg-6-1089-2009>, 2009.
- Lin, X., Ciais, P., Bousquet, P., Ramonet, M., Yin, Y., Balkanski, Y., Cozic, A., Delmotte, M., Evangeliou, N., Indira, N. K., Locatelli, R., Peng, S., Piao, S., Saunois, M., Swathi, P. S., Wang, R., Yver-Kwok, C., Tiwari, Y. K., and Zhou, L.: Simulating CH₄ and CO₂ over South and East Asia using the zoomed chemistry transport model LMDz-INCA, *Atmospheric Chemistry and Physics*, 18, 9475–9497, <https://doi.org/10.5194/acp-18-9475-2018>, 2018.
- Lloret, Z., Cozic, A., Chevallier, F., Remaud, M., & Meurdesoif, Y.: ICOLMDZORINCA CO₂ Transport GMD 2023. Zenodo. <https://doi.org/10.5281/zenodo.10019679>, 2023.
- Lott, F., Fairhead, L., Hourdin, F., and Levan, P.: The stratospheric version of LMDz: dynamical climatologies, arctic oscillation, and impact on the surface climate, *Clim Dyn*, 25, 851–868, <https://doi.org/10.1007/s00382-005-0064-x>, 2005.
- Mellor, G. L. and Yamada, T.: A Hierarchy of Turbulence Closure Models for Planetary Boundary Layers, *Journal of the Atmospheric Sciences*, 31, 1791–1806, [https://doi.org/10.1175/1520-0469\(1974\)031<1791:AHOTCM>2.0.CO;2](https://doi.org/10.1175/1520-0469(1974)031<1791:AHOTCM>2.0.CO;2), 1974.
- Miles, N.L., S.J. Richardson, K.J. Davis, and B.J. Haupt, 2017. In-situ tower atmospheric measurements of carbon dioxide, methane and carbon monoxide mole fraction for the Indianapolis Flux (INFLUX) project, Indianapolis, IN, USA. Data set. Available on-line <http://datacommons.psu.edu> from The Pennsylvania State University Data Commons, University Park, Pennsylvania, USA. <http://dx.doi.org/10.18113/D37G6P>
- Miles, N.L., S.J. Richardson, D.K. Martins, K.J. Davis, T. Lauvaux, B.J. Haupt, and S.K. Miller. 2018. ACT-America: L2 In Situ CO₂, CO, and CH₄ Concentrations from Towers, Eastern USA. ORNL DAAC, Oak Ridge, Tennessee, USA. <https://doi.org/10.3334/ORNLDAAC/1568>
- Mlawer, E. J., Taubman, S. J., Brown, P. D., Iacono, M. J., and Clough, S. A.: Radiative transfer for inhomogeneous atmospheres: RRTM, a validated correlated-k model for the longwave, *Journal of Geophysical Research: Atmospheres*, 102, 16663–16682, <https://doi.org/10.1029/97JD00237>, 1997.

- Niwa, Y., Tomita, H., Satoh, M., Imasu, R., Sawa, Y., Tsuboi, K., Matsueda, H., Machida, T., Sasakawa, M., Belan, B., and Saigusa, N.: A 4D-Var inversion system based on the icosahedral grid model (NICAM-TM 4D-Var v1.0) – Part 1: Offline forward and adjoint transport models, *Geosci. Model Dev.*, 10, 1157–1174, <https://doi.org/10.5194/gmd-10-1157-2017>, 2017.
- Remaud, M., Chevallier, F., Cozic, A., Lin, X., and Bousquet, P.: On the impact of recent developments of the LMDz atmospheric general circulation model on the simulation of CO₂ transport, *Geosci. Model Dev.*, 25, 2018.
- Rio, C., Hourdin, F., Grandpeix, J.-Y., and Lafore, J.-P.: Shifting the diurnal cycle of parameterized deep convection over land, *Geophysical Research Letters*, 36, <https://doi.org/10.1029/2008GL036779>, 2009.
- Rochetin, N., Grandpeix, J.-Y., Rio, C., and Couvreux, F.: Deep Convection Triggering by Boundary Layer Thermals. Part II: Stochastic Triggering Parameterization for the LMDZ GCM, *Journal of the Atmospheric Sciences*, 71, 515–538, <https://doi.org/10.1175/JAS-D-12-0337.1>, 2014.
- Sakaguchi, K., Harrop, B. E., Leung, L. R., Jang, J., Gettelman, A., Zarzycki, C. M., Lin, G., Zhang, K., and Skamarock, W. C.: The CAM-MPAS model simulations for the HighResMIP experiments, 2020, A090-0014, 2020.
- Sepulchre, P., Caubel, A., Ladant, J.-B., Bopp, L., Boucher, O., Braconnot, P., Brockmann, P., Cozic, A., Donnadiou, Y., Dufresne, J.-L., Estella-Perez, V., Ethé, C., Fluteau, F., Foujols, M.-A., Gastineau, G., Ghattas, J., Hauglustaine, D., Hourdin, F., Kageyama, M., Khodri, M., Marti, O., Meurdesoif, Y., Mignot, J., Sarr, A.-C., Servonnat, J., Swingedouw, D., Szopa, S., and Tardif, D.: IPSL-CM5A2 – an Earth system model designed for multi-millennial climate simulations, *Geoscientific Model Development*, 13, 3011–3053, <https://doi.org/10.5194/gmd-13-3011-2020>, 2020.
- Shepherd, T. G., Semeniuk, K., and Koshyk, J. N.: Sponge layer feedbacks in middle-atmosphere models, *Journal of Geophysical Research: Atmospheres*, 101, 23447–23464, <https://doi.org/10.1029/96JD01994>, 1996.
- Staniforth, A. and Thuburn, J.: Horizontal grids for global weather and climate prediction models: a review, *Quarterly Journal of the Royal Meteorological Society*, 138, 1–26, <https://doi.org/10.1002/qj.958>, 2012.
- Tort, M., Dubos, T., Bouchut, F., and Zeitlin, V.: Consistent shallow-water equations on the rotating sphere with complete Coriolis force and topography, *Journal of Fluid Mechanics*, 748, 789–821, <https://doi.org/10.1017/jfm.2014.172>, 2014.
- Ullrich, P. A., Jablonowski, C., Kent, J., Lauritzen, P. H., Nair, R., Reed, K. A., Zarzycki, C. M., Hall, D. M., Dazlich, D., Heikes, R., Konor, C., Randall, D., Dubos, T., Meurdesoif, Y., Chen, X., Harris, L., Kühnlein, C., Lee, V., Qaddouri, A., Girard, C., Giorgetta, M., Reinert, D., Klemp, J., Park, S.-H., Skamarock, W., Miura, H., Ohno, T., Yoshida, R., Walko, R., Reinecke, A., and Viner, K.: DCMIP2016: a review of non-hydrostatic dynamical core design and intercomparison of participating models, *Geoscientific Model Development*, 10, 4477–4509, <https://doi.org/10.5194/gmd-10-4477-2017>, 2017.
- Van Leer, B.: Towards the ultimate conservative difference scheme. IV. A new approach to numerical convection, *Journal of Computational Physics*, 23, 276–299, [https://doi.org/10.1016/0021-9991\(77\)90095-X](https://doi.org/10.1016/0021-9991(77)90095-X), 1977.
- Wagenhäuser, T., Engel, A., and Sitals, R.: Testing the altitude attribution and vertical resolution of AirCore measurements with a new spiking method, *Atmospheric Measurement Techniques*, 14, 3923–3934, <https://doi.org/10.5194/amt-14-3923-2021>, 2021.
- Zheng, T., Feng, S., Davis, K. J., Pal, S., and Morguá, J.-A.: Development and evaluation of CO₂ transport in MPAS-A v6.3, *Geoscientific Model Development*, 14, 3037–3066, <https://doi.org/10.5194/gmd-14-3037-2021>, 2021.

Prerequisite for high-resolution modeling

Contents

1	Motivation	108
2	Outline of the developments	108
	2.1 Preliminary developments	108
	2.2 Summary of the developments	109
3	Introduction	110
4	Model description	111
	4.1 Variational framework	111
	4.2 Transport model	111
	4.3 Direct model	112
	4.4 Tangent-linear and adjoint versions	113
	4.5 Parallelization strategies	113
5	Exploiting GPU hardware	114
6	Evaluation	114
7	Discussion and conclusions	117
8	Data availability statement	117
9	References	118

THIS CHAPTER PRESENTS THE PRELIMINARY DEVELOPMENTS that were necessary to increase the horizontal resolution of the inverse model presented in [chapter VI](#).

The motivation for these developments is presented in [section 1](#). A short summary and outline of this work and of the article that resulted from it are then presented in [section 2](#). The complete article is then presented in the final sections of this chapter and contains examples of the Fortran code highlighting the parallelization.

1 Motivation

In [chapter III](#) of [Part I](#), we established the importance of improving the spatial resolution of atmospheric transport models while highlighting the increase in computational cost associated with it.

This is especially true when using these transport models for atmospheric inversions at a global scale. A single inversion requires many repeated lengthy simulations, running not only the direct model but also the tangent-linear and adjoint version. This makes any major increase in computational time quickly unsustainable to stay within operational requirements.

As part of the early reflection to make the high resolution inversions described in [chapter X](#) possible, we wanted to explore new possibilities which were opened at the start of this thesis. Namely, that the V100 partition of supercomputer Irene of TGCC was created at the end of 2019 with NVIDIA V100 devices; similarly, the partition of Topaze from Centre de calcul recherche et technologie (CCRT) with the more recent NVIDIA A100 devices was opened at the end of 2021. GPUs looked attractive in principle but we did not know neither the height of the required engineering step for the code porting nor whether we could even efficiently formulate the transport model in terms of independent parallel execution entities, as needed by the GPUs. It was a technological bet that was in practice supported by Frédéric Chevallier for the sake of time: he ported the Fortran code on GPU and did most of the writing of the following paper. The implementation and evaluation of the code once ported on GPU were done jointly. These elements motivate the inclusion of the paper here.

2 Outline of the developments

2.1 Preliminary developments

Initially, the CAMS/LSCE inverse system was only running at a resolution of 3.75° in longitude by 1.90° in latitude with 39 vertical layers. This corresponds to a regular grid of 96 by 95 cells per layer. As we explained in [chapter VII](#), our inverse systems runs on pre-computed mass fluxes taken from the output of a direct simulation from the master GCM. By this point the atmospheric model LMDZ was already routinely used at a resolution of 2.50° longitude by 1.27° latitude for many other applications, but the inverse system was not yet ready to use inputs of mass fluxes at this higher resolution. We previously gave an overview of the history of the CAMS/LSCE inverse model in [chapter VII](#), and we saw that its development was disjointed from that of LMDZ and instead had to be periodically be brought back up to date to match the master GCM. This is not inherently a negative thing, but with the upcoming major changes coming to LMDZ and the GCM (The implementation of DYNAMICO, and horizontal resolution increase) a reevaluation of the inverse system was due.

The 40-year-long simulation of the master GCM on the regular latitude-longitude grid of LMDZ run as part of the comparison between two configurations in [chapter VIII](#) was also used to generate new up-to-date mass fluxes at this higher resolution. The simulation was also extended to the most recent months available from the ERA5 reanalysis to provide up-to-date estimations of CO_2 surface fluxes for CAMS.

In theory, running the inverse system at this new resolution is a straightforward process since the

physics did not change, the biggest adjustments needing to be done being the change in the number of vertical layers and some parameters. However this first increase in resolution multiplied the number of 3D cells by 4.5 which immediately makes the limitations in computing efficiency of the model appear. To be used in an operational setup, the inverse system has strict time-to-solution requirements, but the cost in total computing hours also has to be kept relatively small since new inversions have to regularly be run and consume large amounts of resources. The cost of running the inverse system at this new resolution was already too high, and further increases in spatial resolution would be completely unattainable. A complete refactoring of the model was necessary.

2.2 Summary of the developments

The first step to accelerate the model was simply to remove unneeded code that became obsolete across the different updates over almost two decades. The mass fluxes from the master GCM that are used as inputs greatly grow in size when the resolution increases, but they have been efficiently compressed since the biggest file of convective mass fluxes is a matrix mostly made of zeros.

The essential part of the development was to replace the original MPI parallelization with a parallelization across GPU threads. This includes the transport model itself, but also the tangent-linear and adjoint codes. The basic structure of the models coded in Fortran was kept intact and distribution of parallel calculation on GPU threads was handled by OpenACC directives. The loops structure had to be changed to optimally run on GPU threads. Parallel threading on GPUs consists in running code that can modify different parts of the memory by executing the same instructions on each thread. Therefore loops had to sometimes be adapted so that they each only modified one 3D cell at a time. Certain loops were split into two different ones if possible. And when that was not possible, some recursive loops could be threaded on the whole vertical axis for each cell.

The code was optimized to reduce the time spent transferring data between CPU and GPU since it is a costly operation. It has been limited to only transferring the input data from CPU to GPU, and only output data is transferred out of the GPU. The code running on CPU is not parallelized anymore and the system runs on a single CPU core and one GPU.

The GPU version at the new resolution takes 15 s per month for the direct model, 30 s for the tangent linear and 1.3 min for the adjoint. The reading of the input mass fluxes also takes 25 s per month. The original MPI parallelized version at this resolution running on 15 CPU cores takes 5 min for the direct model, 7.5 min for the tangent linear and 16 min for the adjoint. This makes the new GPU version 8 times faster than the equivalent MPI parallelization running on CPU.

It is also interesting to note that the GPU version of the inverse model at this new resolution is still 2.6 times faster than the CPU version of the old resolution. This means that refactoring the code managed to still accelerate the inverse model despite increasing the number of 3D cells by 4.5 at the same time. This not only entirely solved the issue of computing time for this first increase in resolution but made future ones immediately achievable.

Geophysical Research Letters[®]

RESEARCH LETTER

10.1029/2022GL102135

Key Points:

- The workload in a Eulerian transport model on a longitude-latitude grid can be embarrassingly parallel enough to run efficiently on a Graphics Processing Unit
- The transport model of the Laboratoire de Météorologie Dynamique was run in a new grid of 1,626,768 3D cells
- The calculation time in the direct version of the model is now less than the time to read the input meteorological data

Correspondence to:

F. Chevallier,
frederic.chevallier@lscce.ipsl.fr

Citation:

Chevallier, F., Lloret, Z., Cozic, A., Takache, S., & Remaud, M. (2023). Toward high-resolution global atmospheric inverse modeling using graphics accelerators. *Geophysical Research Letters*, 50, e2022GL102135. <https://doi.org/10.1029/2022GL102135>

Received 17 NOV 2022
Accepted 5 JAN 2023


Author Contributions:

Conceptualization: Zoé Lloret
Methodology: Zoé Lloret, Anne Cozic, Marine Remaud
Software: Zoé Lloret, Anne Cozic, Sakina Takache, Marine Remaud
Writing – original draft: Sakina Takache
Writing – review & editing: Zoé Lloret, Anne Cozic, Sakina Takache, Marine Remaud

© 2023. The Authors.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs License](https://creativecommons.org/licenses/by-nc-nd/4.0/), which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

Toward High-Resolution Global Atmospheric Inverse Modeling Using Graphics Accelerators

Frédéric Chevallier¹ , Zoé Lloret¹ , Anne Cozic¹, Sakina Takache¹ , and Marine Remaud¹ 

¹Laboratoire des Sciences du Climat et de l'Environnement, LSCE/IPSL, CEA-CNRS-UVSQ, Université Paris-Saclay, Gif-sur-Yvette, France

Abstract The ability of global transport models to go up in resolution becomes discriminating for greenhouse gas atmospheric inversions. This paper describes the porting on Graphics Processing Units of the global transport model currently used in the European operational Copernicus Atmosphere Monitoring Service (CAMS) for CO₂ and N₂O inversions. It represents an important milestone to achieve sub-degree resolution. The code includes not only the direct model but also its tangent-linear and its adjoint versions which are needed in variational inversions. Tests were carried out for CO₂ at a resolution of 2.50° in longitude, 1.27° in latitude and 79 layers in the vertical, corresponding to 1,626,768 3D cells, 4.5 times more than the current standard resolution of the model used in the CAMS reanalyses. A month's worth of computation of the tangent-linear and of the adjoint versions now takes 2.5 min, including 50 s for reading meteorological data.

Plain Language Summary Atmospheric transport models are intensively used to infer global greenhouse gas emissions and removals, from atmospheric measurements: a single global analysis involves repeated and long transport simulations. This intensive use has limited the spatial resolution of such analyses despite an increasing need for national greenhouse gas budgets, despite an increasing number of corresponding space-based observations at kilometer resolution, and despite an increasing number of high-quality surface measurements made in sites marked by strong local influences. For the global transport model currently used in the European operational atmosphere monitoring service, our objective within the next 5 years is to make it reach a resolution of about 50 km over the whole globe. This paper describes an important milestone in this direction with the porting of this transport model on hardware components initially developed for video display but now used for high performance computing. Tests were carried out at a resolution of 2.50° in longitude, 1.27° in latitude and 79 layers in the vertical, corresponding to 4.5 times more 3D cells than the current standard resolution of the model. The direct model itself now takes less time than reading the input meteorological data.

1. Introduction

The recent focus on national greenhouse budgets for the preparation of the first Global StockTake of the United Nations Framework Convention on Climate Change (UNFCCC) has renewed the incentive toward higher spatial resolutions of the transport models embedded in atmospheric inversions (Chevallier, 2021; Deng et al., 2022). For instance, despite its two parallelization layers, the transport model in the LSCE inversion system, which derives from Remaud et al. (2018), currently achieves a modest 3.75° in longitude × 1.90° in latitude. This limits the comparison of the inversion results with UNFCCC national inventory reports to large countries or groups of countries (Chevallier, 2021). A specific feature of the LSCE system up to now has been its use in the time-critical environment of the operational Copernicus Atmosphere Monitoring Service (CAMS, <https://atmosphere.copernicus.eu/>) for carbon dioxide (CO₂) and nitrous oxide (N₂O): the “time to solution” of the LSCE system, and therefore its spatial resolution, have been constrained by the desire to keep its latest products to a maximum of a few months from real time. However, this resolution challenge is shared by the current generation of global transport models used for inverse modeling. In version 10 of the Model Intercomparison Project (MIP) of the second Orbiting Carbon Observatory (OCO-2), six out of the 11 participating models were run at 5° in longitude × 4° in latitude or coarser; the CAMS/LSCE system at its standard 3.75° × 1.90° resolution was another participant; the resolution of three other participants was between 2.5° and 3.0° in longitude × 2.0° in latitude globally (with a regional zoom for one of them) (<https://ceos.org/gst/carbon-dioxide.html>, accessed 8 January 2023).

The last participant to the OCO-2 MIP (NIES, Maksyutov et al., 2021) was the only one to generate sub-degree resolution results. To achieve this performance, this contributor coupled a global Eulerian model at medium resolution $3.75^\circ \times 3.75^\circ$ with a $0.1^\circ \times 0.1^\circ$ Lagrangian model, which directly parallelizes along a large number of independent air particle trajectories generated for each assimilated observation. Such an “embarrassingly parallel” workload illustrates a major software trend in high-performance computing (HPC) in general, but the ensemble of air parcel trajectories run for each observation mobilizes large Central Processing Unit (CPU) resources. In order to increase the resolution of the transport model in the LSCE inversion system, a more compact and economical approach is favored here which keeps the model within the Eulerian framework using Graphics Processing Units (GPUs).

While semiconductor miniaturization is stalling, accelerators like GPUs are currently sustaining the development of supercomputer performance (e.g., Leiserson et al., 2020). The word “graphics” in GPU refers to the history of this hardware component for video display, but GPUs are now more generally *manycore processors*, that is processors containing thousands of cores which are dedicated to synchronously running the same series of instructions on different input data and on independent parallel execution entities (the *Single-Instruction Multiple Thread* concept). This specialization allows smaller device size and much faster execution times than the general-purpose CPUs for repetitive and computation-intensive tasks. However, the independency requirement is limiting and favors some types of applications, like deep learning. In the case of our subject here atmospheric global transport models, many instructions may be structured in loops over the model 3D cells. However, the latter are not independent of each other because they exchange tracer mass at each model time step: the adaptation to GPU is therefore not straightforward. The specialization of the GPUs also means that they cannot work alone: the different tasks of an entire program must be distributed between CPU and GPU according to their respective capacity, adding another layer of complexity. The increasing investment of HPC centers in hybrid partitions including GPUs forces the adaptation of computer programs that were written for the general-purpose CPU to this custom hardware, at least partially. On the positive side, this development also represents a strategic opportunity to increase the complexity of models like the one in the LSCE inversion system without increasing execution time.

This paper describes the adaptation of the transport model in the LSCE inversion system to the GPU environment at the Centre de Calcul Recherche et Technologie (CCRT, <https://www-ccrt.cea.fr/>) supercomputing facility. Tests have been made at the resolution of 2.50° in longitude times 1.27° in latitude, with 79 layers in the vertical, hence with 4.5 times more 3D cells than the existing standard configuration. This resolution corresponds to that of its reference general circulation model (GCM) used in the sixth phase of the international Coupled MIP as the atmospheric component of the Institut Pierre Simon Laplace coupled model (Hourdin et al., 2020). The text is structured as follows: the next section describes the transport model; Section 3 describes the algorithm adaptation to GPU; Section 4 presents the validation of the adaptation; conclusions and prospects are drawn in the last section.

2. Model Description

2.1. Variational Framework

In order to meet one of the needs of what was to become the European CAMS operational service, the LSCE developed in 2004 a “variational” global atmospheric inversion system for long-lived trace gases like CO_2 , and N_2O (Chevallier et al., 2005; Thompson et al., 2014). Outside CAMS, this same system was used for methane (CH_4) and related molecules (Berchet et al., 2021; Pison et al., 2009). The advantage of this variational system compared to the previous analytical approach developed at LSCE that explicitly used a matrix of derivatives of the whole transport model (Bousquet et al., 2000), lies in its scalability as to the size of the observation vector (i.e., the number of observations to assimilate) and of the control vector (i.e., the number of variables to be optimized). Indeed, with a variational formulation of the Bayesian inversion problem, the calculation of derivatives can be made at the level of the lines of code, using the chain rule, and not at the level of the whole model, thus allowing a considerable gain of computer time and memory.

2.2. Transport Model

Most of the computation time spent by the LSCE global variational system is invested in three configurations of its Eulerian atmospheric transport model: a standard direct version, a tangent-linear version which implements

a first-order Taylor expansion of the model, and an adjoint version which propagates sensitivities backward in time in the model. These three versions are based on the GCM of the Laboratoire de Météorologie Dynamique (LMDz, Hourdin et al., 2020) nudged to a Numerical Weather Prediction (NWP) re-analysis and reduced to tracer transport equations through the use of meteorological variables precomputed with the full GCM run at the needed resolution (Hourdin et al., 2006).

The model addresses tracer transport by large-scale advection, deep convection, thermal plumes and boundary layer turbulence with the same equations used in the full GCM for tracer transport and using a series of meteorological variables computed at exactly the same spatial resolution from the full LMDz model: the 3D mass of air, 3D large-scale atmospheric mass fluxes for the advection, vertical turbulent exchange coefficients and temperature for the boundary-layer turbulence, vertical mass fluxes for the thermal scheme, various updraft and downdraft mass fluxes and exchange coefficients to drive the deep convection scheme. The values of these variables are 3-hr averages stored in single precision in NetCDF4 format, while the model uses double precision to obtain accurate gradients for the variational application. An illustration for the need of accurate gradients is given by Figure 2 from Chevallier (2013), with an outlier point (in 1990) which disappears when the gradient is filtered.

2.3. Direct Model

The horizontal grid is regular in both latitudinal and longitudinal degrees, with grid points duplicated for computational convenience at the poles along the longitudes, and at the change-of-date line. The pressures of the vertical levels combine a term that varies with the surface pressure and a fixed term that dominates above the tropopause.

The dynamic part contains the numerical resolution of the fluid mechanics equations on a three-dimensional grid, described by Hourdin and Armengaud (1999). The tracer transport equations are iteratively solved over the longitude-latitude grid following the Van Leer I (1977) scheme. Paired with a slope-limiting algorithm, where slopes are computed by finite differences of the tracer distribution, this scheme preserves tracer monotonicity and does not incur spurious numerical oscillations. However, it is non-linear by construction. While the Van Leer scheme is one-dimensional in space, it can be used over a 3D grid by treating the flux in each of the three dimensions alternately. The advance in time is accounted for by integrating the fluxes in each dimension over a half or full time step. Hourdin and Armengaud (1999) use the sequence proposed by Russell and Lerner (1981) which consists of undertaking half a time step in the longitudinal direction, half a time step in the latitudinal direction, then a full time step in the vertical direction, followed again by half a time step in latitude, then half in longitude. This sequence of operations has the advantage of preserving a uniform tracer field regardless of the divergent character of the wind field (Lin & Rood, 1996).

In each direction, the concentration slopes are first calculated, then slope-limiting is applied. The air mass fluxes in the considered direction are then calculated based on the wind field. The mass fluxes and limited slopes are then used to evaluate the tracer concentration reaching the cell boundary in the given direction following the Van Leer I scheme. The boundary value, multiplied by the mass flux, gives the tracer flux traversing the given edge. The tracer amount entering the cell is added to the existing amount in the cell for an updated value in the given direction in time (see, e.g., Figure 1 in Hourdin & Armengaud, 1999). The flux contribution is accumulated in the cell in three dimensions from the three upwind directions.

A special treatment is required in the high latitudes as the cells become thinner with the converging longitudes by nature of the longitude-latitude grid. Given constant time steps δt , the Courant number $u \cdot \delta t / \delta x$ — u being the speed of wind in the longitudinal direction x —may exceed 1 due to the diminishing value of δx . This means that the wind field along x traverses multiple cells in one time step. The cell reached at the end of this time step thus has accumulated tracer mass from multiple upwind cells. The tracer fluxes from the traversed cells are summed to form the flux traversed through the edge of the destination cell following Equation 17 of Hourdin and Armengaud (1999) which is non-linear. In practice, this treatment reduces the effective number of grid cells toward the pole.

The representation of unresolved subgrid scales in the model is based on Hourdin et al. (2006). Tracer transport by turbulent diffusion in the planetary boundary layer is represented by vertical exchange coefficients (Equation 6 in Hourdin et al., 2006). It is completed by a thermal model, formulated in terms of mass fluxes, for the case of convective boundary layers (Rio & Hourdin, 2008). Transport by deep convection is derived from the Emanuel (1991) scheme (see, e.g., Pilon et al., 2015, and references therein), which, in practice, links all vertical levels between the cloud base and the level of neutral buoyancy together in a convective system. This link makes

the data volume of convective mass fluxes quadratically dependent on the number of vertical levels of the model. However, in practice, most of the global exchange matrix is made of zeros. In order to save disk space and reduce the input data flow, the matrix is stored in compressed sparse row format. The three subgrid-scale parameterizations redistribute tracer mass in the vertical in a strictly linear way.

The code of the transport model is written in Fortran and is run in chunks of 1 month. The inversion system, which is a code in Python language (see Berchet et al., 2021, for the latest version), manages the connection between the months for the three model versions.

2.4. Tangent-Linear and Adjoint Versions

As expected for a Bayesian inversion system, the LSCE system minimizes the usual Bayesian cost function, which represents the balance between observations and prior information in the inversion solution (see, e.g., Equation 9 in Rayner et al., 2019). The LSCE system is “variational” in that it exploits the gradient of the cost function in an iterative way, rather than by an analytical expression or by a Monte Carlo search. The minimizer itself is either a limited memory quasi-Newton method, the MIQN3 software from Gilbert and Lemaréchal (1989) or, systematically for CO₂, a Lanczos version of the conjugate gradient algorithm developed at ECMWF (Fisher, 1998). If the operator that links the variables to estimate (greenhouse gas surface fluxes) and the observations to assimilate (atmospheric concentrations) is linear, the cost function is quadratic and the conjugate gradient converges super-linearly, thereby reducing the number of needed iterations and saving computing time. For this motif, an exact tangent-linear and its transpose (adjoint) were developed for the LMDz off-line transport model.

The tangent-linear is the first-order Taylor expansion of the model. It is used to compute the difference between the model state and the observations in the Bayesian cost function, with the linearization point obtained from the prior value of the control vector. Its transpose, or adjoint, appears in the expression of the gradient of the cost function and is therefore also important for a variational inversion system (see, e.g., Equation 6 of Chevallier et al., 2005).

In practice, the tangent-linear and adjoint codes were derived manually from the reference direct model line by line, in order to directly control the computational efficiency. The adjoint derivation was particularly delicate as it requires the availability of the forward state each time there is a non-linearity in the model: state-dependent conditional instructions coming from the slope limiters or from the special treatment of advection in the high latitudes; updated values of the mass of the atmosphere to multiply with; a profile scaling when assimilating satellite retrievals in order to conserve the column-average mixing ratio while interpolating to the retrieval grid (Chevallier, 2015). This availability of the forward state is not straight-forward because the adjoint runs backward in time. In order to save computer memory, a mixed approach which combines checkpoint files—containing the tracer state every 3 hr—and re-computations in-between was chosen. When entering a 3-hr slot (from its end), the adjoint code reads the tracer state at the start of the slot, runs the forward subroutines for 3 hr forward, while keeping in memory the tracer state at the start of each call to an advection subroutine within the Russell and Lerner (1981) sequence described above (60 times with a 15-min time step for the 3D advection).

2.5. Parallelization Strategies

Since the inception of the LSCE variational system, various opportunities have been taken to increase the speed of the three configurations of this offline model to allow for longer assimilation windows, increased spatial resolutions, and increased model accuracy.

In 2008, the transport model was parallelized with Message Passing Interface (MPI), allowing multi-year inversions (Chevallier et al., 2010). The globe was split into a series of latitude bands that run on different cores. The bands communicate through MPI for the advection along the latitudes. For the two horizontal resolutions discussed here (using respectively 96 and 143 grid cells along the latitudes), the best speed is achieved already at about 10 cores. Every 3 hr, each core reads the meteorological variables of its own latitude band.

In 2012, another layer of parallelization, based on physical considerations, was introduced in the tangent-linear and adjoint versions (Chevallier, 2013). The tangent-linear and the adjoint versions are run in parallel temporal segments which partially overlap. The tracer increments in the tangent-linear and the adjoint sensitivities are carried out from one parallel segment to the next through a global bias term. The simulation that provides the linearization point is only parallelized through MPI and does not use the bias-term simplification. This bias term assumes that all mole fraction increments are uniformly mixed in the global atmosphere. Unwanted side

effects of this approximation are damped by a months-long overlap period and by an update of the linearization point during the minimization. The increased speed was successively invested into decadal inversions (Le Quéré et al., 2015), then into increased spatial resolution (2.5 more model 3D cells) and increased complexity of the subgrid parameterizations (Locatelli et al., 2015).

The parallelization with MPI in the LMDz transport model could be extended to longitudes and the vertical axis. However, the distributed-memory approach of MPI does not leverage the shared intra-node memory of current multicore nodes. The parallelization of the LMDz transport model therefore needs to be upgraded for better computational efficiency. A solution which takes a single CPU core but parallelizes across GPU threads using OpenACC directives is described in the next section.

3. Exploiting GPU Hardware

We have chosen to keep the LMDz transport model in Fortran but we have restructured and augmented it with OpenACC directives throughout the code, particularly before each loop. These generic directives guide the compiler in the management of the GPU, specifying each step of the data flow between the memory of the CPU, the memory of the GPU and the GPU threads.

As explained in the introduction, codes running on GPUs primarily modify different locations of the computer memory in parallel threads that execute the same series of instructions. The LMDz code has therefore been restructured in order to expose such situations as much as possible. For many loops, independent parallel threads can be identified across all 3D cells of the global atmosphere, as with initializations. Where a single loop modified the variables from two cells at once, we applied loop fission and split it into two loops (Case 1 in Table 1). Recursive statements exist in the sub-grid-scale parameterizations or for pressure computations (from 3D mass to 2D surface pressure, or from 2D surface pressure to 3D pressures), but only along the vertical axis. These recursive loops were placed inside the loops across the horizontal directions so that they would be executed sequentially within the threads, thus avoiding the latency incurred by a loop on the GPU instructions (Case 2 in Table 1). Similarly, averages are performed in the two polar rows after the advection along the latitudes: the corresponding loop over the polar longitudes was also set as an inner sequential one. Along the longitude axis, the advection can happen over more than one cell in a time step with the formulation mentioned above. This is not an issue for the forward and tangent-linear versions because the initial code was already localized, leaving only a sum over several grid cells (in Equation 17 of Hourdin & Armengaud, 1999). This was already the inner loop, but had simply to be made sequential within each thread. However, in the adjoint version, the process is reversed: the adjoint variable of this sum feeds tracer-mass adjoint variables in up to several neighboring cells at once. We have reformulated this part by temporarily storing the tracer-mass adjoint variables in a matrix. This matrix is of dimension longitudes \times longitudes $\times N$, with N being the number of cells in the 3D atmosphere, in which the tracer mass is fully advected in less than one time step. Once the matrix is built, the tracer-concentration adjoint variables can be updated by the GPU, with the loop over N performed sequentially on each thread (Case 3 in Table 1).

For the linearization points in the adjoint computations, we noticed that recomputing the 3-hr tracer states is slightly faster than reading them from the disk, so we chose this option. The memory volume needed to save the tens of intermediate states between two consecutive 3-hr states (see Section 2d) would exceed resources so we keep the re-computation in 3-hr chunks there. In total, running the adjoint code therefore implies two full direct simulations, in addition to the adjoint computations themselves.

The transfer time between CPU and GPU is significant in general, so we structured the code in order to avoid back-and-forth transfers: only the input data of the code is pushed from the CPU to the GPU and only the output data comes out of the GPU (Case 4 in Table 1). Since the CPU role is essentially reduced to input/output tasks, we removed all MPI instructions and code instructions related to latitude banding. The code is made much simpler, but we lose the parallel (but memory-intensive) reading of the meteorological variables which was distributed along the cores. The data flow therefore directly uses a single CPU core and one GPU.

4. Evaluation

Our tests were performed on the hybrid partition of the Topaze supercomputer of CCRT. Topaze is a BullSequana XH2000 made by Atos and delivered in 2021. Its hybrid partition is composed of 48 nodes, each of which has four

Table 1
Four Transformations of the Code for an Efficient Use by the GPU

Case	Initial code	New code
1. Loop fission in vlx ad() – to update q_ad_glo(ij,l-1,iq) and q_ad_glo(ij,l,iq) separately. Subroutine vlx_ad() deals with the adjoint of the vertical advection.	<pre> DO l=2,llm DO ij=ijb,ije IF (dzqw(ij,l) > 0.) THEN dzqw_ad(ij,l)= dzqw_ad(ij,l) + adzqw_ad(ij,l) adzqw_ad(ij,l) = 0. ELSE dzqw_ad(ij,l)= dzqw_ad(ij,l) - adzqw_ad(ij,l) adzqw_ad(ij,l) = 0. ENDIF q_ad(ij,l-1) = q_ad(ij,l-1) + dzqw_ad(ij,l) q_ad(ij,l) = q_ad(ij,l) - dzqw_ad(ij,l) dzqw_ad(ij,l) = 0. ENDDO ENDDO </pre>	<pre> !\$acc parallel loop collapse(2) DO l=2,llm DO ij=1,ip1jmp1 IF (dzqw(ij,l) > 0.) THEN dzqw_ad(ij,l)= dzqw_ad(ij,l) + adzqw_ad(ij,l) ELSE dzqw_ad(ij,l)= dzqw_ad(ij,l) - adzqw_ad(ij,l) ENDIF q_ad_glo(ij,l-1,iq) = q_ad_glo(ij,l-1,iq) + dzqw_ad(ij,l) ENDDO ENDDO !\$acc parallel loop collapse(2) DO l=2,llm DO ij=1,ip1jmp1 q_ad_glo(ij,l,iq) = q_ad_glo(ij,l,iq) - dzqw_ad(ij,l) ENDDO ENDDO </pre>
2. Sequential inner loop in timeloop() – to minimize the number of GPU calls. Subroutine timeloop() manages the time axis.	<pre> DO l=1,llm DO j=jj_begin,ji_end DO i=1,iip1 smass(i,j)=smass(i,j)+masse(i,j,l) ENDDO ENDDO DO j=jj_begin,ji_end DO i=1,iip1 ps(i,j)=smass(i,j)/aire(i,j)*g END DO END DO </pre>	<pre> !\$acc parallel loop collapse(2) DO j=1,jip1 DO i=1,iip1 smass=0. !\$acc loop seq DO l=1,llm smass=smass+masse_glo(i,j,l) ENDDO ps_glo(i,j)=smass/aire_glo(i,j)*g END DO END DO </pre>
3. Temporary transfer matrix in vlx ad() – to create independent parallel threads for the update of q_ad_glo(ijq,l,iq) and dxq_ad(ijq,l). Subroutine vlx_ad() deals with adjoint of the longitudinal advection.	<pre> DO l=llm,1,-1 IF(nl(l)>0) THEN DO iju=1,niju(l) ij=induc(iju,l) j=(j-1)/iip1+1 zu_m=q_m(i,j,l) IF(zu_m > 0.) THEN ijq=ij i=ijq-(j-1)*iip1 DO WHILE(zu_m > masse(ijq,l)) q_ad(ijq,l) = q_ad(ijq,l) + u_mq_ad(ijq,l)*masse(ijq,l) zu_m=zu_m-masse(ijq,l) i=MOD(i-2+iim,iim)+1 ijq=(j-1)*iip1+i ENDDO q_ad(ijq,l) = q_ad(ijq,l) + zu_m * u_mq_ad(ijq,l) dxq_ad(ijq,l) = dxq_ad(ijq,l) & + zu_m*masse(ijq,l)*u_mq_ad(ijq,l) ELSE ijq=ij+1 i=ijq-(j-1)*iip1 DO WHILE(-zu_m > masse(ijq,l)) q_ad(ijq,l) = q_ad(ijq,l) - u_mq_ad(ijq,l)*masse(ijq,l) zu_m=zu_m+masse(ijq,l) i=MOD(i,iim)+1 ijq=(j-1)*iip1+i ENDDO q_ad(ijq,l) = q_ad(ijq,l) + zu_m * u_mq_ad(ijq,l) dxq_ad(ijq,l) = dxq_ad(ijq,l) & - zu_m*0.5*(1+zu_m/masse(ijq,l))*u_mq_ad(ijq,l) ENDIF u_mq_ad(ijq,l) = 0. ENDDO ENDIF ENDDO </pre>	<pre> (...) !\$acc parallel loop collapse(3) DO l=1,llm DO j=2,jip1-1 DO il=1,iip1 IF (indj(i,j) > 0.) THEN ijq=(j-1)*iip1+1 !\$acc loop seq DO i2=1,iip1 q_ad_glo(ijq,l,iq) = q_ad_glo(ijq,l,iq) + qtmp_ad(i1,i2,indj(i,j,l)) dxq_ad(ijq,l) = dxq_ad(ijq,l) + dxqtmp_ad(i1,i2,indj(i,j,l)) ENDDO ENDIF ENDDO ENDDO </pre>
4. All arrays are on the GPU in vlx() – to minimize transfer times to and from the GPU. Subroutine vlx() deals with the longitudinal advection.	<pre> SUBROUTINE vlx_p(q,pente_max,masse,u_m,ijb_x,ije_x) USE Parallel IMPLICIT NONE INCLUDE "dimensions.h" INCLUDE "paramet.h" INCLUDE "logic.h" INCLUDE "convert.h" INCLUDE "comeconst.h" REAL :: masse(ip1jmp1,llm),pente_max REAL :: u_m(ip1jmp1,llm) REAL :: q(ip1jmp1,llm) INTEGER :: ij,l,i,iju,ijq,indu(ip1jmp1),niju INTEGER :: n0,iadvplus(ip1jmp1,llm),nl(llm) REAL :: new_m_zu_m_zdum(ip1jmp1,llm) REAL :: dxq(ip1jmp1,llm),dxqu(ip1jmp1,llm) REAL :: zz(ip1jmp1) REAL :: adxqu(ip1jmp1),dxqmax(ip1jmp1,llm) REAL :: u_mq(ip1jmp1,llm) REAL :: SSUM EXTERNAL :: SSUM INTEGER :: ijb,ije,ijb_x,ije_x </pre>	<pre> SUBROUTINE vlx(iq,pente_max,zzpbar) USE modglob ! contains the arrays already on the GPU IMPLICIT NONE INCLUDE "dimensions.h" INCLUDE "paramet.h" INCLUDE "logic.h" INCLUDE "convert.h" INCLUDE "comeconst.h" REAL :: pente_max REAL :: u_m(ip1jmp1,llm) REAL :: zzpbar,zzw INTEGER :: iq INTEGER :: ij,l,i,ijq REAL :: new_m(ip1jmp1,llm),zu_m_zdum(ip1jmp1,llm) REAL :: dxq(ip1jmp1,llm),dxqu(ip1jmp1,llm) REAL :: adxqu(ip1jmp1,llm),dxqmax(ip1jmp1,llm) REAL :: u_mq(ip1jmp1,llm) !\$acc create(u_m_zdum,u_mq,new_m,dxqu,adxqu,dxqmax,dxq) </pre>

Note. Some loop indices have been changed because the initial subroutines are called per latitude band whereas the new code, which no longer uses MPI, processes the entire globe at once. MPI, Message Passing Interface; GPU, Graphics Processing Unit.




Figure 1. The top row shows differences between the Message Passing Interface-based simulation of CO₂ and the corresponding Graphics Processing Unit-based one at model vertical levels 10 and 50 (around 980 and 7 hPa, respectively). The two simulations at level 10 are shown in the bottom row. The simulations started on 1 January 1979 at 00:00.

NVIDIA A100 GPUs alongside two AMD EPYC Milan 7763 processors. We worked in batch mode on a large “SCRATCH” file system with data transfer rates of about 4 Gb per second per node, by the virtue of Solid-State Drive disks. The compiler was nvfortran version 22.2.

The model was run in a new grid of 2.50° in longitude, 1.27° in latitude and 79 layers in the vertical—that is, 1,626,768 3D cells. The volume of the input meteorological data at that resolution for 1 month is 15 GB, which can be read on a node of Topaze in about 25 s. For tests over a single month, even if the calculations require a single CPU core, it was necessary to reserve four CPU cores from the same node in order to have enough memory. For regular applications over longer times, job submission rules on Topaze require reserving half a processor (32 cores) to access 1 GPU. Note that for the tests, the model was also fed with the location and time of surface measurements so that it simulates them while running.

The direct model was validated against the previous MPI version. Figure 1 shows an MPI-based simulation of CO₂ and the corresponding GPU-based one after 1 month at vertical level 10 of the model (around 980 hPa). It also displays the differences between the two simulations at levels 10 and 50 (around 980 and 7 hPa, respectively): they are marginal, at most of the order of 10⁻⁶% close to the surface, as a result of a different data flow in the two codes, in particular in the sub-grid-scale parameterizations.

The GPU version was then modified in order to make the values of the grid-cell surface area consistent throughout the code, which was hardly possible in the previous code structure. The convergence of the tangent-linear version was checked for varying perturbation sizes and the adjoint was validated against the tangent-linear—by developing the expression of the norm of the tangent-linear output in order to make the adjoint appear. An accuracy of 100 times the epsilon of the machine was found.

In terms of pure computing time, the direct model took 15 s for 1 month with the GPU; the tangent-linear version took 30 s and the adjoint 1.3 min. These times are in addition to the 25 s of the mass flux reading in each version. They hardly varied when repeated multiple times. In comparison, the code parallelized with MPI on 15 CPU cores needed 5, 7.5, and 16 min (time to solution) in its direct, tangent-linear and adjoint versions, respectively, reading included. Adding more cores did not make the code run faster. For a variational system that exploits the quadraticity of the cost function with the Lanczos version of the conjugate gradient, one iteration uses one tangent-linear simulation and one adjoint simulation for each month processed: in this case, the gain in speed

for the GPU code is eight-fold. Interestingly, the GPU code at this resolution is even 2.6 times faster than the MPI code at the previous resolution of $3.75^\circ \times 1.90^\circ \times 39$ layers.

The new direct model can also be run serially on 1 CPU core by ignoring the OpenACC directives when compiling. In this case, it gives the same results as with the GPU but 10 times slower.

5. Discussion and Conclusions

The growing need for information on regional greenhouse-gas fluxes, for example, within the framework of the Global StockTake of the Paris agreement, the kilometeric resolution now achieved by satellite observations of CO_2 and CH_4 , the strong local influences at many new greenhouse-gas measurement stations, all constitute a strong incentive toward global transport models in atmospheric inversions at higher resolution than the existing ones. Higher resolutions do not systematically reduce transport model errors and their benefit in inversion systems is damped by prior error correlations, but they allow a finer representation of orography and coastlines and a finer description of emission hotspots, when available. However, for long-lived greenhouse gases like CO_2 , N_2O and CH_4 , years-long or even decades-long reanalyses need to be updated frequently as more observations or improved satellite retrievals become available (e.g., Friedlingstein et al., 2022).

For the LMDz transport model currently used in the CAMS operational service for CO_2 and N_2O inversions, our goal is to dramatically increase the rate of resolution upgrades over time, with a target resolution of around 1° over the whole globe at the end of 2023 and around 50 km in less than 5 years for routine inversions without lengthening production time. This ambition is supported by the corresponding development of HPC which is moving toward exascale capability, but must be accompanied by a parallelization strategy that fits the corresponding structure of the hardware resources. We showed here that the workload in a Eulerian transport model on a three-dimensional longitude-latitude grid can be made “embarrassingly parallel” enough to run efficiently on a GPU, without any simplification. This efficiency is evidenced by the fact that the calculation time itself in the direct version of the model is now less than the time to read the input meteorological data. The calculation time of the tangent-linear version is comparable to the reading time for the meteorological data, while the adjoint is more than twice as slow, penalized by the management of the linearization points. The code performance not only optimizes the GPU resources, but also leaves room to run the Monte Carlo ensembles of inversions that are critical for the computation of uncertainty in variational systems (Chevallier et al., 2007).

Reading the meteorological data on the node is now the main computation bottleneck. The fact that the transport calculation is taken by the GPU leaves little activity on the CPU: the reading and writing tasks, and the management of the GPU. The code is therefore most suitable for computer partitions with a small ratio of CPU cores to GPU device per node and with large node memory, as is the case for deep learning applications. In comparison, the Topaze machine we tested on is more hybrid and requires reserving tens of cores to access a single GPU.

The next development step for the LMDz transport model is the application to higher-resolution horizontal grids. We expect the speed of the new code to be proportional to the number of atmospheric columns, but note that the time step must be reduced with the grid refinement for the advection along latitudes. In the course of these resolution upgrades, our strategy is to closely follow the scientific and technical developments in the full LMDz GCM, which is part of the Earth system model of Institut Pierre-Simon-Laplace. In particular, the advection scheme of the LMDz GCM is now moving away from regular longitude-latitude grids, in order to avoid their computationally-expensive polar singularity (Dubos et al., 2015). The transport model described here is already evolving accordingly in order to be able to exploit the meteorological variables generated on the new grid, currently made of hexagons. Technically, its developments in the coming years will follow the evolution of Fortran-compiler support for GPU parallelism and will be guided by the evolution of HPC resources toward exascale computing, particularly in France.

Data Availability Statement

This version of the LMDz global transport model, v3.1, is publicly available from <https://doi.org/10.5281/zenodo.7324039> (Chevallier, 2022) under the Creative Commons Attribution 4.0 International licence.

Acknowledgments

This work was granted access to the HPC resources of CCRT under the allocation CEA/DRF, and of TGCC under the allocation A0110102201 made by GENCI. It was funded by the Copernicus Atmosphere Monitoring Service, implemented by ECMWF on behalf of the European Commission (Grant: CAMS2 55). The authors would like to thank Yann Meurdesoif (LSCE) for his long-standing support with the model developments and his advice about GPUs in particular.

References

Berchet, A., Sollum, E., Thompson, R. L., Pison, I., Thanwerdas, J., Broquet, G., et al. (2021). The community inversion framework v1.0: A unified system for atmospheric inversion studies. *Geoscientific Model Development*, *14*(8), 5331–5354. <https://doi.org/10.5194/gmd-14-5331-2021>

Bousquet, P., Peylin, P., Ciais, P., Quere, C., Friedlingstein, P., & Tans, P. (2000). Regional changes in carbon dioxide fluxes of land and oceans since 1980. *Science*, *290*(5495), 1342–1346. <https://doi.org/10.1126/science.290.5495.1342>

Chevallier, F. (2013). On the parallelization of atmospheric inversions of CO₂ surface fluxes within a variational framework. *Geoscientific Model Development*, *6*(3), 783–790. <https://doi.org/10.5194/gmd-6-783-2013>

Chevallier, F. (2015). On the statistical optimality of CO₂ atmospheric inversions assimilating CO₂ column retrievals. *Atmospheric Chemistry and Physics*, *15*(19), 11133–11145. <https://doi.org/10.5194/acp-15-11133-2015>

Chevallier, F. (2021). Fluxes of carbon dioxide from managed ecosystems estimated by national inventories compared to atmospheric inverse modeling. *Geophysical Research Letters*, *48*(15), e2021GL093565. <https://doi.org/10.1029/2021GL093565>

Chevallier, F. (2022). LMDz transport model (3.1). *Zenodo*. <https://doi.org/10.5281/zenodo.7324039>

Chevallier, F., Bréon, F.-M., & Rayner, P. J. (2007). The contribution of the Orbiting Carbon Observatory to the estimation of CO₂ sources and sinks: Theoretical study in a variational data assimilation framework. *Journal of Geophysical Research*, *112*(D9), D09307. <https://doi.org/10.1029/2006JD007375>

Chevallier, F., Ciais, P., Conway, T. J., Aalto, T., Anderson, B. E., Bousquet, P., et al. (2010). CO₂ surface fluxes at grid point scale estimated from a global 21-year reanalysis of atmospheric measurements. *Journal of Geophysical Research*, *115*(D21), D21307. <https://doi.org/10.1029/2010JD013887>

Chevallier, F., Fisher, M., Peylin, P., Serrar, S., Bousquet, P., Bréon, F.-M., et al. (2005). Inferring CO₂ sources and sinks from satellite observations: Method and application to TOVS data. *Journal of Geophysical Research*, *110*(D24), D24309. <https://doi.org/10.1029/2005JD006390>

Deng, Z., Ciais, P., Tzompa-Sosa, Z. A., Saunio, M., Qiu, C., Tan, C., et al. (2022). Comparing national greenhouse gas budgets reported in UNFCCC inventories against atmospheric inversions. *Earth System Science Data*, *14*(4), 1639–1675. <https://doi.org/10.5194/essd-14-1639-2022>

Dubos, T., Dubey, S., Tort, M., Mittal, R., Meurdesoif, Y., & Hourdin, F. (2015). DYNAMICO-1.0, an icosahedral hydrostatic dynamical core designed for consistency and versatility. *Geoscientific Model Development*, *8*(10), 3131–3150. <https://doi.org/10.5194/gmd-8-3131-2015>

Emanuel, K. A. (1991). A scheme for representing cumulus convection in large-scale models. *Journal of the Atmospheric Sciences*, *48*(21), 2313–2329. [https://doi.org/10.1175/1520-0469\(1991\)048<2313:ASFRCC>2.0.CO;2](https://doi.org/10.1175/1520-0469(1991)048<2313:ASFRCC>2.0.CO;2)

Fisher, M. (1998). Minimization algorithms for variational data assimilation. In *Proceedings of seminar on recent developments in numerical methods for atmospheric modelling, 7–11 September 1998* (pp. 364–385). ECMWF. Retrieved from <https://www.ecmwf.int/node/9400>

Friedlingstein, P., Jones, M. W., O'Sullivan, M., Andrew, R. M., Bakker, D. C. E., Hauck, J., et al. (2022). Global carbon budget 2021. *Earth System Science Data*, *14*(4), 1917–2005. <https://doi.org/10.5194/essd-14-1917-2022>

Gilbert, J.-C., & Lemaréchal, C. (1989). Some numerical experiments with variable-storage quasi-Newton algorithms. *Mathematical Programming*, *45*(1–3), 407–435. <https://doi.org/10.1007/bf01589113>

Hourdin, F., & Armengaud, A. (1999). The use of finite-volume methods for atmospheric advection of trace species. Part I: Test of various formulations in a general circulation model. *Monthly Weather Review*, *127*(5), 822–837. [https://doi.org/10.1175/1520-0493\(1999\)127<0822:tuofvm>2.0.co;2](https://doi.org/10.1175/1520-0493(1999)127<0822:tuofvm>2.0.co;2)

Hourdin, F., Rio, C., Grandpeix, J.-Y., Madeleine, J.-B., Cheruy, F., Rochetin, N., et al. (2020). LMDZ6A: The atmospheric component of the IPSL climate model with improved and better tuned physics. *Journal of Advances in Modeling Earth Systems*, *12*(7), e2019MS001892. <https://doi.org/10.1029/2019MS001892>

Hourdin, F., Talagrand, O., & Idelkadi, A. (2006). Eulerian backtracking of atmospheric tracers. II: Numerical aspects. *Quarterly Journal of the Royal Meteorological Society*, *132*(615), 585–603. <https://doi.org/10.1256/qj.03.198.B>

Leiserson, C. E., Emer, J. S., Kuszmaul, B. C., Lampson, B. W., Sanchez, D., & Schardl, T. B. (2020). There's plenty of room at the top: What will drive computer performance after Moore's law? *Science*, *368*, 6495. <https://doi.org/10.1126/science.aam9744>

Le Quéré, C., Moriarty, R., Andrew, R. M., Peters, G. P., Ciais, P., Friedlingstein, P., et al. (2015). Global carbon budget 2014. *Earth System Science Data*, *7*(1), 47–85. <https://doi.org/10.5194/essd-7-47-2015>

Lin, S., & Rood, R. B. (1996). Multidimensional flux-form semi-Lagrangian transport schemes. *Monthly Weather Review*, *124*(9), 2046–2070. [https://doi.org/10.1175/1520-0493\(1996\)124<2046:mffslt>2.0.co;2](https://doi.org/10.1175/1520-0493(1996)124<2046:mffslt>2.0.co;2)

Locatelli, R., Bousquet, P., Hourdin, F., Saunio, M., Cozic, A., Couvreux, F., et al. (2015). Atmospheric transport and chemistry of trace gases in LMDz5B: Evaluation and implications for inverse modelling. *Geoscientific Model Development*, *8*(2), 129–150. <https://doi.org/10.5194/gmd-8-129-2015>

Maksyutov, S., Oda, T., Saito, M., Janardanan, R., Belikov, D., Kaiser, J. W., et al. (2021). Technical note: A high-resolution inverse modelling technique for estimating surface CO₂ fluxes based on the NIES-TM-FLEXPART coupled transport model and its adjoint. *Atmospheric Chemistry and Physics*, *21*(2), 1245–1266. <https://doi.org/10.5194/acp-21-1245-2021>

Pilon, R., Grandpeix, J.-Y., & Heinrich, P. (2015). Representation of transport and scavenging of trace particles in the Emanuel moist convection scheme. *Quarterly Journal of the Royal Meteorological Society*, *141*(689), 1244–1258. <https://doi.org/10.1002/qj.2431>

Pison, I., Bousquet, P., Chevallier, F., Szopa, S., & Hauglustaine, D. A. (2009). Multi-species inversion of CH₄, CO, and H₂ emissions from surface measurements. *Atmospheric Chemistry and Physics*, *9*(14), 5281–5297. <https://doi.org/10.5194/acp-9-5281-2009>

Rayner, P. J., Michalak, A. M., & Chevallier, F. (2019). Fundamentals of data assimilation applied to biogeochemistry. *Atmospheric Chemistry and Physics*, *19*(22), 13911–13932. <https://doi.org/10.5194/acp-19-13911-2019>

Remaud, M., Chevallier, F., Cozic, A., Lin, X., & Bousquet, P. (2018). On the impact of recent developments of an atmospheric general circulation model on the simulation of CO₂ transport. *Geoscientific Model Development*, *11*, 4489–4513. <https://doi.org/10.5194/gmd-11-4489-2018>

Rio, C., & Hourdin, F. (2008). A thermal plume model for the convective boundary layer: Representation of cumulus clouds. *Journal of the Atmospheric Sciences*, *65*(2), 407–425. <https://doi.org/10.1175/2007jas2256.1>

Russell, G. L., & Lerner, J. A. (1981). A new finite-differencing scheme for the tracer transport equation. *Journal of Applied Meteorology*, *20*(12), 1483–1498. [https://doi.org/10.1175/1520-0450\(1981\)020<1483:anfdfs>2.0.co;2](https://doi.org/10.1175/1520-0450(1981)020<1483:anfdfs>2.0.co;2)

Thompson, R. L., Chevallier, F., Crotwell, A. M., Dutton, G., Langenfelds, R. L., Prinn, R. G., et al. (2014). Nitrous oxide emissions 1999–2009 from a global atmospheric inversion. *Atmospheric Chemistry and Physics*, *14*(4), 1801–1817. <https://doi.org/10.5194/acp-14-1801-2014>

Van Leer, B. (1977). Towards the ultimate conservative difference scheme. Part IV: A new approach to numerical convection. *Journal of Computational Physics*, *23*(3), 276–299. [https://doi.org/10.1016/0021-9991\(77\)90095-x](https://doi.org/10.1016/0021-9991(77)90095-x)

High-resolution inversions

Contents

1	Motivation and strategy for the evolution of the inverse system	120
2	Outline of the study	121
2.1	Setup of the study	121
2.2	Summary of the results	122
2.3	Conclusion of the study	123
3	Key points	124
4	Abstract	125
5	Plain Language Summary	125
6	Introduction	125
7	Model and inversion setup	127
7.1	Inversion system	127
7.2	Inversion configuration	128
7.3	Evaluation	129
7.3.1	CO ₂ data for evaluation	129
7.3.2	Processing of the surface stations	131
7.3.3	Processing of the column-averaged CO ₂ and vertical profiles	132
7.3.4	Processing of the surface flux estimates	132
8	Results and discussion	133
8.1	Surface stations	133
8.2	TCCON observations	136
8.3	Vertical profiles	138
8.4	Regional fluxes	142
8.5	Local fluxes	145
9	Conclusion	146
10	Data availability statement	148
11	References	148
12	Appendix A : Observation datasets	156

THIS CHAPTER CONSISTS OF A STUDY that examines the differences between two inversions of atmospheric CO₂, each using a different horizontal resolution. Two inversions were run over the years 2015-2016, assimilating data from the OCO-2 satellite and were then validated against independent observations.

In [section 1](#) we present the motivation for this study as well as the development strategy for the CAMS/LSCE system. A short summary and outline of the study and article is then presented in [section 2](#). The complete article that was created from this study is then presented in the final sections of this chapter.

1 Motivation and strategy for the evolution of the inverse system

In [chapter IX](#), we presented the refactoring of the inverse model that was done in order to run it on a GPU. As part of this work, the inverse model resolution was updated from the outdated 96×96×39 resolution to the 144 ×143 ×79 resolution used as a reference in [chapter VIII](#). Thanks to the GPU version being much faster than the original CPU version of the model, this first resolution increase actually ended up accelerating the inverse system. The mass fluxes generated by the 40-year long simulation run in [chapter VIII](#) could then easily be used to run large-scale inversions for CAMS.

The massive gains in computational speed thanks to the GPU therefore gave us the opportunity to immediately increase the horizontal resolution of our model for a second time. The model was then run on a regular latitude-longitude grid with 256 ×256 ×79 cells. This could be done for a minimal engineering cost compared to porting the inverse model to run on the DYNAMICO icosahedral grid. This porting required a complete rewrite of the inverse transport model, tangent linear and adjoint and was being carried out by other members of the team without a clear outlook on the time still needed to make it operational. Therefore, the choice was made to first increase the resolution of the regular latitude-longitude grid, run the very first inversions with it and evaluate the impact of this change before proceeding with inversions using DYNAMICO.

To run inversions at this new high-resolution we needed to provide mass fluxes as an input at the corresponding resolution. The 40-year long simulation run in [chapter VIII](#) that provided mass fluxes for what became our new reference resolution came at a significant computational cost. This was acceptable since a long running and high-frequency output was in any case necessary to compare the reference GCM with the new configuration using DYNAMICO. But a simulation could not be run in an identical manner at an even higher resolution for such a long time. The original configuration of the master GCM was coupled with a chemistry model (INCA) which could not easily be run at a higher resolution without significant computational cost. Thankfully, the generation of atmospheric mass fluxes does not necessarily need the fully coupled configuration of the master GCM since they do not depend on the atmospheric chemistry model. We could therefore modify our version of the master GCM to only use the atmospheric component of LMDZ and the land model ORCHIDEE. The land model is still necessary to maintain the stability of the surface temperature and comes at only a small computational cost.

The cost of running the GCM to generate these mass fluxes is however still very high, and this could not be mitigated with the use of GPUs since this model could only be run on CPUs and refactoring it was impossible for this thesis given its scale compared to the extremely compact code of the inverse model. A team of 10 people are currently working on achieving this, and it is expected to still take months or even years. As we saw in [subsection 2.3 of chapter VIII](#), this model is limited in scaling and in how many CPUs it can use in parallel, which places an absolute limit on how fast it can generate inputs for the inverse system. Furthermore, the generated mass fluxes files have a very large size on disk, since they are not compressed when outputted by the GCM, but instead by a dedicated script afterward. This means that each simulated month takes up about 300 GB of disk space, and that for future continued use of this GCM to inverse system pipeline, the compression algorithm must be

implemented directly in the outputs of the GCM. With the limited time available to us we therefore decided to limit the simulation length of this first inversion to two years.

In particular, we wanted to study the impact of an increase in spatial resolution on an inversion assimilating satellite data. As discussed in [section 1 of chapter III](#), the assimilation of very high resolution satellite data for atmospheric inversions of greenhouse gases is becoming ubiquitous and future missions will only make this more prevalent. Since these observations are much more spatially precise than the coarse grids of the transport model, it is hoped that increasing the spatial resolution of the models will allow better assimilation of these data.

With this study, we aim to answer several questions:

- What is the computational performance of our high-resolution inverse model in operational conditions?
- Does a higher horizontal resolution give a benefit in the assimilation of OCO-2 satellite data?
- Does the higher resolution improve the representation of global atmospheric transport of CO₂?
- Does the increase in resolution significantly change the results of the estimates of CO₂ surface fluxes?

2 Outline of the study

To evaluate the performance of our inverse system at this new high resolution, we compared an inversion at two different resolutions over the years 2015 and 2016 while assimilating retrieval of the column-averaged dry-air mole fraction of CO₂ from the OCO-2 satellite. We then compared the result of these inversions to independent observations of CO₂, and also directly compared the results of the two inversions.

2.1 Setup of the study

We performed two CO₂ inversions at the global scale over the years 2015 and 2016 with an extra 3 months for spin-up in 2014 and 3 months for spin-down in 2017. The two inversions differed by their horizontal resolution. The reference one, or low-resolution (LR) model was run on a 144 × 143 × 79 regular latitude-longitude grid corresponding to 1.27° in latitude and 2.50° in longitude. The new high-resolution (HR) model was run on a 256 × 256 × 79 regular latitude-longitude grid corresponding to 0.70° in latitude and 1.41° in longitude. The time step of the HR model had to be reduced to 3 min for the horizontal advection, 6 min for the vertical advection and 12 min for the subgrid processes.

The same prior fluxes of CO₂ (Ocean, biomass burning, biosphere, and fossil emissions) were used for both inversions and interpolated to their respective horizontal resolution. Only midday clear-sky observations over land of XCO₂ from the OCO-2 satellite were assimilated (Atmospheric Carbon Observations from Space, ACOS v1.1.1). Both inversions were done in 40 iterations on 1 CPU and 1 NVIDIA A100 GPU.

The final state of the inversion was compared to independent measurements of CO₂ mole fraction from surface stations and AirCore flights and to observations of XCO₂ from TCCON stations. The method for this comparison was similar to the seasonal decomposition in the previous study in [chapter VIII](#).

The final estimates of global CO₂ surface fluxes as well as the horizontal and vertical distribution of CO₂ concentration were aggregated by month and then directly compared between the two resolutions.

Regional estimates of CO₂ surface fluxes were aggregated per Transcom3 region and compared between resolutions.

2.2 Summary of the results

The inversions took 4 days and 4 hours for the LR model and 9 days and 15 hours for the HR model on one GPU and CPU each. This increase in computing time is to be contrasted with the equivalent sixfold increase in the number of computing operations required for the HR model.

Both resolutions capture the average seasonal cycle well with an average correlation coefficient of 0.90 across all surface stations. The amplitude of this seasonal cycle is also adequately modeled by both resolutions. The performance in terms of synoptic variability at surface stations is slightly better at the higher resolution. The stations where the HR model performs better are located on the coast or in mountainous regions. Some urban stations located in two cities in the continental USA however perform worse with the HR model.

When comparing the final state of the inversions to XCO₂ observations from TCCON stations, both resolutions perform almost exactly the same, whether in terms of direct comparison with the observations or when looking at the modeled seasonal cycle.

Both resolutions have a similar vertical profile of CO₂ when compared to measurements from AirCore flights. But under 3 km altitude the HR model performs better and is closer to measurements. Above 22 km, both resolutions differ from measurements and show a positive bias. In general the difference between model and AirCore measurements shows a lower bias and standard deviation with the HR model than the LR model.

A direct comparison of zonal vertical profiles of CO₂ and XCO₂ averaged per year and season show that the largest difference in concentration between the two resolutions is in the tropics with the HR model having higher values of XCO₂. It is the opposite for latitudes over 60°N, particularly in boreal winter. This difference in XCO₂ is less than 0.10 ppm when averaged per longitude band.

The estimates of CO₂ surface fluxes at the global scale after the inversion are similar between the resolutions, with the HR model showing a slightly lower carbon sink than the LR model for both years (-3.31 vs -3.22 GtC/year for 2015 and -3.88 vs -3.71 GtC/year for 2016). The HR model assigns a higher sink to the land and a lower one to the ocean.

When comparing surface fluxes per Transcom3 region, only a few regions exhibit significant differences between resolutions. In the North American Boreal region, the higher carbon sink of the HR model is paralleled by a better assimilation of the satellite data in this region. This is not the case for the other regions or at the global scale.

At the local scale, we see that the improvement in the representation of CO₂ transport at coastal stations is driven by a better representation of the coastline, with the higher resolution changing many areas previously modeled by mixed land-ocean cells into a finer grid that more clearly defines the coast.

2.3 Conclusion of the study

- The resolution of the CAMS/LSCE inversion system was successfully increased to a global resolution of 0.7° latitude and 1.4° longitude.
- The computational time required for the inversion was only doubled when using the higher-resolution, showcasing the strong scaling of the GPU parallelization.
- The high-resolution model shows a small benefit in CO_2 transport for some coastal and mountainous stations. This is mostly driven by a better orography and coastal definition thanks to the higher-resolution grid.
- The high-resolution model performs better in terms of vertical profile of CO_2 when compared to AirCore measurements, particularly under 3 km.
- The high-resolution inversion estimates a slightly higher global yearly carbon sink than the low-resolution inversion. The high-resolution inversion attributes a higher portion of this sink to the land rather than the ocean.
- A difference in the estimate of the surface flux in the North American Boreal region is paralleled by a better fit of the high-resolution model to the assimilated satellite data, suggesting that in this region with sparse retrievals the high-resolution provides a benefit.

Refining the Global Picture: the Impact of Increased Resolution on CO₂ Atmospheric Inversions using OCO-2 XCO₂ retrievals

Zoé Lloret¹, Frédéric Chevallier¹ and Anne Cozic¹

¹Laboratoire des Sciences du Climat et de l'Environnement, LSCE/IPSL, CEA-CNRS-UVSQ, Université Paris-Saclay, Gif-sur-Yvette, France

Corresponding author: Zoé Lloret (zoe.lloret@lsce.ipsl.fr)

Key Points

- We upgraded our global atmospheric inverse system to 0.7 degree latitude by 1.4 degree longitude with a modest computational overhead.
- The resolution increase improves CO₂ transport representation, benefiting coastal stations the most.
- Global flux estimates in 2015-2016 were similar across resolutions, the high-resolution attributes higher sinks to land than to the ocean.

Abstract

The threat posed by the increasing concentration of carbon dioxide (CO₂) in the atmosphere motivates a detailed and precise estimation of CO₂ emissions and removals over the globe. This study refines the spatial resolution of the CAMS/LSCE inversion system, achieving a global resolution of 0.7° latitude and 1.4° longitude, or three times as many grid boxes as the current operational setup. In a two-year inversion assimilating the midday clear-sky retrievals of the column-average dry-air mole fraction of carbon dioxide (XCO₂) from NASA's second Orbiting Carbon Observatory (OCO-2), the elevated resolution demonstrates an improvement in the representation of atmospheric CO₂, particularly at the synoptic time scale, as validated against independent surface measurements. Vertical profiles of the CO₂ concentration differ slightly above 22 km between resolutions compared to AirCore profiles, and highlight differences in the vertical distribution of CO₂ between resolutions. However, this disparity is not evident for XCO₂, as evaluated against independent reference ground-based observations. Global and regional estimates of natural fluxes for 2015-2016 are similar between the two resolutions, but with North America exhibiting a higher natural sink at high-resolution for 2016. Overall, both inversions seem to yield reasonable estimates of global and regional natural carbon fluxes. The increase in calculation time is less than the increase in the number of operations and in the volume of input data, revealing greater efficiency of the code executed on a Graphics Processing Unit. This allows us to make this higher resolution the new standard for the CAMS/LSCE system.

Plain Language Summary

Human activities have significantly increased the amount of carbon dioxide (CO₂) in the atmosphere, a major driver of climate change. Accurately quantifying CO₂ emissions and removals, known as fluxes, is crucial for implementing effective mitigation strategies. Inverse models are computer programs that analyze large amounts of CO₂ observations to estimate surface fluxes that best match these observations in space and time. While satellites provide extremely precise CO₂ observations all around the Earth, most inverse models lack the resolution to fully utilize this data at a large scale. Our study doubled the horizontal resolution of our inverse model, enhancing its performance and spatial precision when using data from the OCO-2 satellite. Thanks to Graphics Processing Units (GPU) acceleration, the computational cost remained manageable. This improved resolution is now being implemented in the European Copernicus Atmosphere Monitoring Service, with ongoing efforts to further improve the resolution. This advancement promises a more detailed understanding of global CO₂ dynamics, supporting climate change mitigation efforts.

1 Introduction

The escalating carbon dioxide (CO₂) concentration in the atmosphere, driven by anthropogenic emissions, is a primary catalyst for climate change. Notably, the Intergovernmental Panel on Climate Change (IPCC) estimates a global mean surface temperature increase of approximately 1.07°C during the period 2011-2019 compared to the pre-industrial era (1850-1900) (IPCC, 2019), underscoring the urgency of addressing greenhouse gas emissions, particularly CO₂, to damp climate variations. Precise spatio-temporal estimations of these emissions are imperative for effective mitigation strategies.

While direct measurements of carbon fluxes provide essential insights for that goal, their spatial coverage remains limited for mapping extensive regions globally. However, contemporary direct measurements of CO₂ mole fractions are abundant in numerous regions worldwide, complemented by valuable satellite observations offering a macroscopic view of CO₂ distribution. Leveraging this wealth of information, inverse atmospheric transport systems within a Bayesian framework enable the inference of CO₂ sources

and sinks by optimizing surface fluxes based on observed CO₂ mole fractions and analyzed meteorological variables.

These inversions, whether conducted at a global or regional scale, grapple with inherent uncertainties, particularly at finer scales. Notably, the Global Carbon Budget 2023 of the Global Carbon Project (Friedlingstein et al., 2023) revealed significant spread across inversions, with estimates of the net atmosphere-to-surface sink in the northern latitudes (>30° N) from 2013 to 2022 ranging between 1.7 and 3.3 GtC yr⁻¹. Much of this spread is attributed to errors in the transport models (Basu et al., 2018). A notable limitation in the current global models employed in the Global Carbon Budget is actually their coarse horizontal resolution, averaging only 2.80° in latitude and 2.93° in longitude in the 2023 edition. The same issue was present in the v10 Model Intercomparison Project (MIP) of the second Orbiting Carbon Observatory (OCO-2) aimed to characterize the influence of transport model and inversion methods on flux estimates: the average resolution of all the global transport models employed in the v10 OCO-2 MIP intercomparison was only 3.4° latitude by 4.4° longitude (Byrne et al., 2023).

Augmenting the resolution of transport models holds promise, even at large scale (Liu et al., 2024), reducing numerical errors and thereby fostering convergence among different models (Prather, 2008). Increasing the horizontal resolution presents an opportunity for mitigating the representativeness error (Tolk et al., 2008). However, this effect is not universally applicable across all resolutions and does not follow a linear trend. Notably, while kilometer-scale resolutions have demonstrated positive impacts, particularly in regions with complex terrain (Hedelius et al., 2017), the same does not hold true at the scale of hundreds of kilometers, where an increase in horizontal resolution may not necessarily diminish this error (Lin et al., 2018).

Interestingly, the few inversions driven by OCO-2 satellite data in the Global Carbon Budget 2023 show a smaller difference between the latitudes north of 30°N and those further south in their estimates of the net atmosphere-land flux compared to inversions driven by surface observations. This could be due to additional information obtained when using the spatially dense OCO-2 retrievals (Friedlingstein et al., 2023) and such a benefit of the retrievals would be better exploited at higher model resolution.

The needs of the United Nations Framework Convention on Climate Change (UNFCCC), recommending the evaluation of national emission inventories compared to atmospheric inversions (IPCC, 2006, 2019), further reinforces the necessity of this resolution increase (Chevallier, 2021). While this makes high-resolution targets likely in the future for most inverse systems, it remains of crucial scientific interest to judiciously evaluate the costs and benefits associated with augmenting the horizontal resolution of atmospheric models, in order to optimize computing resources, energy use and processing times.

Indeed, resolution enhancement comes at a considerable computational cost given the intricate demands of global inverse models involving prolonged data assimilation windows, complex statistical inversion schemes, and stable atmospheric modeling under the Courant–Friedrichs–Lewy condition (Courant et al., 1928). This condition imposes that for a given velocity field, when the resolution of the spatial discretization increases, the time step of the simulation must be reduced to maintain stability. The quadratic growth in the size of modeled 3D atmospheric fields with horizontal resolution necessitates a judicious balance between resolution increments and expected performance gains.

The transport model used in the CO₂ inversion system of the European operational Copernicus Atmosphere Monitoring Service (CAMS) (<https://atmosphere.copernicus.eu/>) underwent a first horizontal resolution increase back in 2015, doubling the number of vertical layers from 19 to 39 (Locatelli et al., 2015), and a substantial upgrade of the physics in 2018 (Remaud et al., 2018). Tests at higher spatial and vertical resolutions (another doubling of the vertical layers to 79, and a doubling of the number of horizontal boxes to reach a resolution around 2 degrees over the whole globe) proved inadequate for accurately simulating atmospheric dynamics in regions characterized by complex topography, such as mountainous areas (Remaud et al. 2018): the increased 3D resolution did not yield a significant improvement compared to observational data, underscoring the need for further refinement, particularly to show improvement at the synoptic timescale (Agustí-Panareda et al., 2019). The vertical profiles of CO₂ concentration were not significantly affected by changes in resolution unlike the XCO₂ fields, especially around emission hotspots. The high computing cost associated with this resolution increase delayed its implementation in the production chain of the CAMS CO₂ inversion product until the code was ported on Graphics Processing Units (GPUs) in 2023 (Chevallier et al., 2023). The migration also opened the possibility of further resolution increases while maintaining a processing time, or "time to solution", compatible with operational constraints.

This study investigates the effect of enhancing horizontal resolution on global-scale CO₂ inversion to about 1 degree. The comparison entails evaluating the outcomes of a two-year inversion at an increased resolution, assimilating OCO-2 data, against a reference configuration and independent observations. The choice of the OCO-2 data, rather than surface or other satellite measurements, is linked to their global coverage, rapid availability and exceptional quality, making them a backbone of low-latency carbon cycle monitoring. The study examines both the influence of horizontal resolution on atmospheric CO₂ transport and the overall impact on the final estimates of carbon fluxes. The subsequent section delineates the inverse system and the experimental setup, followed by a presentation of results compared to independent observations between low and high resolutions in Section 3. Section 4 succinctly summarizes the findings and concludes with insights derived from this resolution increase.

2 Model and inversion setup

2.1 Inversion system

The inversion system that is used to perform global CO₂ and N₂O atmospheric inversions for CAMS has been developed in the LSCE since 2004 (Chevallier et al. 2005). The same system has also been used outside CAMS for other tracers, such as methane (Berchet et al., 2021), carbon monoxide, or nitrogen oxides (Fortems-Cheiney et al., 2021).

This inverse system is based on a 4D variational approach of the Bayesian inversion problem: assimilating observational data of CO₂ concentrations to derive an optimal state of CO₂ fluxes given a prior estimate of the CO₂ fluxes.

Mathematically, this consists in iteratively minimizing a cost function J which is defined as follows:

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}^b)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^b) + \frac{1}{2}(\mathbf{H}\mathbf{x} - \mathbf{y})^T \mathbf{R}^{-1}(\mathbf{H}\mathbf{x} - \mathbf{y}) \quad (1)$$

Here, \mathbf{x} represents the vector of the variables being optimized, which, in this case, corresponds to successive global maps of the CO₂ fluxes at the resolution of the transport model, weekly throughout the inversion window, and to the 3D state of CO₂ at the start of the inversion window. \mathbf{x}^b means the vector of the prior state of \mathbf{x} , and \mathbf{y} represents the assimilated observations. The matrices \mathbf{R} and \mathbf{B} correspond to the

error covariance matrices associated with the uncertainty of the assimilated observations, as defined from the transport model, and of the prior fluxes, respectively. The linearized operator \mathbf{H} projects the control vector \mathbf{x} into the observation space: it is primarily based on the transport model. In our case, the transport model is an off-line version of the general circulation model (GCM) of the Laboratoire de Météorologie Dynamique (LMDZ) in its latest version, LMDZ6A (Remaud et al., 2018; Hourdin et al., 2020). The off-line version only solves tracer transport equations, driven by pre-computed air mass fluxes from a reference run of the full GCM nudged to the 3-hourly horizontal winds from the fifth generation ECMWF reanalysis (ERA5). The code of the off-line transport model corresponds to the one made public by Chevallier et al. (2023) with some memory optimizations in order to accommodate the larger arrays of the new resolution. The inversion system, coded in Python and run on CPUs, orchestrates the connection across monthly runs of the transport model, coded in Fortran and basically run on GPUs, ensuring the coherence and continuity of the inversion process.

The minimization of J is done iteratively by calculating its gradient using the adjoint version of the transport model and a conjugate gradient algorithm (Fisher, 1998; Chevallier et al., 2005).

2.2 Inversion configuration

To assess the impact of the resolution increase on our inverse system, we conducted two global-scale CO₂ inversions around years 2015 and 2016, incorporating three months for spin-up in 2014 and three months for spin-down in 2017, at two distinct horizontal resolutions. The inversion of reference, referred to as the low-resolution (LR) model throughout the text, operates on a latitude-longitude grid with dimensions of 1.27° in latitude, 2.50° in longitude, and 79 vertical layers, totaling 1,626,768 cells with each cell of size 140 km by 278 km at the equator. The new resolution, designated as the high-resolution (HR) model hereafter, utilizes a latitude-longitude grid with dimensions of 0.70° in latitude, 1.41° in longitude, and 79 vertical layers, resulting in a total of 5,177,344 cells with each cell of size 78 km by 157 km at the equator. The model time step of the LR is 5 minutes for horizontal advection, 10 minutes for vertical advection and 20 minutes for subgrid processes. In order to respect the Courant–Friedrichs–Lewy condition for stability in the HR, it has to go down to 3 minutes for horizontal advection and 6 minutes for vertical advection; for subgrid processes, we reduce it as well to 12 minutes. In both LR and HR configurations, the pre-computed air mass fluxes are 3-hourly averages.

Both inversions share identical prior states for CO₂ fluxes, which are interpolated onto their respective grids, incorporating the following data sources with their native resolution:

- CO₂ fluxes over the ocean are based on the CMEMS-LSCE-FFN 2022 estimates at a native monthly 0.25° resolution (Chau et al. 2022).
- CO₂ biomass burning emissions are from the GFED4.1s inventory at a native monthly 0.25° resolution. No atmospheric source of CO₂ is considered.
- CO₂ fossil emissions are based on GCP-GridFEDv2023.1 estimates at a native monthly 0.1° resolution (Jones et al., 2021).
- Natural fluxes of CO₂ from the biosphere are based on a climatology of 3-hourly averaged estimates from the ORCHIDEE model, version 2.2, revision 7262 (Krinner et al., 2005 ; Friedlingstein et al., 2022) at a 0.5° resolution.

Observations of midday clear-sky total column-averaged CO₂ concentrations over land from the OCO-2 satellite were assimilated, specifically NASA’s Atmospheric CO₂ Observations from Space (ACOS) bias-

corrected nadir and glint land retrievals of X_{CO_2} , version 11.1 (OCO-2/OCO-3 Science Team et al., 2022, O'Dell et al., 2018, , 2023). OCO-2 ocean observations were not used in this study, neither were observations over mixed land-water surfaces. Only data flagged as "good" were used, as 10-second averages, i.e. about 67 km along the orbit track, with an averaging procedure implemented at LSCE and similar to the one defined in the OCO-2 MIP (Crowell et al., 2019). In order to account for likely correlations between the transport model errors at the sub-grid scale, we de-weighted the OCO-2 binned retrievals that fall within a same LMDz grid box for a same orbit by inflating the assigned error variance by the number of retrievals in the box.

The retrievals initially adhered to the X2007 scale of the World Meteorological Organization (WMO). We converted them to the X2019 scale following Hall et al., (2021):

$$X_{2019} = 1.00079 \cdot X_{2007} - 0.142 \text{ ppm} \quad (2)$$

When assimilating the satellite retrievals, the prior and averaging kernel of each retrieval were used in the model. No other data was assimilated so that flasks, in-situ and ground-based X_{CO_2} observations are fully independent.

The spatial correlations of the prior uncertainty, which drive the off-diagonal terms of \mathbf{B} in Equation 1, decay exponentially with a length of 500 km over land and 1000 km over sea for both resolutions. The standard deviations over land are proportional to the climatological daily-varying heterotrophic respiration flux simulated by ORCHIDEE and are constant in $\text{gC}\cdot\text{m}^{-2}$ per day over the ocean. They were tuned at each resolution so that over a full year, the total 1-sigma uncertainty for the prior land fluxes amounts to $2.9 \text{ GtC}\cdot\text{yr}^{-1}$, and for the open ocean to a global air-sea flux 1-sigma uncertainty of $0.2 \text{ GtC}\cdot\text{yr}^{-1}$.

Both inversions were performed over 40 iterations, on 1 CPU and 1 NVIDIA A100 GPU as in Chevallier et al., (2023). The inversion system may be accelerated with a physical parallelization in which the years are run in parallel on different GPUs with a spin-up period for each (Chevallier, 2013), but this possibility has not been exploited here.

The inversions took 4 days and 4 hours for the LR model and 9 days and 15 hours for the HR model. This twofold increase in overall inversion computing time is much smaller than the sixfold increase in the number of operations within the transport model: threefold for the number of global grid cells and an additional twofold for the number of time steps. It is less than the extra computations induced by the ninefold increase in the dimension of the prior error covariance matrix \mathbf{B} . It is also relatively less than what the threefold increase in the volume of transport model input data implies on reading time. Since the computer code is the same between the two resolutions, the relatively modest increase in calculation time reveals better efficiency of our code with increased resolution, which is not unexpected with GPUs, since higher resolutions allow larger loops that better keep the GPUs busy.

2.3 Evaluation

We evaluated the two inversions by directly comparing their final state and estimates of CO_2 fluxes at the global, regional, and local scales. We also compared them to independent observations of CO_2 concentrations.

2.3.1 CO_2 data for evaluation

To assess the agreement between our simulated tracer concentrations and observed data, we sampled mole fraction fields at the nearest cell center, model level (when relevant), and timestamp for each data point. We

utilized high-quality measurements from the CO₂ GLOBALVIEWplus v8.0_2022-08-27 ObsPack database (Schuldt et al., 2022, Miles et al., 2017, Miles et al., 2018, ICOS RI, et al., 2023, Lan et al., 2023) on the WMO CO₂ X2019 scale (Hall et al., 2021). For AirCore, we used Version 20230831 of the dataset from NOAA (Baier et al., 2021). We also exploited ground-based XCO₂ retrievals from the Total Carbon Column Observing Network (TCCON, Wunch et al., 2011) from which we selected in 2015 and 2016 twenty Fourier transform spectrometers around the globe (Buschmann et al., 2022, C et al., 2022, Deutscher et al., 2023, Dubey et al., 2022, Iraci et al., 2022, Kivi et al., 2022, Maziere et al., 2022, Morino et al., 2022a, Morino et al., 2022b, Notholt et al., 2022, Sherlock et al., 2022, Shiomi et al., 2022, Strong et al., 2022, Sussmann and Rettinger, 2017, Te et al., 2022, Warneke et al., 2022, Wennberg et al., 2022a, Wennberg et al., 2022b, Wennberg et al., 2022c, Wunch et al., 2022).

Similar to prior studies involving inverse modeling with LMDZ, we only selected measurements that could be well modeled by a transport model, particularly avoiding tracer accumulation at low altitudes. For in-situ surface stations located under 1000 m above sea level (a.s.l.), we only considered data from 12:00 to 16:00 local time, for in-situ stations above 1000 m a.s.l., only nighttime data from 00:00 to 4:00 local time were retained. We kept all flask measurements.

The observations were categorized into three groups: surface in-situ and flask measurements, AirCore flight measurements, and remote-sensing observations from the OCO-2 mission and TCCON sites. Vertical profiles of CO₂ mole fraction were obtained using AirCore, an atmospheric sampling system that collects successive samples of ambient air (Karion et al. 2010, Baier et al., 2021). From the Obspack dataset, 112 surface stations were selected for analysis, excluding those with fewer than 1200 measurement points over the 2-year study period that passed the initial data selection criteria. The full list of Obspack and TCCON stations used is available as a Supplement. All samples from AirCore data were retained.

The uncertainty associated with the in-situ and flask CO₂ mole fraction measurements used in this study is approximately 0.1 micromol per mol (or part per million, ppm), as detailed in Crotwell et al. (2020) for systematic errors and Hazan et al. (2016) for standard deviation. This uncertainty is considered negligible compared to the model uncertainty stemming from transport errors, estimated to be around 1 ppm under 3000 m (Lauvaux et al., 2009). The altitude determination error for AirCore measurements due to storage diffusion can be substantial, ranging from approximately 250 m below 20 km to 1 km above that altitude (Wagenhäuser et al., 2021). The uncertainty of the measurements of the AirCore sample itself is under 0.1 ppm on average. The precision of TCCON measurements varies by site but generally remains below 0.25% (1-sigma) for individual measurements of XCO₂ under clear or partly cloudy skies.

The spatial distribution of these observation sites is as expected very unequal across the globe (Fig. 1), with the majority of them situated between 30°N and 60°N. TCCON sites provide nonetheless a good overview of different latitudes, but AirCore flights for this time period are only limited to 4 different areas.

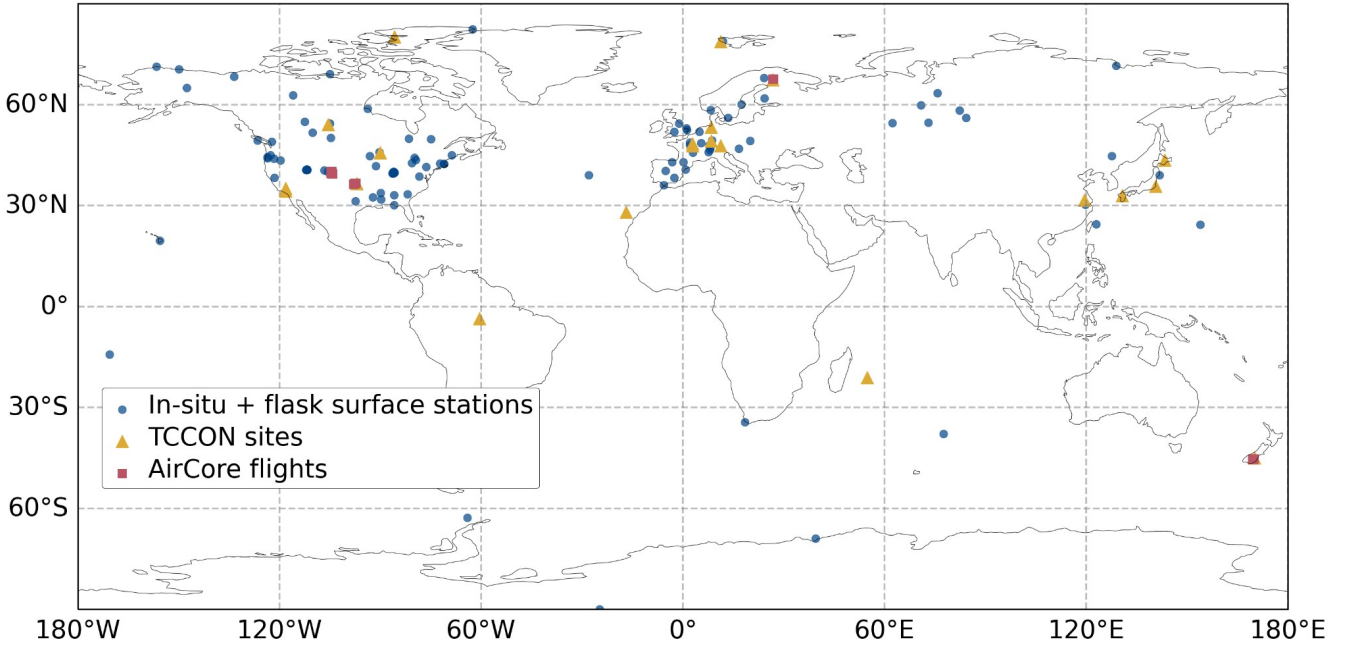


Figure 1. Map of the location of selected surface stations (blue dots), TCCON sites (yellow triangles), and AirCore flights (red squares).

2.3.2 Processing of the surface stations

To compare the results of our inversions with measurements from surface stations, we employed a curve-fitting methodology to extract the annual mean, seasonal cycle, and synoptic variability of the CO₂ mole fraction from the time-series of measurements and the model. The function used for fitting consists of a second-degree polynomial and eight harmonics (Eq. 3 below). The polynomial characterizes the background growth rate in CO₂ concentration, although this aspect is not the focus of our study due to the limited duration of our inversions. The harmonics capture the seasonal variability of CO₂ concentrations, while the synoptic variability is obtained by subtracting the fitted curve from the raw measurements or model values (Eq. 4 hereafter).

$$f(t) = p_0 + p_1 \cdot t + p_2 \cdot t^2 + \sum_{k=3}^{10} p_k \cdot \sin(2\pi kt) \quad (3)$$

$$r(t) = x(t) - f(t) \quad (4)$$

To study the seasonal cycle we quantify the correlation of the phase between model and measurements as well as the normalized peak-to-peak amplitude of the harmonics. For the synoptic variability, we look at the correlation coefficient between model and measurements and at the normalized standard deviation of the values. An example of this seasonal decomposition is presented as supplementary material in Fig. S1.

The metrics are denoted by the corresponding abbreviation of the model resolution when appropriate, for example RLR for the correlation between the low-resolution model and the measurements. When comparing a metric between the two resolutions, it is always calculated by subtracting the LR value from the HR value such as in Eq. 5:

$$\Delta R = \Delta_{\text{HR}} - \Delta_{\text{LR}} \quad (5)$$

The normalization of a metric in our case refers to the division of the model metric by the one of the observations. For the normalized peak-to-peak amplitude (NPtP) and normalized standard deviation (NSD) of the LR model for example in Eq. 6:

$$NPtP_{LR} = \frac{PtP_{LR}}{PtP_{obs}}; NSD_{LR} = \frac{SD_{LR}}{SD_{obs}} \quad (6)$$

2.3.2 Processing of the column-averaged CO₂ and vertical profiles

In evaluating the vertical profiles of CO₂ mole fractions, we employed a binning and averaging approach to organize the data from AirCore measurements and our models into 21 altitude bins of varied sizes between 500 m and the maximum altitude of 26 km. The height of each of these bins is shown together with the results of Fig. 6 in Section 3.3. They were chosen to be more refined at the altitudes with the most differences between model and measurements.

We looked at the direct values and gradients of these vertical profiles as well as the distribution of the median bias per altitude bin.

To compare our model to independent TCCON observations on the X2019 scale, we computed the column-averaged CO₂ mole fraction at each observation location and time with their respective averaging kernel and prior profile. We then computed the difference between observations and models, and in particular looked at the mean bias, correlation and normalized standard deviation (as defined in the previous subsection). In addition, we also applied the seasonal decomposition analysis described above to the TCCON observations.

2.3.3 Processing of the surface flux estimates

To study the regional distribution of the CO₂ fluxes, we divided the domain into the 22 Transcom3 regions of Gurney et al. (2002) and computed the CO₂ monthly fluxes of the two inversions in each one. From this subpartition, we also calculated annual fluxes at a global scale and a land or ocean partition.

We also compared the differences at a smaller scale by generating maps that averaged CO₂ fluxes in each cell per season, providing insights into local variations.

3 Results and discussion

3.1 Surface stations




Figure 2. Pearson correlation coefficient (a) and average normalized peak-to-peak amplitude (b) of the modeled vs. measured CO₂ mole fraction seasonal cycle for each surface station studied for the years 2015-2016. Blue circles are for the LR model and red circles are for the HR model. The stations are represented by their code in the ObsPack database. The average correlation coefficient for each resolution is in the corresponding color as a solid or dotted line in panel (a). The black dashed line in (b) corresponds to the ideal normalized peak-to-peak amplitude of 1. The stations are ordered on the abscissa by increasing latitude from -90° to +90°.

The mean correlation coefficient of the seasonal cycle across all stations studied is 0.90 for both resolutions (Fig. 2a). The average normalized peak-to-peak amplitude is 1.08 for the LR and 1.07 for the HR. The standard deviation for the normalized peak-to-peak amplitude is 0.52 for the LR and 0.42 for the HR (Fig. 2b). Both resolutions therefore capture the seasonal cycle similarly well in general, and only a few stations show large differences between the two resolutions. The HR shows a significantly lower spread of the peak-to-peak amplitude, indicating an improvement in modeling the seasonal variability.

The better-performing stations in terms of seasonal cycle correlation ($\Delta R > 0.1$) and normalized peak-to-peak amplitude ($\Delta NPtP > 0.3$) for the HR model compared to the LR model are the following ones: DEC, PV, BU, CPT and SGP, CIT, BRM, OWA, WAO, LAN, HNP. The stations that perform worse with the HR model while still capturing the seasonal cycle well in the LR model ($\Delta R < 0.1$, $R_{LR} > 0.7$ and $\Delta NPtP < 0.3$, $|NPtP_{LR} - 1| < 0.5$) are: BIR, UTSUG, UTMSA and BAO, INX06, INX07. Their locations and characteristics are presented in Table 1.

There is no general trend directly linking these results to the latitude of the studied stations.




Figure 3. Same as Fig. 2 but for the Pearson correlation coefficient (a) and the normalized standard deviation (b) of the daily average residue between our modeled and measured CO_2 mole fraction at the surface stations averaged for the years 2015-2016.

Figure 3 (a) shows that the mean synoptic variability correlation slightly improves at the higher resolution, going from 0.36 for the LR to 0.38 for the HR. The average normalized standard deviation is 1.33 for the LR model, and reduced to 1.29 for the HR model. This shows a small but significant overall improvement regarding the synoptic variability of surface stations when increasing the resolution of our model. The improvement is actually pronounced at the lower end (mean improvement of 0.03 for $R_{LR} < 0.4$) while correlations are hardly changing at the higher end (mean improvement of 0.002 for $R_{LR} > 0.4$).

The better-performing stations in terms of synoptic variability correlation ($\Delta R > 0.1$) and normalized standard deviation ($\Delta \text{NSD} > 1.0$) for the HR model compared to the LR model are the following ones: DEC, PV, BU, WAO, HNP, OMP, SGP and CIT, BRM. The stations that perform worse with the HR model while still capturing the synoptic variability well in the LR model ($\Delta R < 0.1$, $R_{LR} > 0.3$ and $\Delta \text{NSD} < 1.0$, $|\text{NSD}_{LR} - 1| < 1.0$) are CRV, INU, UTMSA and BAO. Their locations and characteristics are also presented in Table 1.

Most of the better-performing stations at the HR are coastal or next to areas with sharp elevation changes, while the worse-performing ones correspond to two cities, Salt Lake City and Indianapolis. The coastal and mountainous stations already perform better in the HR prior simulation than in the LR prior simulation (not shown), because the better coastline definition is hardly exploited in the assimilation of CO_2 column retrievals.

Station code	Type	Country	Seasonal better-performing version	Synoptic better-performing version
BAO	Urban, mountainous	USA	LR	LR
BIR	Coastal	Norway	LR	None

BRM	Mountainous	Switzerland	HR	HR
BU	Coastal, urban	USA	HR	HR
CIT	Coastal	USA	HR	HR
CPT	Coastal	South Africa	HR	None
CRV	Boreal	USA	None	LR
DEC	Coastal	Spain	HR	HR
HNP	Urban, lake	Canada	HR	HR
INU	Boreal	Canada	None	LR
INX06	Urban	USA	LR	None
INX07	Urban	USA	LR	None
LAN	Coastal, mountainous	China	HR	None
OMP	Coastal, mountainous	USA	None	HR
OWA	Coastal, mountainous	USA	HR	None
PV	Coastal	USA	HR	HR
SGP	Plains	USA	HR	HR
UTMSA	Urban	USA	LR	LR
UTSUG	Urban	USA	LR	None
WAO	Coastal, mountainous	UK	HR	HR

Table 1. Notable Stations Identified by Seasonal and Synoptic Variability Performance. A station is identified as better performing for a certain resolution if the difference in metric between the resolutions is superior to a threshold as defined in Section 3.1. The metrics for the seasonal cycle are the correlation and normalized peak-to-peak amplitude. For the synoptical variability it is the correlation and normalized standard deviation.

3.2 TCCON observations




Figure 4. Correlation (a) and normalized standard deviation (b) of the difference between the model XCO_2 and remotely-sensed XCO_2 from TCCON stations averaged over the years 2015-2016 for each station and then averaged across the 25 stations. Blue circles are for the LR model and red circles are for the HR model. The average correlation and normalized standard deviation for each resolution are in the corresponding color as a solid or dotted line in panels (a,b). The black dashed line in (b) corresponds to the ideal normalized standard deviation of 1. The stations are ordered on the abscissa by increasing latitudes. The y-axis on panel (b) is in log scale.

When comparing XCO_2 between the final state of our inversion and independent observations from TCCON, we see that the mean difference between the model and observations is almost identical for both resolutions, at 0.06 ppm for the LR and 0.08 ppm for the HR (not shown). Figure 4 shows that the average correlation is 0.88 for the LR and 0.89 for the HR. The average normalized standard deviation is 0.53 for both resolutions. When looking at the behavior of individual stations, the result is very different, with both the general bias and normalized standard deviation varying widely for different stations, without any obvious link with the station location. However, both resolutions behave similarly to each other at each station, with the worst-performing stations being identical for both resolutions. The two urban stations of Hefei and Tsukuba show a better correlation at HR, but this improvement is small and to be contrasted with the relatively lower performance of some urban in-situ stations at HR as shown in Section 3.1. These two TCCON stations can therefore not be taken as a general trend showing a better modelization of urban stations by the HR model.

When analyzing the seasonal fit of the observations and model at TCCON sites in Fig. 5 we see that all stations perform well in terms of correlation of the seasonal cycle, with both resolutions having a mean correlation of 0.92. They also perform almost identically regarding the modelization of the peak-to-peak amplitude of the seasonal cycles.

The simulation of column-averaged CO_2 is in principle not as sensitive to resolution increase of the transport model as for surface CO_2 (Rayner and O'Brien, 2001) and this can explain the marginal difference

between the resolutions with respect to TCCON observations.




Figure 5. Same as Fig. 4 but for the Pearson correlation coefficient (a) and the normalized peak-to-peak amplitude (b) of the fitted seasonal cycle between our modeled and measured XCO₂ mole fraction at the TCCON stations averaged for the years 2015-2016.

The difference in bias and standard deviation between the two resolutions compared to already assimilated OCO-2 retrievals is negligible at the global scale (mean bias of 0.05 ppm and standard deviation of 0.84 ppm for both resolutions).

3.3 Vertical profiles

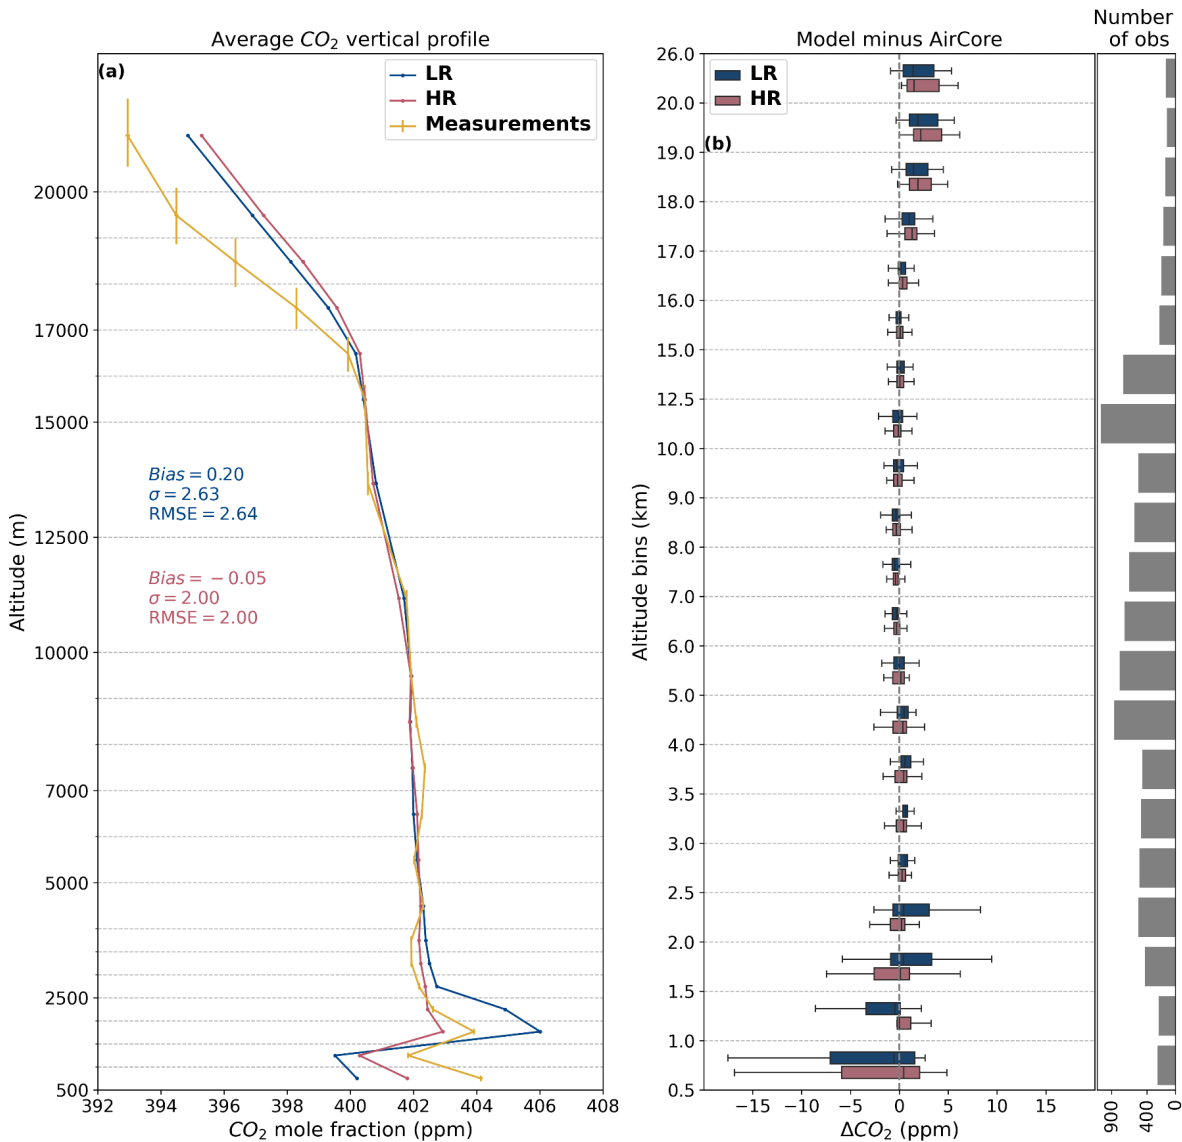


Figure 6. Panel (a) shows the CO₂ mole fraction vertical profile in ppm for the two resolutions of the model (blue for LR, red for HR) and all valid AirCore sample measurements (yellow). The lines were generated by averaging the data over all altitude bins. Error bars of the measurements correspond to the altitude determination uncertainty of the sample and to the uncertainty of the measurement itself. The dotted grey horizontal lines show the altitude of each bin. The values of the bias, standard deviation and root-mean-square deviation of the difference between the raw data of the models and measurements are presented for each resolution in their respective color (blue for LR and red for HR). Panel (b) shows the median bias (model minus measurements) and its probability distribution

over each altitude bin averaged for all valid AirCore sample measurements. The number of measurements in each altitude bin is indicated on the right.

We utilized AirCore flight data to compare the CO₂ mole fractions of our model with measurement data, obtaining vertical profiles extending to the low stratosphere. This analysis aimed to investigate the impact of increasing resolution on vertical transport. The measurements were limited in latitudes and the results may be different in the tropics, the majority of the measurements coming from conterminous United States (see Fig. 1).

As depicted in Fig. 6 (a), under 3 km, in and just above the boundary layer, the HR model performs better and shows a better agreement with measurements. When looking at the probability distribution of the bias at this altitude (b) we see that the HR model has a lower spread than the LR model for most altitude bins, indicating a better representation of the boundary layer. After 3 km, both resolutions of the model exhibit good agreement with measurements up to around 16 km. Beyond that, up to 22 km, both resolutions differ from measurements, showing a positive bias. This leads to a lower general bias for the HR model compared to measurements (-0.05 ppm vs 0.20 ppm) and a lower spread of the difference between model and measurements (standard deviation of 2.00 ppm vs 2.63 ppm).

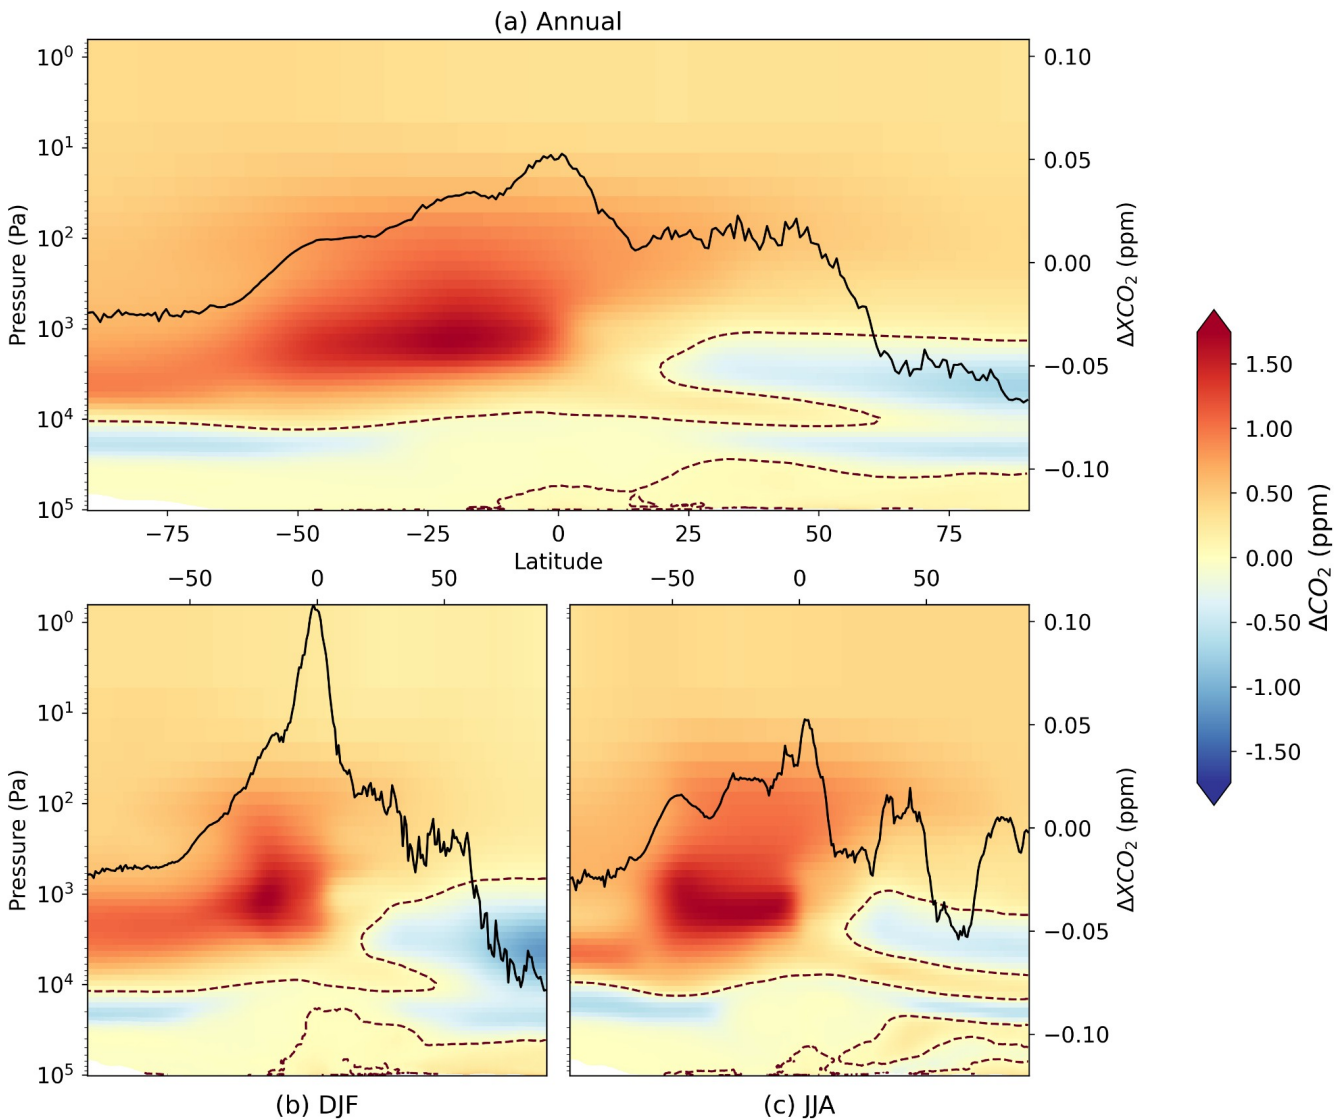


Figure 7. Difference in CO₂ mole fraction in ppm between the HR and LR models after inversion (HR-LR), averaged per month over the two years and per longitude band. The results are then averaged again either annually (a) or per season: DJF (b) and JJA (c). The data of the LR model was interpolated on the latitudes of the HR model before computing the difference. The dark brown dashed contour line shows where the value of the difference equals zero. In each panel, the black line corresponds to the difference in XCO₂ mole fraction in ppm between the HR and LR models averaged and interpolated in time and space in the same manner. The scale of this difference (HR-LR) in ppm is on the right y-axis.

When looking at the time-averaged zonal vertical profiles of CO₂ mole fraction, we can see that the distribution is different between the resolutions and is on the order of -0.7 to +1.7 ppm (Fig. 7). These variations vary both in latitude and in altitude, and the previous comparison to AirCore data only gave a limited view into these differences. The HR model shows a higher concentration of CO₂ in the upper atmosphere in general, but at these high altitudes the total mass of CO₂ is very low. The black lines in Fig. 7 corresponding to the zonal mean of the difference in XCO₂ between the resolutions show that this vertically integrated difference behaves very differently depending on the latitude and season. The

difference in XCO_2 is most important in the -25° to 0° latitude band, particularly in winter. For other latitudes, the concentration difference in the high atmosphere is largely compensated by an opposite difference at lower altitudes. This suggests that in these cases the difference in vertical profile of atmospheric CO_2 is mostly driven by a higher vertical transport speed in our HR model, whereas around the 0° latitude band, the higher atmospheric CO_2 concentration is present irrespective of altitude.

Figure 8 shows the 2D spatial distribution of the difference in XCO_2 between the two resolutions. The difference remains small and mostly under 0.2 ppm, but a pattern still emerges. All year round, the HR model has a higher XCO_2 mole fraction over both land and ocean in the tropics than the LR model but a lower one over land in the high northern latitudes. The overall sign opposition between northern extratropical lands and the extratropical Pacific ocean may be favored by the exclusion of ocean satellite retrievals from the assimilation system (see Section 2.2)

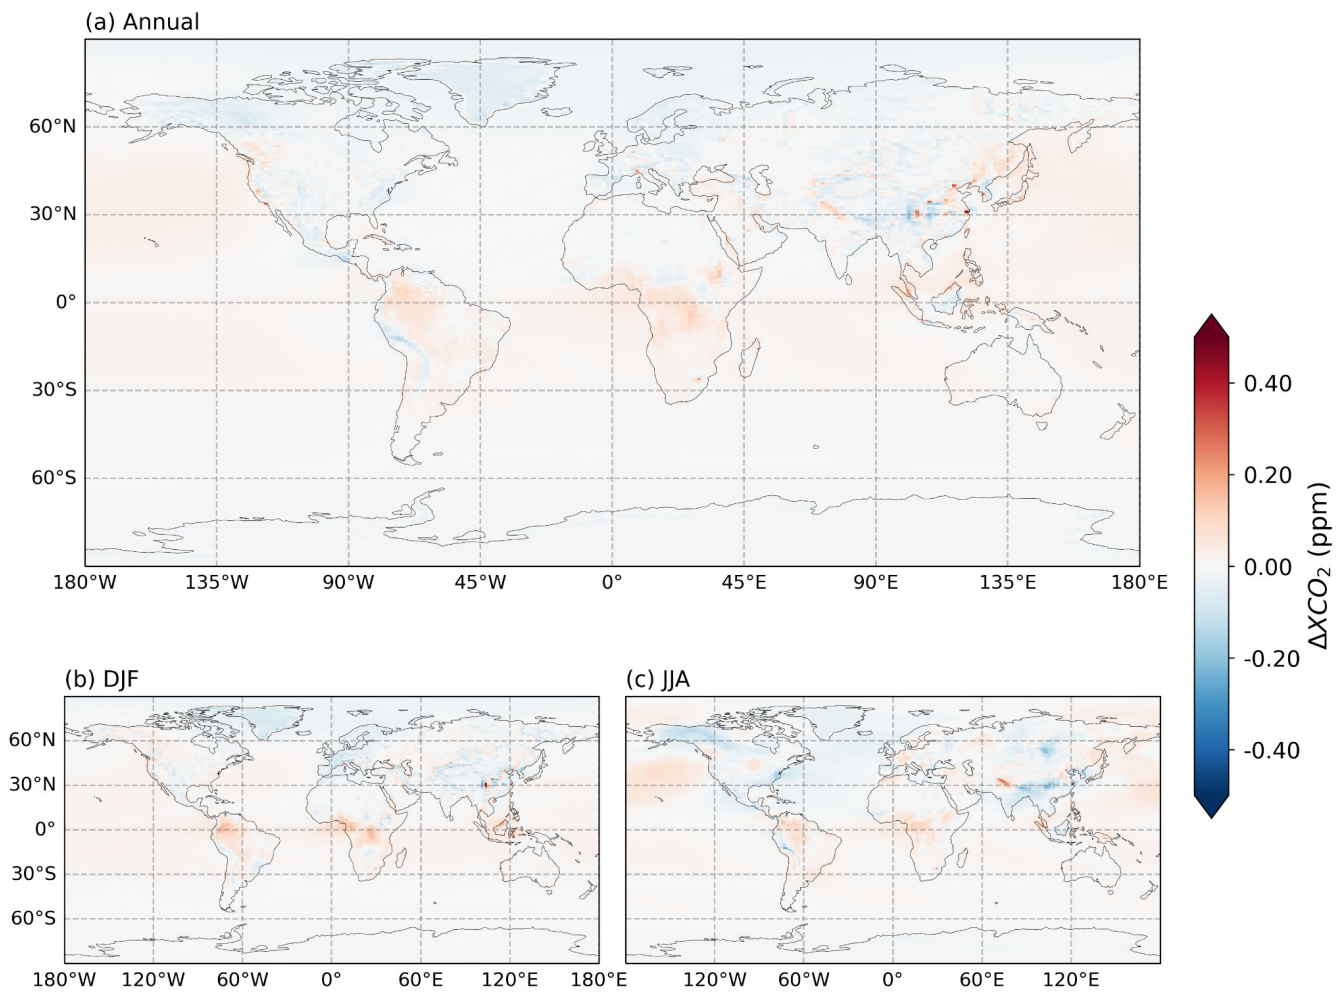


Figure 8. Maps of the difference in XCO_2 mole fraction in ppm between the HR and LR models (HR-LR) after the inversion, averaged per year and season over the 2015-2016 period. The data of the LR model was interpolated on the longitudes and latitudes of the HR model before computing the difference.

3.4 Regional fluxes

Table 2 shows the global estimates of natural carbon fluxes after inversion for our model at both resolutions. The results are very similar to each other for both years at the global scale, with the HR model giving a slightly stronger sink. This is due to a stronger land sink at the HR, which is not fully compensated by the relatively weaker ocean sink. This difference in the total land flux is not equally distributed across space.

Global flux estimates are in line with estimates from atmospheric inversion results using the v9 OCO-2 retrievals for 2015 (Peiro et al., 2022), but this is true at the global scale because the land sink is higher while the ocean sink is lower. This trend is similar in 2016, but the lower ocean sink makes the global sink lower than the range of v9 OCO-2 estimates (Peiro et al., 2022).

Year	Model	Land flux (GtC yr ⁻¹)	Ocean flux (GtC yr ⁻¹)	Global flux (GtC yr ⁻¹)
2015	LR	-1.29	-1.94	-3.22
	HR	-1.54	-1.77	-3.31
2016	LR	-1.53	-2.18	-3.71
	HR	-1.87	-2.00	-3.88

Table 2. Estimation of the natural carbon fluxes per year after inversion for each model resolution, at the global scale and partitioned by land and ocean.




Figure 9. Total annual surface emissions minus the fossil fuel emissions for LR and HR (in blue and red respectively) in GtC for each Transcom3 region, for the year 2015 with squares above the black dotted line, and with circles for the year 2016 below the line.

Figure 9 shows the annual net surface flux in GtC per year minus the fossil fuel emissions per Transcom3 region for each year of our inversion and both resolutions. This information, combined with some monthly estimates of CO₂ fluxes from Fig. 10 inform us about when and where surface fluxes estimated by the inversions differ depending on the corresponding model resolution.

A few Transcom3 regions exhibit notable differences in CO₂ flux dynamics (Fig. 10). This is particularly the case for North American boreal forests, where the HR model suggests substantially more sink in both

years. This difference in regional carbon flux between the two models is not paralleled by notable discrepancies in the seasonal cycle of CO₂ concentrations compared to independent measurements from surface stations. The CRV and INU stations situated in this region only perform worse with the HR model in terms of synoptical variability, not seasonality (as noted in Section 3.1). However, the bias of the model XCO₂ with the assimilated OCO-2 retrievals in this region is significantly smaller ($t = 5.3$, $p < 0.001$) for the HR model ($M = 0.19$, $SD = 0.97$) compared to the LR one ($M = 0.35$, $SD = 1.03$). This result indicates a more efficient assimilation of the satellite data in this region. For the South American Tropical and North American Temperate regions, the HR model has a bigger carbon sink, particularly in the year 2016. This is in line with the higher global land sink of the HR model. The Eurasian Boreal region on the opposite has higher emissions in the HR model. Fig. 10 (d) shows that this is only limited to the beginning and end of a year. And given the large size of this region, the overall impact at the global scale is minor.




Figure 10. Monthly averaged surface flux minus the fossil fuel emissions for LR and HR model in GtC per month (blue and red respectively), for 2015 and 2016 (solid lines and dashed lines respectively) in Transcom3 regions North American Boreal (a), South American Tropical (b), North American Temperate (c), and Eurasian Boreal (d). These regions show the greatest relative difference in estimated annual flux between the two resolutions of our model.

3.5 Local fluxes

When looking at fluxes at the local scale, we can directly see the benefit of the high resolution with respect to coastal definition, in particular in areas with complex coastlines. Figure 11 shows maps of the increments of the surface fluxes, i.e. the correction of the prior fluxes by the posterior ones, averaged for winter and summer between 2015 and 2016. Some regional scale patterns discussed in section 3.5 can be immediately seen, such as the higher summer sink of carbon for the HR model in boreal North America. The general patterns of surface fluxes for the HR model are similar to the LR model but provide much more spatial details.

The surface stations that the HR model fits better and, therefore that benefit the most from the increased resolution as discussed in Section 3.1, are situated either in continental North America, near large population centers with complex orography, or near the coast (listed in Table 1 and visible in Fig. 11). This indicates that the improvement we see is not primarily caused by fine-scale changes in the seasonal flux patterns but more so by the improved orography and wind fields used to drive the model.

The zoom of Figure 12 exemplifies the improvement gained by the increase in resolution around Southeast China and Taiwan. The Taiwan Strait at HR is represented with some pure marine pixels in contrast to LR. Conversely, the LAN station in the North East of the figure is in a mixed cell at LR with both land and sea surfaces, but is clearly inland at HR. Such a behavior can be seen across the globe in particular around large islands or straits. This benefit from the HR model does not come through a better assimilation of the OCO-2 data, but is inherent to the resolution of the transport model itself.




Figure 11. Surface flux increments between the prior and posterior state of the inversion for the LR

(a,b) and HR (c,d) versions, in $\text{kg}/\text{m}^2/\text{month}$. The fluxes are averaged over the corresponding months for the 2 years of inversion. December, January and February (a,c), June, July and August (b,c). The dots correspond to the surface stations that each resolution improve the most compared to the other one in terms of seasonal cycle and synoptic variability, as listed in Table 1 (blue for stations performing better in LR, red in HR).




Figure 12. Total monthly surface flux including fossil fuel emissions averaged over the period 2015-2016 for the LR (a) and HR (b) versions, in $\text{kg}/\text{m}^2/\text{month}$, zoomed around the area near the station LAN in China. The lines show the edge of the cells of each model, highlighting the difference in resolution, particularly along the coastline.

4 Conclusion

We successfully increased the resolution of the CAMS/LSCE inversion system, tripling the number of global grid points and reaching a global resolution of 0.7° latitude and 1.4° longitude. This was made possible thanks to recent developments in the model, allowing it to run on GPUs and limiting the necessary higher computational cost than the previous resolution to twice without increasing the number of devices. While this study focused on an inversion over two years and only assimilating OCO-2 data over land, larger and longer-lasting inversions are now possible and will be part of future operational work within CAMS.

As seen in the previous sections, the increase in resolution of our inverse model leads to a small but significant overall improvement in the representation of atmospheric CO_2 compared to independent measurements from surface stations, particularly at the synoptic time scale. The stations where the benefit of the new resolution is seen the most were situated primarily near coasts or large cities. This gain was primarily due to the resolution increase of the transport model, leading to a better orography and coastal definition. This is promising for the quality of future surface-driven inversions run at the new resolution.

The vertical profiles of CO_2 concentration are different between the two resolutions when compared to AirCore measurements, particularly for altitudes above 22 km. The HR model performs better under 3 km

which leads to a lower general bias and spread of the difference with measurements. This difference can also be seen when looking at zonal averages of the vertical profile of CO₂. This disparity between resolutions is however not evidenced when looking at XCO₂ globally, whether when comparing the final inversion product to already assimilated OCO-2 observations or to independent TCCON observations.

The global and regional estimates of the natural fluxes for the years 2015 and 2016 are similar for our two resolutions but the HR model shows a consistently higher land sink and lower ocean sink than the LR model (without assimilating satellite ocean retrievals). The largest regional difference is a higher natural sink in North America for the HR model during the year 2016. Both inversions offer valid options for global and regional estimates of natural carbon fluxes and we cannot directly demonstrate the expected superiority of the higher resolution ones.

Further enhancement in horizontal resolution holds the potential for increased benefits in atmospheric transport, with a critical threshold being the attainment of full cloud resolution rather than relying on subgrid parameterization (Schneider et al., 2017). Upcoming missions such as the Carbon Dioxide Monitoring (CO2M) and the Global Observing SATellite for Greenhouse gases and Water cycle (GOSAT-GW) will use wide-swath sensors which will provide a much higher observation density. How well higher resolution inverse models will be able to leverage this increase in observation density is still not clear, the demonstrated better assimilation of OCO-2 data by our HR model being only restricted to a limited area in the North American Boreal region. Furthermore, conventional latitude-longitude grids may encounter computing bottlenecks when scaling up in resolution, particularly due to clustering issues at the poles. The proposed strategy for the CAMS/LSCE inversion system to address this challenge involves adopting a new dynamical core operating on an icosahedral grid (Dubos et al., 2015). Ongoing development efforts aim to bring such a core in the CAMS/LSCE inversion system in order to reach sub-degree resolutions.

Acknowledgments

This work was granted access to the HPC resources of CCRT under the allocation CEA/DRF, and of TGCC under the allocation A0130102201 made by GENCI. It was funded by the Copernicus Atmosphere Monitoring Service, implemented by ECMWF on behalf of the European Commission (Grant: CAMS2 55). The XCO₂ retrieval data used for the inversion was produced by the OCO-2 project at the Jet Propulsion Laboratory, California Institute of Technology, and obtained from the OCO-2 data archive maintained at the NASA Goddard Earth Science Data and Information Services Center. We gratefully acknowledge the many people who contributed atmospheric observations. The TCCON data were obtained from the TCCON Data Archive hosted by CaltechDATA at <https://tccodata.org>. The AirCore data were obtained from NOAA at <https://gml.noaa.gov/ccgg/arc/?id=144>. We thank Bianca Baier for her help in using this data and representing it.

The authors acknowledge the PIs and contributors related to the operations in the compilations of Obspack CO₂ dataset (obspack_co2_1_GLOBALVIEWplus v8.0_2022-08-27) and ICOS network. The contributions from the following people and institutions are thankfully acknowledged. A. Cox, B. Paplawsky, E. Gloor, E. Kort, F. Apadula, M. Kumar Sha, M. De Mazière, P. Trisolino, S. Walker, S. Piper and T. Biermann; A. Giorgio di Sarra and S. Piacentino (ENEA); A. Vermeulen (LU); A. Manning (METOFFICE); A. Beyersdorf (CSUSB); A. Zahn, F. Obersteiner, H. Boenisch and T. Gehrlein (KIT/IMK-IFU); A. Manning, G. Forster and R. A. F. de Souza (UEA); A. Karion (NIST); A. Hoheisel, I. Levin, J. Della Coletta and S. Hammer (UHEI-IUP); A. Leskinen, J. Hatakka, K. Lehtinen, O. Peltola, T. Laurila and T. Aalto (FMI); A. Hensen, A. Frumau, D. van Dinter and P. van den Bulk (ECN); A. Andrews, B. Baier, C. Sweeney, E. Dlugokencky, E. Hintsa, F. Moore, J. Peischl, J. B. Miller, J. Mund, K. McKain, K. Aikin, K. N. Schuldt, K. Thoning, P. Tans, S. Montzka and X. Lan (NOAA); A. Jordan, C.

Gerbig, H. Moossen, J. Lavric, M. Heimann, S. Zaehle and W. A. Brand (MPI-BGC); A. Colomb and J. Marc Pichon (OPGC); B. Scheeren, H. Meijer and H. Chen (RUG); B. Law and C. Hanson (OSU); B. Munger, M. Sargent and S. Wofsy (HU); B. Viner (SRNL); B. Stephens (NCAR); C. Labuschagne (SAWS); C. Lund Myhre, C. René Lunder, K. Tørseth, O. Hermanssen and S. Matthew Platt (NILU); C. Couret (UBA); C. E. Miller (NASA-JPL); C. Plass-Duelmer, C. Plass-Duelmer, D. Kubistin, M. Schumacher, M. Lindauer and T. Kneuer (DWD); C. D. Sloop and S. Prinzivalli (EN); D. Jaffe (UofWA); D. Heltai (RSE); D. Bowling, J. Lin and L. Mitchell (U-ATAQ); D. H.Y. Lam and O. S.M. Lee (HKO); D. Munro (NOAA - CIRES); D. Young, J. Pitt and S. O'Doherty (UNIVBRIS); D. Worthy (ECCC); E. Kozlova (CEDA); E. Cuevas, E. Reyes-Sanchez and P. P. Rivas (AEMET); E. Morgan, J. Kim, L. Merchant, R. Keeling, R. Weiss and S. Clark (SIO); F. Meinhardt (UBA-SCHAU); G. Vitkova, K. Kominkova and M. V. Marek (CAS); G. Chen and M. Shook (NASA LaRC); G. A. Martins (FDB); G. Manca and P. Bergamaschi (JRC); G. Brailsford and S. Nichol (NIWA); H. Riris, J. Brice Abshire and S. Randolph Kawa (NASA-GSFC); H. Matsueda (MRI); I. T. Lujikx (WU AND ICOS-CP); I. Lehner and M. Heliasz (LUND-CEC); I. Mammarella, J. Levula, P. Kolari and P. Keronen (UHELS); J. W. Elkins (HATS); J. Necki, L. Chmura, M. Galkowski and M. Zimnoch (AGH); J. Müller-Williams (HPB); J. Turnbull (GNS); J. Lee (UofME); J. Morgui, R. Curcoll and S. Climadat (ICTA-UAB); J. P. DiGangi (NASA-LaRC); J. Holst and M. Mölder (LUND-NATEKO); K. Saito (JMA); K. Davis, N. Miles, S. Richardson and T. Lauvaux (PSU); L. Lotte Sørensen (AU); L. V. Gatti (INPE); L. Emmenegger (EMPA); L. Haszpra (RCAES); M. Delmotte, M. Schmidt, M. Ramonet, M. Lopez and V. Kazan (LSCE); M. L. Fischer and M. Torn (LBNL); M. Leuenberger (KUP); M. Steinbacher (empa); M. Sasakawa, T. Machida and Y. Niwa (NIES); O. Laurent (ICOS-ATC); P. Cristofanelli (CNR-ISAC); P. Krummel, R. Langenfelds and Z. Loh (CSIRO); P. Shepson (PU); P. Smith (SLU); S. Newman (CIT); S. C. Biraud (LBNL-ARM); S. Morimoto and S. Aoki (TU); S. Fang (CMA); S. De Wekker (UofVA); S. Conil (Andra); T. Schuck (IAU); T. Griffis (uminn); V. Ivakhov (MGO)

Conflict of Interest

The authors declare that they have no conflict of interest.

Open Research

Data Availability Statement

The LMDZ off-line transport model v3.1 is publicly available from <https://doi.org/10.5281/zenodo.7324039> (Chevallier, 2022) under the Creative Commons Attribution 4.0 International license.

The inverse system in Python is available as part of the CIF at <https://git.nilu.no/VERIFY/CIF>.

References

- Agustí-Panareda, A., Diamantakis, M., Massart, S., Chevallier, F., Muñoz-Sabater, J., Barré, J., Curcoll, R., Engelen, R., Langerock, B., Law, R. M., Loh, Z., Morguí, J. A., Parrington, M., Peuch, V.-H., Ramonet, M., Roehl, C., Vermeulen, A. T., Warneke, T., & Wunch, D. (2019). Modelling CO₂ weather – why horizontal resolution matters. *Atmospheric Chemistry and Physics*, 19(11), 7347–7376. <https://doi.org/10.5194/acp-19-7347-2019>
- Baier, B., Sweeney, C., Newberger, T., Higgs, J., Wolter, S., & NOAA Global Monitoring Laboratory. (2021). NOAA AirCore atmospheric sampling system profiles (Version 20230831) [Data set]. NOAA GML. <https://doi.org/10.15138/6AV0-MY81>.

- Basu, S., Baker, D. F., Chevallier, F., Patra, P. K., Liu, J., & Miller, J. B. (2018). The impact of transport model differences on CO₂ surface flux estimates from OCO-2 retrievals of column average CO₂. *Atmospheric Chemistry and Physics*, 18(10), 7189–7215. <https://doi.org/10.5194/acp-18-7189-2018>
- Berchet, A., Sollum, E., Thompson, R. L., Pison, I., Thanwerdas, J., Broquet, G., Chevallier, F., Aalto, T., Berchet, A., Bergamaschi, P., Brunner, D., Engelen, R., Fortems-Cheiney, A., Gerbig, C., Groot Zwaaftink, C. D., Haussaire, J.-M., Henne, S., Houweling, S., Karstens, U., ... Zhao, Y. (2021). The Community Inversion Framework v1.0: A unified system for atmospheric inversion studies. *Geoscientific Model Development*, 14(8), 5331–5354. <https://doi.org/10.5194/gmd-14-5331-2021>
- Buendia, E., Tanabe, K., Kranjc, A., Jamsranjav, B., Fukuda, M., Ngarize, S., Osako, A., Pyrozhenko, Y., Shermanau, P., & Federici, S. (2019). 2019 Refinement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories.
- Buschmann, M., Petri, C., Palm, M., Warneke, T., & Notholt, J. (2022). TCCON data from Ny-Ålesund, Svalbard (NO), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tcon.ggg2020.nyalesund01.R0>
- Liu, C., Wang, W., Sun, Y., & Shan, C. (2023). TCCON data from Hefei (PRC), Release GGG2020.R1 (Version R1) [Data set]. CaltechDATA. <https://doi.org/10.14291/tcon.ggg2020.hefei01.R1>
- Chau, T. T. T., Gehlen, M., & Chevallier, F. (2022). A seamless ensemble-based reconstruction of surface ocean pCO₂ and air–sea CO₂ fluxes over the global coastal and open oceans. *Biogeosciences*, 19(4), 1087–1109. <https://doi.org/10.5194/bg-19-1087-2022>
- Chevallier, F., Fisher, M., Peylin, P., Serrar, S., Bousquet, P., Bréon, F.-M., Chédin, A., & Ciais, P. (2005). Inferring CO₂ sources and sinks from satellite observations: Method and application to TOVS data. *Journal of Geophysical Research: Atmospheres*, 110(D24). <https://doi.org/10.1029/2005JD006390>
- Chevallier, F. (2013). On the parallelization of atmospheric inversions of CO₂ surface fluxes within a variational framework. *Geoscientific Model Development*, 6(3), 783–790. <https://doi.org/10.5194/gmd-6-783-2013>
- Chevallier, F. (2021). Fluxes of Carbon Dioxide From Managed Ecosystems Estimated by National Inventories Compared to Atmospheric Inverse Modeling. *Geophysical Research Letters*, 48(15), e2021GL093565. <https://doi.org/10.1029/2021GL093565>
- Chevallier, F. (2022). LMDz transport model (3.1). [Software] Zenodo. <https://doi.org/10.5281/zenodo.7324039>
- Chevallier, F., Lloret, Z., Cozic, A., Takache, S., & Remaud, M. (2023). Toward High-Resolution Global Atmospheric Inverse Modeling Using Graphics Accelerators. *Geophysical Research Letters*, 50(5), e2022GL102135. <https://doi.org/10.1029/2022GL102135>
- [Courant, R., Friedrichs, K., & Lewy, H. \(1928\). Über die partiellen Differenzgleichungen der mathematischen Physik. *Mathematische Annalen*, 100\(1\), 32–74. https://doi.org/10.1007/BF01448839](https://doi.org/10.1007/BF01448839)
- Crotwell, A., Lee, H., & Steinbacher, M. (2020). “20th WMO/IAEA Meeting on Carbon Dioxide, Other Greenhouse Gases and Related Measurement Techniques (GGMT-2019),” GAW REPORT No. 255. <https://library.wmo.int>

- Deutscher, N. M., Griffith, D. W. T., Paton-Walsh, C., Velazco, V. A., Wennberg, P. O., Blavier, J.-F., Washenfelder, R. A., Yavin, Y., Keppel-Aleks, G., Toon, G. C., Jones, N. B., Kettlewell, G. C., Connor, B. J., Macatangay, R. C., Wunch, D., Roehl, C., & Bryant, G. W. (2023). TCCON data from Darwin (AU), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.darwin01.R0>
- Dubey, M. K., Henderson, B. G., Allen, N. T., Blavier, J.-F., Roehl, C. M., & Wunch, D. (2022). TCCON data from Manaus (BR), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.manaus01.R0>
- Dubos, T., Dubey, S., Tort, M., Mittal, R., Meurdesoif, Y., & Hourdin, F. (2015). DYNAMICO-1.0, an icosahedral hydrostatic dynamical core designed for consistency and versatility. *Geoscientific Model Development*, 8(10), 3131–3150. <https://doi.org/10.5194/gmd-8-3131-2015>
- Eggleston, H. S., Buendia, L., Miwa, K., Ngara, T., & Tanabe, K. (2006). 2006 IPCC Guidelines for National Greenhouse Gas Inventories. <https://www.osti.gov/etdeweb/biblio/20880391>
- Fisher, M. (1998). Minimization algorithms for variational data assimilation. In *Proceedings of seminar on recent developments in numerical methods for atmospheric modelling*. 7–11 September 1998 (pp. 364–385). ECMWF. Retrieved from <https://www.ecmwf.int/node/9400>
- Fortems-Cheiney, A., Pison, I., Broquet, G., Dufour, G., Berchet, A., Potier, E., Coman, A., Siour, G., & Costantino, L. (2021). Variational regional inverse modeling of reactive species emissions with PYVAR-CHIMERE-v2019. *Geoscientific Model Development*, 14(5), 2939–2957. <https://doi.org/10.5194/gmd-14-2939-2021>
- Friedlingstein, P., O’Sullivan, M., Jones, M. W., Andrew, R. M., Gregor, L., Hauck, J., Le Quéré, C., Luijkx, I. T., Olsen, A., Peters, G. P., Peters, W., Pongratz, J., Schwingshackl, C., Sitch, S., Canadell, J. G., Ciais, P., Jackson, R. B., Alin, S. R., Alkama, R., ... Zheng, B. (2022). Global Carbon Budget 2022. *Earth System Science Data*, 14(11), 4811–4900. <https://doi.org/10.5194/essd-14-4811-2022>
- Friedlingstein, P., O’Sullivan, M., Jones, M. W., Andrew, R. M., Bakker, D. C. E., Hauck, J., Landschützer, P., Le Quéré, C., Luijkx, I. T., Peters, G. P., Peters, W., Pongratz, J., Schwingshackl, C., Sitch, S., Canadell, J. G., Ciais, P., Jackson, R. B., Alin, S. R., Anthoni, P., ... Zheng, B. (2023). Global Carbon Budget 2023. *Earth System Science Data*, 15(12), 5301–5369. <https://doi.org/10.5194/essd-15-5301-2023>
- Gaubert, B., Stephens, B. B., Basu, S., Chevallier, F., Deng, F., Kort, E. A., Patra, P. K., Peters, W., Rödenbeck, C., Saeki, T., Schimel, D., Van der Laan-Luijkx, I., Wofsy, S., & Yin, Y. (2019). Global atmospheric CO₂ inverse models converging on neutral tropical land exchange, but disagreeing on fossil fuel and atmospheric growth rate. *Biogeosciences*, 16(1), 117–134. <https://doi.org/10.5194/bg-16-117-2019>
- Geels, C., Gloor, M., Ciais, P., Bousquet, P., Peylin, P., Vermeulen, A. T., Dargaville, R., Aalto, T., Brandt, J., Christensen, J. H., Frohn, L. M., Haszpra, L., Karstens, U., Rödenbeck, C., Ramonet, M., Carboni, G., & Santaguida, R. (2007). Comparing atmospheric transport models for future regional inversions over Europe – Part 1: Mapping the atmospheric CO₂ signals. *Atmospheric Chemistry and Physics*, 7(13), 3461–3479. <https://doi.org/10.5194/acp-7-3461-2007>
- Gurney, K. R., Law, R. M., Denning, A. S., Rayner, P. J., Baker, D., Bousquet, P., Bruhwiler, L., Chen, Y.-H., Ciais, P., Fan, S., Fung, I. Y., Gloor, M., Heimann, M., Higuchi, K., John, J., Maki, T., Maksyutov, S.,

Masarie, K., Peylin, P., ... Yuen, C.-W. (2002). Towards robust regional estimates of CO₂ sources and sinks using atmospheric transport models. *Nature*, 415(6872), Article 6872.

<https://doi.org/10.1038/415626a>

Hall, B. D., Crotwell, A. M., Kitzis, D. R., Mefford, T., Miller, B. R., Schibig, M. F., & Tans, P. P. (2021). Revision of the World Meteorological Organization Global Atmosphere Watch (WMO/GAW) CO₂ calibration scale. *Atmospheric Measurement Techniques*, 14(4), 3015–3032. <https://doi.org/10.5194/amt-14-3015-2021>

Hazan, L., Tarniewicz, J., Ramonet, M., Laurent, O., & Abbaris, A. (2016). Automatic processing of atmospheric CO₂ and CH₄ mole fractions at the ICOS Atmosphere Thematic Centre. *Atmospheric Measurement Techniques*, 9(9), 4719–4736. <https://doi.org/10.5194/amt-9-4719-2016>

Hedelius, J. K., Feng, S., Roehl, C. M., Wunch, D., Hillyard, P. W., Podolske, J. R., Iraci, L. T., Patarasuk, R., Rao, P., O’Keefe, D., Gurney, K. R., Lauvaux, T., & Wennberg, P. O. (2017). Emissions and topographic effects on column CO₂ () variations, with a focus on the Southern California Megacity. *Journal of Geophysical Research: Atmospheres*, 122(13), 7200–7215.

<https://doi.org/10.1002/2017JD026455>

Hourdin, F., Rio, C., Grandpeix, J.-Y., Madeleine, J.-B., Cheruy, F., Rochetin, N., Jam, A., Musat, I., Idelkadi, A., Fairhead, L., Foujols, M.-A., Mellul, L., Traore, A.-K., Dufresne, J.-L., Boucher, O., Lefebvre, M.-P., Millour, E., Vignon, E., Jouhaud, J., ... Ghattas, J. (2020). LMDZ6A: The Atmospheric Component of the IPSL Climate Model With Improved and Better Tuned Physics. *Journal of Advances in Modeling Earth Systems*, 12(7), e2019MS001892. <https://doi.org/10.1029/2019MS001892>

ICOS RI, Bergamaschi, P., Colomb, A., De Mazière, M., Emmenegger, L., Kubistin, D., Lehner, I., Lehtinen, K., Lund Myhre, C., Marek, M., Platt, S.M., Plaß-Dülmer, C., Schmidt, M., Apadula, F., Arnold, S., Blanc, P.-E., Brunner, D., Chen, H., Chmura, L., Conil, S., Couret, C., Cristofanelli, P., Delmotte, M., Forster, G., Frumau, A., Gheusi, F., Hammer, S., Haszpra, L., Heliasz, M., Henne, S., Hoheisel, A., Kneuer, T., Laurila, T., Leskinen, A., Leuenberger, M., Levin, I., Lindauer, M., Lopez, M., Lunder, C., Mammarella, I., Manca, G., Manning, A., Marklund, P., Martin, D., Meinhardt, F., Müller-Williams, J., Necki, J., O’Doherty, S., Ottosson-Löfvenius, M., Philippon, C., Piacentino, S., Pitt, J., Ramonet, M., Rivas-Soriano, P., Scheeren, B., Schumacher, M., Sha, M.K., Spain, G., Steinbacher, M., Sørensen, L.L., Vermeulen, A., Vítková, G., Xueref-Remy, I., di Sarra, A., Conen, F., Kazan, V., Roulet, Y.-A., Biermann, T., Heltai, D., Hensen, A., Hermansen, O., Komínková, K., Laurent, O., Levula, J., Pichon, J.-M., Smith, P., Stanley, K., Trisolino, P., ICOS Carbon Portal, ICOS Atmosphere Thematic Centre, ICOS Flask And Calibration Laboratory, ICOS Central Radiocarbon Laboratory, 2023. European Obspack compilation of atmospheric carbon dioxide data from ICOS and non-ICOS European stations for the period 1972-2023; obspack_co2_466_GLOBALVIEWplus_v8.0_2023-04-26. <https://doi.org/10.18160/CEC4-CAGK>

IPCC, Eggleston, H S, Buendia, L, Miwa, K, Ngara, T, and Tanabe, K. (2006). IPCC Guidelines for National Greenhouse Gas Inventories, prepared by the National Greenhouse Gas Inventories Programme. IGES

IPCC, Buendia, E., Tanabe, K., Kranjc, A., Jamsranjav, B., Fukuda, M., Ngarize, S., Osako, A., Pyrozhenko, Y., Shermanau, P., and Federici, S. (2019). 2019 Refinement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories

- Iraci, L. T., Podolske, J. R., Roehl, C., Wennberg, P. O., Blavier, J.-F., Allen, N., Wunch, D., & Osterman, G. B. (2022). TCCON data from Edwards (US), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.edwards01.R0>
- Jones, M. W., Andrew, R. M., Peters, G. P., Janssens-Maenhout, G., De-Gol, A. J., Ciais, P., Patra, P. K., Chevallier, F., & Le Quéré, C. (2021). Gridded fossil CO₂ emissions and related O₂ combustion consistent with national inventories 1959–2018. *Scientific Data*, 8(1), Article 1. <https://doi.org/10.1038/s41597-020-00779-6>
- Karion, A., Sweeney, C., Tans, P., & Newberger, T. (2010). AirCore: An Innovative Atmospheric Sampling System. *Journal of Atmospheric and Oceanic Technology*, 27(11), 1839–1853. <https://doi.org/10.1175/2010JTECHA1448.1>
- Kenneth N. Schuldt, John Mund, Ingrid T. Lujkx, Tuula Aalto, James B. Abshire, Ken Aikin, Arlyn Andrews, Shuji Aoki, Francesco Apadula, Bianca Baier, Peter Bakwin, Jakub Bartyzel, Gilles Bentz, Peter Bergamaschi, Andreas Beyersdorf, Tobias Biermann, Sebastien C. Biraud, Harald Boenisch, David Bowling, Gordon Brailsford, Willi A. Brand, Huilin Chen, Gao Chen, Lukasz Chmura, Shane Clark, Sites Climadat, Aurelie Colomb, Roisin Commane, Sébastien Conil, Cedric Couret, Adam Cox, Paolo Cristofanelli, Emilio Cuevas, Roger Curcoll, Bruce Daube, Kenneth Davis, Martine De Mazière, Stephan De Wekker, Julian Della Coletta, Marc Delmotte, Joshua P. DiGangi, Ed Dlugokencky, James W. Elkins, Lukas Emmenegger, Shuangxi Fang, Marc L. Fischer, Grant Forster, Arnoud Frumau, Michal Galkowski, Luciana V. Gatti, Torsten Gehrlein, Christoph Gerbig, Francois Gheusi, Emanuel Gloor, Vanessa Gomez-Trueba, Daisuke Goto, Tim Griffis, Samuel Hammer, Chad Hanson, László Haszpra, Juha Hatakka, Martin Heimann, Michal Heliasz, Daniela Heltai, Arjan Hensen, Ove Hermanssen, Eric Hintsa, Antje Hoheisel, Jutta Holst, Viktor Ivakhov, Dan Jaffe, Armin Jordan, Warren Joubert, Anna Karion, Stephan R. Kawa, Victor Kazan, Ralph Keeling, Petri Keronen, Jooil Kim, Tobias Kneuer, Pasi Kolari, Katerina Kominkova, Eric Kort, Elena Kozlova, Paul Krummel, Dagmar Kubistin, Casper Labuschagne, David H. Lam, Xin Lan, Ray Langenfelds, Olivier Laurent, Tuomas Laurila, Thomas Lauvaux, Jost Lavric, Bev Law, Olivia S. Lee, John Lee, Irene Lehner, Kari Lehtinen, Reimo Leppert, Ari Leskinen, Markus Leuenberger, Ingeborg Levin, Janne Levula, John Lin, Matthias Lindauer, Zoe Loh, Morgan Lopez, Chris R. Lunder, Toshinobu Machida, Ivan Mammarella, Giovanni Manca, Alistair Manning, Andrew Manning, Michal V. Marek, Melissa Y. Martin, Giordane A. Martins, Hidekazu Matsueda, Kathryn McKain, Harro Meijer, Frank Meinhardt, Lynne Merchant, N. Mihalopoulos, Natasha Miles, John B. Miller, Charles E. Miller, Logan Mitchell, Stephen Montzka, Fred Moore, Heiko Moossen, Eric Morgan, Josep-Anton Morgui, Shinji Morimoto, Bill Munger, David Munro, Cathrine L. Myhre, Meelis Mölder, Jennifer Müller-Williams, Jaroslaw Necki, Sally Newman, Sylvia Nichol, Yosuke Niwa, Simon O'Doherty, Florian Obersteiner, Bill Paplawsky, Jeff Peischl, Olli Peltola, Salvatore Piacentino, Jean M. Pichon, Steve Piper, Joseph Pitt, Christian Plass-Duelmer, Stephen M. Platt, Steve Prinzivalli, Michel Ramonet, Ramon Ramos, Enrique Reyes-Sanchez, Scott Richardson, Haris Riris, Pedro P. Rivas, Michael Rothe, Thomas Ryerson, Kazuyuki Saito, Maryann Sargent, Motoki Sasakawa, Bert Scheeren, Martina Schmidt, Tanja Schuck, Marcus Schumacher, Thomas Seifert, Mahesh K. Sha, Paul Shepson, Michael Shook, Christopher D. Sloop, Paul Smith, Martin Steinbacher, Britton Stephens, Colm Sweeney, Lise L. Sørensen, Pieter Tans, Kirk Thoning, Helder Timas, Margaret Torn, Pamela Trisolino, Jocelyn Turnbull, Kjetil Tørseth, Alex Vermeulen, Brian Viner, Gabriela Vitkova, Stephen Walker, Andrew Watson, Ray Weiss, Steve Wofsy, Justin Worsley, Doug Worthy, Dickon Young, Sönke Zaehle, Andreas Zahn, Miroslaw Zimnoch, Rodrigo A. de Souza, Alcide G. di Sarra, Danielle van Dinther, Pim van den Bulk (2022). Multi-laboratory compilation of atmospheric carbon dioxide data for the period 1957-2021; obspack_co2_1_GLOBALVIEWplus_v8.0_2022-08-27; NOAA Earth System Research Laboratory, Global Monitoring Laboratory. <http://doi.org/10.25925/20220808>

- Kivi, R., Heikkinen, P., & Kyrö, E. (2022). TCCON data from Sodankylä (FI), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.sodankyla01.R0>
- Krinner, G., Viovy, N., de Noblet-Ducoudré, N., Ogée, J., Polcher, J., Friedlingstein, P., Ciais, P., Sitch, S., & Prentice, I. C. (2005). A dynamic global vegetation model for studies of the coupled atmosphere-biosphere system. *Global Biogeochemical Cycles*, 19(1). <https://doi.org/10.1029/2003GB002199>
- Lan, X., J.W. Mund, A.M. Crotwell, M.J. Crotwell, E. Moglia, M. Madronich, D. Neff and K.W. Thoning (2023), Atmospheric Carbon Dioxide Dry Air Mole Fractions from the NOAA GML Carbon Cycle Cooperative Global Air Sampling Network, 1968-2022, Version: 2023-08-28, <https://doi.org/10.15138/wkgj-f215>
- Lauvaux, T., Pannekoucke, O., Sarrat, C., Chevallier, F., Ciais, P., Noilhan, J., & Rayner, P. J. (2009). Structure of the transport uncertainty in mesoscale inversions of CO₂ sources and sinks using ensemble model simulations. *Biogeosciences*, 6(6), 1089–1102. <https://doi.org/10.5194/bg-6-1089-2009>
- Lin, X., Ciais, P., Bousquet, P., Ramonet, M., Yin, Y., Balkanski, Y., Cozic, A., Delmotte, M., Evangeliou, N., Indira, N. K., Locatelli, R., Peng, S., Piao, S., Saunois, M., Swathi, P. S., Wang, R., Yver-Kwok, C., Tiwari, Y. K., & Zhou, L. (2018). Simulating CH₄ and CO₂ over South and East Asia using the zoomed chemistry transport model LMDz-INCA. *Atmospheric Chemistry and Physics*, 18(13), 9475–9497. <https://doi.org/10.5194/acp-18-9475-2018>
- Liu, Z., Zeng, N., Liu, Y., Wang, J., Han, P., & Cai, Q. (2024). Weaker regional carbon uptake albeit with stronger seasonal amplitude in northern mid-latitudes estimated by higher resolution GEOS-Chem model. *Science of The Total Environment*, 912, 169477. <https://doi.org/10.1016/j.scitotenv.2023.169477>
- Locatelli, R., Bousquet, P., Hourdin, F., Saunois, M., Cozic, A., Couvreux, F., Grandpeix, J.-Y., Lefebvre, M.-P., Rio, C., Bergamaschi, P., Chambers, S. D., Karstens, U., Kazan, V., van der Laan, S., Meijer, H. A. J., Moncrieff, J., Ramonet, M., Scheeren, H. A., Schlosser, C., ... Williams, A. G. (2015). Atmospheric transport and chemistry of trace gases in LMDz5B: Evaluation and implications for inverse modelling. *Geoscientific Model Development*, 8(2), 129–150. <https://doi.org/10.5194/gmd-8-129-2015>
- De Mazière, M., Sha, M. K., Desmet, F., Hermans, C., Scolas, F., Kumps, N., Zhou, M., Metzger, J.-M., Dufлот, V., & Cammas, J.-P. (2022). TCCON data from Réunion Island (RE), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.reunion01.R0>
- Miles, N.L., S.J. Richardson, K.J. Davis, and B.J. Haupt, 2017. In-situ tower atmospheric measurements of carbon dioxide, methane and carbon monoxide mole fraction for the Indianapolis Flux (INFLUX) project, Indianapolis, IN, USA. Data set. Available on-line <http://datacommons.psu.edu> from The Pennsylvania State University Data Commons, University Park, Pennsylvania, USA. <http://dx.doi.org/10.18113/D37G6P>
- Miles, N.L., S.J. Richardson, D.K. Martins, K.J. Davis, T. Lauvaux, B.J. Haupt, and S.K. Miller. 2018. ACT-America: L2 In Situ CO₂, CO, and CH₄ Concentrations from Towers, Eastern USA. ORNL DAAC, Oak Ridge, Tennessee, USA. <https://doi.org/10.3334/ORNLLDAAC/1568>
- Morino, I., Ohyama, H., Hori, A., & Ikegami, H. (2022). TCCON data from Rikubetsu (JP), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.rikubetsu01.R0>

- Morino, I., Ohyama, H., Hori, A., & Ikegami, H. (2022). TCCON data from Tsukuba (JP), 125HR, Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.tsukuba02.R0>
- Notholt, J., Petri, C., Warneke, T., & Buschmann, M. (2022). TCCON data from Bremen (DE), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.bremen01.R0>
- OCO-2/OCO-3 Science Team, Vivienne Payne, Abhishek Chatterjee (2022), OCO-2 Level 2 bias-corrected XCO₂ and other select fields from the full-physics retrieval aggregated as daily files, Retrospective processing V11.1r, Greenbelt, MD, USA, Goddard Earth Sciences Data and Information Services Center (GES DISC), Accessed: 5 October 2023, 10.5067/8E4VLCK16O6Qk
- O'Dell, C. W., Eldering, A., Wennberg, P. O., Crisp, D., Gunson, M. R., Fisher, B., Frankenberg, C., Kiel, M., Lindqvist, H., Mandrake, L., Merrelli, A., Natraj, V., Nelson, R. R., Osterman, G. B., Payne, V. H., Taylor, T. E., Wunch, D., Drouin, B. J., Oyafuso, F., ... Velazco, V. A. (2018). Improved retrievals of carbon dioxide from Orbiting Carbon Observatory-2 with the version 8 ACOS algorithm. *Atmospheric Measurement Techniques*, 11(12), 6539–6576. <https://doi.org/10.5194/amt-11-6539-2018>
- O'Dell, C. W., Osterman, G., & team. (2023). Orbiting Carbon Observatory-2 & 3 (OCO-2 & OCO-3) Data Product User's Guide, Operational Level 2 Lite Files – Version 2.0 – Revision A – July 5, 2022 – Data Release: 11 (OCO-2), 10/10.4 (OCO-3), https://docserver.gesdisc.eosdis.nasa.gov/public/project/OCO/OCO2_V11_OCO3_V10_DUG.pdf
- Peiro, H., Crowell, S., Schuh, A., Baker, D. F., O'Dell, C., Jacobson, A. R., Chevallier, F., Liu, J., Eldering, A., Crisp, D., Deng, F., Weir, B., Basu, S., Johnson, M. S., Philip, S., & Baker, I. (2022). Four years of global carbon cycle observed from the Orbiting Carbon Observatory 2 (OCO-2) version 9 and in situ data and comparison to OCO-2 version 7. *Atmospheric Chemistry and Physics*, 22(2), 1097–1130. <https://doi.org/10.5194/acp-22-1097-2022>
- Prather, M. J., Zhu, X., Strahan, S. E., Steenrod, S. D., & Rodriguez, J. M. (2008). Quantifying errors in trace species transport modeling. *Proceedings of the National Academy of Sciences*, 105(50), 19617–19621. <https://doi.org/10.1073/pnas.0806541106>
- Remaud, M., Chevallier, F., Cozic, A., Lin, X., & Bousquet, P. (2018). On the impact of recent developments of the LMDz atmospheric general circulation model on the simulation of CO₂ transport. *Geosci. Model Dev.*, 25.
- Rayner, P. J., & O'Brien, D. M. (2001). The utility of remotely sensed CO₂ concentration data in surface source inversions. *Geophysical Research Letters*, 28(1), 175–178. <https://doi.org/10.1029/2000GL011912>
- Schneider, T., Teixeira, J., Bretherton, C. S., Brient, F., Pressel, K. G., Schär, C., & Siebesma, A. P. (2017). Climate goals and computing the future of clouds. *Nature Climate Change*, 7(1), Article 1. <https://doi.org/10.1038/nclimate3190>
- Sherlock, V., Connor, B., Robinson, J., Shiona, H., Smale, D., & Pollard, D. F. (2022). TCCON data from Lauder (NZ), 125HR, Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.lauder02.R0>

Shiomi, K., Kawakami, S., Ohyama, H., Arai, K., Okumura, H., Ikegami, H., & Usami, M. (2022). TCCON data from Saga (JP), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.saga01.R0>

Strong, K., Roche, S., Franklin, J. E., Mendonca, J., Lutsch, E., Weaver, D., Fogal, P. F., Drummond, J. R., Batchelor, R., Lindenmaier, R., & McGee, E. (2022). TCCON data from Eureka (CA), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.eureka01.R0>

Sussmann, R., & Rettinger, M. (2023). TCCON data from Garmisch (DE), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.garmisch01.R0>

Té, Y., Jeseck, P., & Janssen, C. (2022). TCCON data from Paris (FR), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.paris01.R0>

Tolk, L. F., Meesters, A. G. C. A., Dolman, A. J., & Peters, W. (2008). Modelling representation errors of atmospheric CO₂ mixing ratios at a regional scale. *Atmospheric Chemistry and Physics*, 8(22), 6587–6596. <https://doi.org/10.5194/acp-8-6587-2008>

Wagenhäuser, T., Engel, A., & Sitals, R. (2021). Testing the altitude attribution and vertical resolution of AirCore measurements with a new spiking method. *Atmospheric Measurement Techniques*, 14(5), 3923–3934. <https://doi.org/10.5194/amt-14-3923-2021>

Warneke, T., Petri, C., Notholt, J., & Buschmann, M. (2022). TCCON data from Orléans (FR), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.orleans01.R0>

Wennberg, P. O., Roehl, C. M., Wunch, D., Toon, G. C., Blavier, J.-F., Washenfelder, R., Keppel-Aleks, G., & Allen, N. T. (2022). TCCON data from Park Falls (US), Release GGG2020.R1 (Version R1) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.parkfalls01.R1>

Wennberg, P. O., Roehl, C. M., Wunch, D., Blavier, J.-F., Toon, G. C., Allen, N. T., Treffers, R., & Laughner, J. (2022). TCCON data from Caltech (US), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.pasadena01.R0>

Wennberg, P. O., Wunch, D., Roehl, C. M., Blavier, J.-F., Toon, G. C., & Allen, N. T. (2022). TCCON data from Lamont (US), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.lamont01.R0>

Wunch, D., Mendonca, J., Colebatch, O., Allen, N. T., Blavier, J.-F., Kunz, K., Roche, S., Hedelius, J., Neufeld, G., Springett, S., Worthy, D., Kessler, R., & Strong, K. (2022). TCCON data from East Trout Lake, SK (CA), Release GGG2020.R0 (Version R0) [Data set]. CaltechDATA. <https://doi.org/10.14291/tccon.ggg2020.eastroutlake01.R0>

Appendix A: Observation datasets

Table A1 presents the datasets used from the Obspack database as well as the corresponding abbreviated site code for each station used in the main text.

Table A2 presents in a similar way the list of TCCON sites used in the study.

Table A1. List of datasets used from Obspack for surface stations

Site code	Dataset
AirCoreNOAA	aircorenoaa_aircore_1_allvalid
ABT	abt_surface-insitu_6_allvalid
ALT	alt_surface-flask_426_representative
ALT	alt_surface-insitu_6_allvalid
ALT	alt_surface-flask_1_representative
ALT	alt_surface-flask_2_representative
ALT	alt_surface-flask_4_representative
AMS	ams_surface-flask_1_representative
AMS	ams_surface-insitu_11_allvalid
AMT	amt_tower-insitu_1_allvalid-30magl
AMT	amt_tower-insitu_1_allvalid-12magl
AMT	amt_surface-pfp_1_allvalid-107magl
AMT	amt_tower-insitu_1_allvalid-107magl
AZV	azv_tower-insitu_20_allvalid-29magl
AZV	azv_tower-insitu_20_allvalid-50magl
BAO	bao_tower-insitu_1_allvalid-100magl
BAO	bao_tower-insitu_1_allvalid-300magl

BAO	bao_surface-pfp_1_allvalid-300magl
BAO	bao_tower-insitu_1_allvalid-22magl
BCK	bck_surface-insitu_6_allvalid
BIR	bir_surface-insitu_56_allvalid
BRA	bra_surface-insitu_6_allvalid
BRM	brm_tower-insitu_49_allvalid-12magl
BRM	brm_tower-insitu_49_allvalid-72magl
BRM	brm_tower-insitu_49_allvalid-45magl
BRM	brm_tower-insitu_49_allvalid-212magl
BRM	brm_tower-insitu_49_allvalid-132magl
BRW	brw_surface-insitu_1_allvalid
BRW	brw_surface-flask_4_representative
BRW	brw_surface-flask_1_representative
BRW	brw_surface-flask_426_representative
BRZ	brz_tower-insitu_20_allvalid-20magl
BRZ	brz_tower-insitu_20_allvalid-80magl
BRZ	brz_tower-insitu_20_allvalid-5magl
BRZ	brz_tower-insitu_20_allvalid-40magl
BSD	bsd_tower-insitu_160_allvalid-108magl

BSD	bsd_tower-insitu_160_allvalid-248magl
BSD	bsd_tower-insitu_160_allvalid-42magl
BU	bu_surface-insitu_59_allhours
CBW	cbw_tower-insitu_445_allvalid-27magl
CBW	cbw_tower-insitu_445_allvalid-67magl
CBW	cbw_tower-insitu_445_allvalid-127magl
CBW	cbw_tower-insitu_445_allvalid-207magl
CBY	cby_surface-insitu_6_allvalid
CHL	chl_surface-insitu_6_allvalid
CIT	cit_surface-insitu_115_allhours-200magl
COP	cop_tower-insitu_59_allhours
CPS	cps_surface-insitu_6_allvalid
CPT	cpt_surface-flask_1_representative
CPT	cpt_surface-insitu_36_marine
CRV	crv_tower-insitu_1_allvalid-32magl
CRV	crv_surface-pfp_1_allvalid-32magl
CRV	crv_tower-insitu_1_allvalid-17magl
CRV	crv_tower-insitu_1_allvalid-5magl
DEC	dec_surface-insitu_431_allvalid

DEM	dem_tower-insitu_20_allvalid-45magl
DEM	dem_tower-insitu_20_allvalid-63magl
EEC	eec_surface-insitu_431_allvalid
EGB	egb_surface-insitu_6_allvalid
ENA	ena_surface-insitu_64_allvalid-10magl
ESP	esp_surface-flask_2_representative
ESP	esp_surface-insitu_6_allvalid
EST	est_surface-insitu_6_allvalid
ETL	etl_surface-insitu_6_allvalid
FSD	fsd_surface-insitu_6_allvalid
GCI01	gci01_tower-insitu_60_allvalid
GCI02	gci02_tower-insitu_60_allvalid
GCI03	gci03_tower-insitu_60_allvalid
GCI04	gci04_tower-insitu_60_allvalid
GCI05	gci05_tower-insitu_60_allvalid
GIC	gic_surface-insitu_431_allvalid
GIF	gif_surface-insitu_11_allvalid
GOULD	gould_shipboard-insitu_1_allvalid
HDP	hdp_surface-insitu_3_nonlocal

HEI	hei_surface-insitu_22_allvalid
HFM	hfm_tower-insitu_59_allhours
HNP	hnp_surface-insitu_6_allvalid
HTM	htm_tower-insitu_424_allvalid-70magl
HTM	htm_tower-insitu_424_allvalid-30magl
HTM	htm_tower-insitu_424_allvalid-150magl
HUN	hun_tower-insitu_35_allvalid-48magl
HUN	hun_tower-insitu_35_allvalid-10magl
HUN	hun_tower-insitu_35_allvalid-115magl
HUN	hun_tower-insitu_35_allvalid-82magl
HUN	hun_surface-flask_1_representative
INU	inu_surface-insitu_6_allvalid
INX01	inx01_surface-insitu_60_allhours
INX02	inx02_surface-insitu_60_allhours
INX03	inx03_surface-insitu_60_allhours
INX04	inx04_surface-insitu_60_allhours
INX06	inx06_surface-insitu_60_allhours
INX07	inx07_surface-insitu_60_allhours
INX08	inx08_surface-insitu_60_allhours

INX09	inx09_surface-insitu_60_allhours
INX10	inx10_surface-insitu_60_allhours
INX11	inx11_surface-insitu_60_allhours
INX13	inx13_surface-insitu_60_allhours
JFJ	jfj_surface-insitu_5_allvalid
JFJ	jfj_surface-insitu_49_allvalid
KAS	kas_surface-insitu_53_allvalid
KCMP	kcmp_tower-insitu_102_allhours-200magl
KRS	krs_tower-insitu_20_allvalid-67magl
KRS	krs_tower-insitu_20_allvalid-35magl
LAN	lan_surface-insitu_33_allvalid
LEF	lef_tower-insitu_1_allvalid-244magl
LEF	lef_tower-insitu_1_allvalid-122magl
LEF	lef_surface-pfp_1_allvalid-396magl
LEF	lef_tower-insitu_1_allvalid-30magl
LEF	lef_tower-insitu_1_allvalid-11magl
LEF	lef_tower-insitu_1_allvalid-76magl
LEF	lef_tower-insitu_1_allvalid-396magl
LEF	lef_surface-pfp_1_allvalid-244magl

LFS	lfs_surface-insitu_33_allvalid
LLB	llb_surface-insitu_6_allvalid
LLB	llb_surface-flask_1_representative
MBO	mbo_surface-pfp_1_allvalid-11magl
MBO	mbo_surface-insitu_1_allvalid-11magl
MLO	mlo_surface-flask_1_representative
MLO	mlo_surface-flask_4_representative
MLO	mlo_surface-flask_426_representative
MLO	mlo_surface-flask_2_representative
MLO	mlo_surface-insitu_1_allvalid
MNM	mnm_surface-insitu_19_representative
MRC	mrc_surface-pfp_1_allvalid-south
MRC	mrc_tower-insitu_60_allvalid-south
MRC	mrc_surface-pfp_1_allvalid-east
NOR	nor_tower-insitu_424_allvalid-59magl
NOR	nor_tower-insitu_424_allvalid-100magl
NOR	nor_tower-insitu_424_allvalid-32magl
NOY	noy_tower-insitu_20_allvalid-43magl
NOY	noy_tower-insitu_20_allvalid-21magl

NWR	nwr_surface-pfp_1_allvalid-3magl
NWR	nwr_surface-insitu_3_nonlocal
NWR	nwr_surface-flask_1_representative
OLI	oli_surface-insitu_64_allvalid-10magl
OMP	omp_surface-insitu_68_allhours
ONG	ong_surface-insitu_68_allhours
OPE	ope_tower-insitu_11_allvalid-120magl
OSI	osi_tower-insitu_68_allhours-269magl
OSI	osi_tower-insitu_68_allhours-31magl
OWA	owa_surface-insitu_68_allhours
PAL	pal_surface-flask_1_representative
PAL	pal_surface-insitu_30_nonlocal
PAL	pal_surface-insitu_30_continental
PAL	pal_surface-insitu_30_marine
PDM	pdm_surface-flask_11_representative
PDM	pdm_surface-insitu_11_allvalid
PRS	prs_surface-insitu_21_allvalid
PUY	puy_surface-insitu_11_allvalid
PV	pv_surface-insitu_115_allhours-200magl

RGL	rgl_tower-insitu_160_allvalid-45magl
RGL	rgl_tower-insitu_160_allvalid-90magl
RYO	ryo_surface-insitu_19_representative
SCT	sct_tower-insitu_1_allvalid-61magl
SCT	sct_surface-pfp_1_allvalid-305magl
SCT	sct_tower-insitu_1_allvalid-305magl
SCT	sct_tower-insitu_1_allvalid-31magl
SGP	sgp_surface-insitu_64_allvalid-60magl
SGP	sgp_surface-flask_1_representative
SMO	smo_surface-flask_426_representative
SMO	smo_surface-flask_1_representative
SMO	smo_surface-insitu_1_allvalid
SMO	smo_surface-flask_4_representative
SMR	smr_tower-insitu_421_allvalid-67magl
SMR	smr_tower-insitu_421_allvalid-17magl
SMR	smr_tower-insitu_421_allvalid-125magl
SNP	snp_surface-insitu_1_allvalid-10magl
SNP	snp_surface-insitu_1_allvalid-5magl
SNP	snp_surface-insitu_1_allvalid-17magl

SPL	spl_surface-insitu_3_nonlocal
SPO	spo_surface-flask_4_representative
SPO	spo_surface-flask_2_representative
SPO	spo_surface-insitu_1_allvalid
SPO	spo_surface-flask_426_representative
SPO	spo_surface-flask_1_representative
SSC	ssc_surface-insitu_431_allvalid
SSL	ssl_surface-insitu_107_allvalid
SYO	syo_surface-insitu_8_allvalid
SYO	syo_surface-flask_1_representative
TAC	tac_tower-insitu_160_allvalid-185magl
TAC	tac_surface-flask_1_representative
TAC	tac_tower-insitu_160_allvalid-54magl
TAC	tac_tower-insitu_160_allvalid-100magl
TIK	tik_surface-insitu_30_allvalid
TIK	tik_surface-flask_1_representative
TPD	tpd_surface-insitu_6_allvalid
TRN	trn_tower-insitu_11_allvalid-180magl
UTDBK	utdbk_tower-insitu_432_allvalid

UTMSA	utmsa_tower-insitu_432_allvalid
UTRPK	utrpk_tower-insitu_432_allvalid
UTSUG	utsug_tower-insitu_432_allvalid
UTUOU	utuou_tower-insitu_432_allvalid
VAC	vac_surface-insitu_431_allvalid
VGN	vgn_tower-insitu_20_allvalid-42magl
VGN	vgn_tower-insitu_20_allvalid-85magl
WAO	wao_surface-insitu_13_allvalid
WBI	wbi_tower-insitu_1_allvalid-31magl
WBI	wbi_tower-insitu_1_allvalid-99magl
WBI	wbi_tower-insitu_1_allvalid-379magl
WBI	wbi_surface-pfp_1_allvalid-379magl
WGC	wgc_tower-insitu_1_allvalid-483magl
WGC	wgc_surface-pfp_1_allvalid-91magl
WGC	wgc_tower-insitu_1_allvalid-91magl
WGC	wgc_tower-insitu_1_allvalid-30magl
WGC	wgc_surface-pfp_1_allvalid-483magl
WKT	wkt_tower-insitu_1_allvalid-244magl
WKT	wkt_tower-insitu_1_allvalid-62magl

WKT	wkt_tower-insitu_1_allvalid-457magl
WKT	wkt_tower-insitu_1_allvalid-30magl
WKT	wkt_tower-insitu_1_allvalid-122magl
WKT	wkt_surface-pfp_1_allvalid-122magl
WKT	wkt_tower-insitu_1_allvalid-9magl
WKT	wkt_surface-pfp_1_allvalid-457magl
YON	yon_surface-insitu_19_representative
ZEP	zep_surface-insitu_56_allvalid
ZEP	zep_surface-flask_1_representative

Table A2. List of TCCON sites used and their locations

TCCON code	Location
br	Bremen, Germany
ci	Pasadena, California, USA
db	Darwin, Australia
df	Edwards, USA
et	East Trout Lake, Canada
eu	Eureka, Canada
gm	Garmisch, Germany
hf	Hefei, China

js	Saga, Japan
oc	Lamont, Oklahoma, USA
ll	Lauder, New Zealand
ma	Manaus, Brazil
ny	Ny-Alesund, Svalbard, Norway
or	Orleans, France
pa	Park Falls, Wisconsin, USA
pr	Paris, France
ra	Reunion Island, France
rj	Rikubetsu, Hokkaido, Japan
so	Sodankyla, Finland
tk	Tsukuba, Ibaraki, Japan

Part IV

Conclusions and perspectives



Conclusion

LIMBIC SYSTEM — Beep-beep-beep! The alarm is ringing, Harry. The disco circus goes on and on! You barely slept three hours last night.

VOLITION — You can do it. It's nothing. Do it for the city. Go.

SHIVERS — Do it for the wind.

LOGIC — Do it for the picture puzzle. Put it all together. Solve the world. One conversation at a time.

Disco Elysium

The goal of this thesis was to study and implement different ways of adapting an atmospheric greenhouse gas inverse system to new supercomputing resources. The growing number of processors or threads allows massive parallelization of computations which can be leveraged to improve atmospheric inverse models. These models, combined with observations of atmospheric CO₂, help us understand climate change by enhancing our knowledge of the greenhouse gas sources and sinks that drive and modulate it.

In this thesis, we studied two complementary ways of improving an inverse system:

- **Using a new icosahedral grid from DYNAMICO**
- **Increasing the horizontal resolution of the model**

The CAMS/LSCE inverse system we focus on in this work derives its atmospheric transport model from a GCM that recently integrated a new dynamical core, DYNAMICO, that uses a quasi-uniform icosahedral grid.

We first had to prepare this new version of the GCM since it was not yet ready to be used in a configuration suited for our inverse system. It lacked the capability of nudging the winds of the GCM to wind fields from a reanalysis so that the representation of atmospheric transport could be compared to individual observations meaningfully. The GCM also needed to generate air mass fluxes since they are the input that drives the inverse system.

The impact of the icosahedral grid of DYNAMICO was then studied in the GCM by comparing it to a reference configuration of the GCM using a regular latitude-longitude grid that has comparable horizontal resolution at low latitudes. We compared the outputs of a 40-year-long simulation to independent observations of CO₂ to see if the new grid provided a better representation of tracer

atmospheric transport. We also directly compared the computational performance of the two configurations. This first study demonstrated that the new grid was just as effective as the reference grid in modeling atmospheric transport at the global scale, showing only small differences mainly in terms of vertical transport which can be improved with further tuning. There was no major difference near the poles, where the two grids differed the most. The new configuration was computationally faster, predominantly due to fewer grid cells for an equivalent resolution of the reference configuration at low latitudes. It can also scale much better and use many more processors in parallel compared to the regular grid, providing a clear advantage when increasing the spatial resolution of the grids.

The main advantage of the icosahedral grid compared to a regular latitude-longitude grid would only be evidenced at a higher resolution than the initial one of our inverse system. We therefore decided to first focus on increasing the inverse model's resolution with the regular grid before implementing DYNAMICO in it. We did so by replacing the original MPI parallelization of the inverse model with a new parallelization scheme running on GPU. This allowed us to increase the resolution of the inverse system from 3.75° in longitude by 1.90° in latitude with 39 vertical layers to 2.50° in longitude by 1.27° in latitude with 79 vertical layers, the same resolution used as a reference in the previous study of the GCM. This first increase in resolution was only made possible by the GPU parallelization since it makes the model computation 8 times faster than the equivalent MPI parallelization. This meant that the inverse model at the new resolution running on GPU was even faster than the original version at the old resolution.

Thanks to the unexpected breakthrough in computational speed provided by the GPU parallelization we could immediately increase the horizontal resolution of the inverse model using the regular latitude-longitude grid a second time rather than waiting for a new version using DYNAMICO. We did so and almost doubled the resolution of our inverse model to 1.41° in longitude by 0.70° in latitude. We studied the impact of this resolution increase on a two-year-long global inversion of atmospheric CO_2 assimilating satellite data from the OCO-2 mission. The high-resolution inversion showed a small improvement in the representation of CO_2 transport mainly driven by a better orography and coastal definition. It also improved the vertical transport of CO_2 particularly at low altitudes. The optimized surface fluxes of CO_2 obtained by the high-resolution inversion attribute a slightly higher portion of the global carbon sink to the land rather than to the ocean compared to the lower-resolution inversion.

By comparison, throughout this thesis, the operational CAMS CO_2 forecasting system maintained by ECMWF has had a resolution of 25 km for the assimilation and 9 km for the forecast itself. This high resolution is allowed by the much shorter time scales addressed by the ECMWF system compared to the inversion system: hours vs. decades. However, this thesis triggered a dramatic reduction of the scale gap, with the inversion model at a few degrees at the start and at about 1 degree at the end.

Further increases in resolution of this version of the CAMS/LSCE inverse system are not directly limited by technical constraints of the inverse mode anymore but instead only by the future obsolescence of the regular latitude-longitude version of the master GCM which will not be officially supported at higher resolutions than 512×361 . New versions of this inverse system based on DYNAMICO, whose transport model was validated for long-term tracer transport in the first study of this thesis, have no such limitation and will continue to quickly increase in resolution in the future.



Perspectives

This work is part of a wide effort to increase the spatial resolution of global climate and atmospheric inverse models. Specifically, for inverse models, this effort aims to improve our understanding of the sources and sinks of greenhouse gases. This manuscript presented the technical developments used to make such a resolution increase possible and a scientific evaluation of their impact. Both aspects were studied within specific models, but many of the conclusions drawn from this work are applicable in a broader context.

1 The future of the CAMS/LSCE inverse system

The work presented in this manuscript resulted in a version of the CAMS/LSCE inverse system running on a regular latitude-longitude grid at a resolution of 0.7° latitude by 1.4° longitude. This version is already in use for the production of the CAMS global inversion-optimized greenhouse gas fluxes and concentrations dataset. It also offers new possibilities to improve this inverse system in innovative ways, even in the short term.

1.1 Inverse system using DYNAMICO

As explained throughout this manuscript and particularly in [chapter IV \(Part I\)](#), the initial short-term goal of this work was to implement DYNAMICO into our inverse system after evaluating it in the master GCM. However since it showed only small advantages at low spatial resolution and its implementation required a significant engineering investment, the decision was made to first focus on directly increasing the resolution of the inverse system on the regular grid. In parallel, an offline version of DYNAMICO was being developed and eventually its tangent-linear and adjoint, creating a full inverse system. Since the regular grid is reaching its limits due to the lack of technical and scientific support of the master GCM at high resolutions in particular, switching to DYNAMICO is a necessary next step for the CAMS/LSCE inverse system.

1.1.1 Current state of the inverse system

As of July 2024, a version of the CAMS/LSCE inverse system based on the transport model of DYNAMICO is already close to operational. Though final validation in the full CAMS production environment is still underway, so far its performance and results in terms of estimations of CO_2 surface fluxes seem similar to the version on the regular latitude-longitude grid. It runs on the nbp80 resolution of DYNAMICO: using 64002 hexagonal cells of diagonal 110 km over 79 vertical levels. Interestingly, the number of cells is close to the one of the regular 1.41° in longitude by 0.70° in latitude, which means

that it allows a redistribution of the cell density along the latitudes: more cells at low latitudes, less at high ones (but without the need of a longitudinal filter anymore). It also runs on GPU, taking advantage of the previous developments on the regular version to enhance its computational speed. At this resolution, a global inversion of atmospheric CO₂ assimilating OCO-2 data takes a couple of weeks on the Irene V100 CPU/GPU hybrid supercomputer of the TGCC. These inversions are still driven by pre-generated air mass fluxes from the master GCM in its configuration coupling LMDZ, DYNAMICO, and ORCHIDEE. This configuration of the master GCM only runs on CPU, and while porting the code to run on GPU is a long-term goal (or at least for parts of the code), it requires vast engineering resources and is an ongoing work of the IPSL modeling group. Since the master GCM scales better and can use many more CPUs with DYNAMICO, another increase in horizontal resolution to nbp160 (hexagons of diagonal 55 km or around 0.5° near the equator) should soon be possible. While the computational advantage of DYNAMICO in the master GCM is clear, and therefore directly benefits the inverse system through faster generation of the input mass fluxes, the inverse system itself is not much faster with DYNAMICO. There is a direct computational speed advantage resulting from the lower number of cells at equivalent resolutions in the Tropics, but since each cell of the icosahedral grid has more sides than cells from a rectangular grid, more calculations need to be done for mass exchanges between cells. Further, spatial coordinates are 3D with DYNAMICO while the standard latitude-longitude grid allows for a 2D representation of all variables. This complexity must be kept in mind when deciding on which grid to use for modeling atmospheric tracer transport.

1.1.2 Possible optimization

This version of the CAMS/LSCE inverse system can still be optimized for better computational performance and representation of atmospheric transport. In the inverse system, in contrast to the GCM, the cells are not ordered optimally in the computer memory: neighboring cells on the grid are not necessarily allocated to the same chunks of memory in the GPU. This creates an overhead when computing data exchange between these cells since additional calls have to be made between their memory pools. Properly indexing the cells and optimizing their memory allocations is a difficult endeavor but will be necessary if we want to take full advantage of the potential of the new grid. The parallelization scheme by tiles for the GCM shown in [Figure V.7](#) (presented in [chapter V](#) on page 48) could be an inspiration for a similar scheme in the offline transport model.

The current version of the tangent-linear and adjoint of the DYNAMICO offline transport model did not include slope-limiters for the advection since they introduce non-linearities that greatly slow down the computational speed of the adjoint model. Their impact on the overall representation of atmospheric transport was judged to be small enough to be ignored for now given the first comparisons with the regular grid, but this assumption warrants further investigation in the future, especially at higher resolution.

Finally, the efficiency of the transport computations on the GPU leaves reading the meteorological data to be the main computation bottleneck. The transfer rates between the file systems and the computing nodes remain a limiting factor to completely fill the gap with the 9-km IFS resolution for year-long inversion windows because the input data volume is directly proportional to the number of 3D cells, and it is still 100 times larger in the horizontal for the IFS than our finest resolution. Reducing the volume of the input meteorological dataset could provide larger speed gains than further optimization of specific routines of the code, for instance with some tailored compression that would minimize the loss in accuracy of the simulations, possibly in an adaptive way, i.e. depending on the inversion problem. However, hardware improvements for faster data transfer rates may have more impact on runtime than software ones in general because they do not affect simulation accuracy.

1.2 Bypassing the master GCM

The current process for producing a global inversion of atmospheric CO₂, whether in the regular configuration or the one with DYNAMICO takes multiple steps:

1. Collect wind fields from the ERA5 reanalysis at regular intervals
2. Use them to nudge a simulation of atmospheric transport from the master GCM and generate air mass fluxes
3. Compress the mass fluxes
4. Use these mass fluxes as input for the inverse system

This process has several steps that limit how fast these inversions can be produced. The initial release of ERA5 reanalysis, called ERA5T, lags behind real-time by five days. It gets consolidated after a couple of months, and this is the version that we usually collect at LSCE. The master GCM also needs time to run before the inversion can even be started. If steps 2 and 3 could be bypassed, and the information from a reanalysis directly used to drive the inverse system instead of having to use the master GCM, this process would be much faster. We could also avoid maintaining the full GCM within the SATINV team and focus more resources on the offline model. However, the conversion of wind fields to LMDZ-type air mass fluxes to replace the master GCM is not trivial, and our first try to do so with the regular version of the inverse system was a failure.

While the original inverse system used air mass fluxes as input for the dynamics, as described in Hourdin, 2015, the new version with DYNAMICO simply uses winds generated by the GCM as detailed in [chapter VII of Part III](#). The latter are also directly available from the ECMWF analysis and reanalysis archives, while the former requires a combination of the wind and the pressure fields on the sphere (Equation 2.5 of Hourdin, 2005). A version using the wind fields from the ECMWF Integrated Forecasting System (IFS) is already being developed, and the dynamics of the transport model is functional. Using the IFS provides real-time wind fields compared to the month-long delay of an ERA5 reanalysis, and these fields are available at very high resolution: up to ~ 9 km. The adaptation is less straightforward for the simulation of transport by convection and boundary-layer turbulence for which the input variables for the GCM subroutines used by the inverse system are also not directly available from the ECMWF analysis and reanalysis archives: the GCM routines have therefore been replaced by the IFS ones in the inverse system for that new version.

Importantly, this hybridization of the off-line transport model with IFS components opens the possibility of its connection to the IFS. The IFS data assimilation window for CO₂ and CH₄ could be extended beyond its current 12 or 24 hours at little cost by connecting the tangent-linear and adjoint transport models that are embedded within the computationally-expensive IFS, with our lighter and cheaper transport model. This would be an original building block for the future European anthropogenic CO₂ Monitoring and Verification Support capacity (Janssens-Maenhout et al., 2020).

2 Lessons on increasing resolution

While this work focused on a single analytical inverse system and its associated transport model in a GCM, the problems faced throughout it as well as some of their solutions can be of use to others looking to increase the resolution of their atmospheric model.

2.1 GPU: an opportunity not to be missed

In contrast to our expectations at the start of this work, the biggest change in the inverse system that made any other envisioned improvement technically possible was the implementation of GPU parallelization. This made the model orders of magnitude faster and still leaves room for optimization as discussed in the section above. Only a small engineering cost was necessary compared to the

gains in computational speed, making it a very efficient way of accelerating code. The choice to use OpenACC and a heterogeneous CPU/GPU parallelization instead of rewriting the code from scratch simplified the work and can easily be applied to any model already parallelized with OpenMP and coded in C, C++, or Fortran. This new parallelization method does not necessarily need to be all-encompassing, it can first be focused on only certain parts of the code that are already suited for GPU parallelization before refactoring the rest. GPUs are only going to become more prevalent in future supercomputing architectures and climate or atmospheric inverse models should be at the forefront of these innovations as they have always been in the past (Leiserson et al., 2020).

2.2 How to choose the grid?

Increasing the spatial resolution of an atmospheric model, however, is not only a question of which parallelization scheme to choose. Any kind of parallelization, whether it be on CPU or GPU, will only be useful at higher resolution if it scales efficiently when using more processors. This is in large part dependent on the grid chosen for the model. In our case, the regular latitude-longitude grid was quite limited, and the icosahedral grid of DYNAMICO was a clear improvement in terms of scaling, it will not be the limiting factor for a possible two-fold to four-fold increase in resolution of our inverse system in the future. It is easy to stay with the usual grid when increasing the resolution of a model since in the short term it is often sufficient, but despite the high engineering cost that a grid change entails, it is important to always look ahead and ask the question of whether or not the same grid will still be relevant a couple of years or even a decade from now. Many climate and inverse models have now implemented new grids that answer this exact problem, and all of their experiences can be leveraged by anyone wanting to improve their own models.

However, one somewhat disappointing lesson from our work is that the advantages of a new grid should not be overestimated either. We started this work hoping for great improvements in the representation of atmospheric transport near the poles but ended up not seeing any significant difference in that regard. This might however be improved in the future with further tuning of the models by the IPSL as part of the CMIP7 project. The expected gains in computational speed were also not quite what we expected. While we did see notable gains, most of it was due to the lower number of cells of the icosahedral grid compared to a regular grid at equivalent resolution. The grid itself does not necessarily make the computations any faster, and the higher number of sides of each hexagonal cell, together with the 3D representation of the variables, have a negative effect in particular for the inverse model.

The icosahedral grid introduced some important challenges when processing the inputs and the outputs of the model, even just to visualize them. For developers and model users that study these models in detail, enough public tools exist now that they can be overcome. This might however be a problem for end-users that simply want to exploit the results from global inversion products that are distributed on a non-cartesian grid. The need to invest in new tools or design a different workflow than for other products on classical latitude-longitude grids could negatively impact their adoption. These products might therefore need to be regridded to more common grids at the risk of some degradation in accuracy, at least until unstructured grids become more widespread.

2.3 Technical challenges

The results of this work show that there are no longer any technical limitations to a massive increase in the resolution of global inverse atmospheric models. Enough options exist in terms of parallelization schemes and grids such that most models should be able to find a solution to the specific challenges they face. The technical solutions presented throughout this manuscript can provide some insight into which ones are worth pursuing in priority depending on the desired outcome.

3 Could AI's reign on global atmospheric transport be over before it even started?

The use of machine learning (ML), often referred to as Artificial Intelligence (AI), has become ubiquitous in many scientific applications, promising to revolutionize the quality and speed of modeling through its data-driven approach. In atmospheric modeling, ML approaches have recently demonstrated competitive accuracy in various weather forecasting metrics, delivering forecasts at a computational cost orders of magnitude lower than traditional numerical weather prediction methods (Bouallègue et al., 2024). First results of deep learning approaches at the scale of an emission plume are also promising (Dumont Le Brazidec et al., 2024).

However, the adoption of ML by users has often been limited by an interpretability problem: it can look like a *black box* with no clear understanding of how ML makes its predictions. Various methods exist to help visualize and interpret the relationship between predictors and the final output of the models, but the most important meteorological predictors identified with such methods may sometimes overshadow other physically relevant ones (McGovern et al., 2019).

These ML models rely entirely on large, high-quality datasets for training. For weather forecasting, they have predominantly used ERA5 reanalysis data, which provides a high-resolution, continuous dataset of meteorological variables from 1940 to the present day, benefiting from the data assimilation of numerous observations. The use of ML for global atmospheric transport of greenhouse gases has been hindered by the absence of such high-quality datasets. This challenge complicates the potential for ML-driven inversions of CO₂, although initiatives like AI4Carbon (<https://ai4carbon.github.io/>) aim to address this issue.

The results from our work, and particularly the potential of a hybrid inverse system for CO₂ directly connected to the IFS could however somewhat steal the thunder of these new approaches. Indeed, it would exhibit the main qualities that make ML techniques attractive in the first place: leveraging GPUs to be much faster than previous inverse models and making use of high quality meteorological data via its coupling to the IFS.

Moreover, this approach does not encounter the interpretability issues associated with ML. It conserves mass with high accuracy and at local scale, a property which is at the root of global atmospheric inverse modelling but which may be challenging for ML, and it can be directly related to any atmospheric observation through its generation of full 3D tracer fields as a by-product of the inversion process. The challenge of finding a suitable dataset for training an ML model for atmospheric transport is completely bypassed, as the traditional technique of data assimilation of CO₂ observations provides a very rich and continuously expanding dataset of high quality.

A potential collaboration between our approach and ML could lie in the compression of the input meteorological dataset, as previously mentioned in [subsection 1.1.2](#). Autoencoders routinely used in deep learning for dimension reduction, for instance, could be an efficient alternative to, e.g., a Principal Component Analysis of the wind fields to reduce the volume of the input data. The idea here is to replace some of the reading time in the model by computing time to decode the input data, as computing time on GPU may be negligible.

4 Is increasing the resolution worth the pain?

This work started from the assumption that increasing the horizontal spatial resolution would greatly benefit the representation of atmospheric tracer transport and the subsequent estimation of CO₂ surface fluxes by our inverse model. This assumption however warrants a second look, firstly as a result of our scientific inquiry but also in a more fundamental way by questioning the very underlying concepts behind this assumption.

4.1 Modest benefits so far

The increase in horizontal resolution of our inverse system was entirely successful and even exceeded our expectations of how fast we managed to reach our first target of 1° resolution. However, while this was an undeniable technical success, the benefits of these resolution increases were less significant than initially expected. The representation of atmospheric transport in our inverse model did see a small benefit, but it was relatively minor compared to the computational cost increase required. We saw some differences in the estimates of natural surface fluxes of CO_2 with the high-resolution inversion but these were also relatively small.

Some of the motivation behind a resolution increase was also to make better use of satellite data, and the results on that front were also modest. We did see an improvement in a high northern latitude region where OCO-2 XCO_2 retrievals are scarce, but there was no significant difference at the global scale. It must be noted however that high-resolution inversions might benefit from a change in the averaging frequency of the assimilated satellite retrievals and in optimizing the correlation matrices.

4.2 A non-negligible ecological cost

One aspect of atmospheric modeling that is rarely discussed is the energy cost of the simulations and by extension their ecological impact. We often measure the total CPU hours consumed, and try to optimize it, but most of the time it comes from a concern about the monetary cost and time taken by these simulations. We rarely take into account the energy consumed by the supercomputers we use. Their impact depends on their efficiency but also on the energy mix that supplies their electricity which depends on the country. A study of the energy usage and subsequent carbon footprint of CMIP6 simulations estimated a total carbon footprint of around 1700 t CO_2 for all the experiments (Acosta et al., 2024). This figure only includes the *useful* experiments of direct scientific value. It does not take into account the numerous tests or failed experiments for example which can also have a significant cost. This is only one calculation of a subset of climate experiments, and it is very difficult to extrapolate this data to estimate the carbon footprint of the global greenhouse gas inversion modeling community. Nonetheless, with the spatial resolution of models increasing at a very fast pace their carbon footprint will also grow in the same manner. The main argument for carrying out these experiments is that their contribution to climate science can help us combat climate change and reduce greenhouse gas emissions, but this surely cannot be done at any cost either. We must ask ourselves if it is even possible to truly weigh the impact of these simulations against their potential future benefit. Some work on this subject must be done by the climate science community either way, as this issue will only become more pressing with time.

4.3 A paradigm change in the future?

One advantage of a resolution increase of inverse models that is evident but hard to quantify the benefit of, is the fact that the smaller grid size will now allow global inversions to produce data on the scale of small countries. This scale was until now reserved for limited-area models which come with their own drawbacks. Such products will be very useful in creating the national greenhouse gas inventories required by the United Nations Framework Convention on Climate Change as part of the Paris Agreement or in monitoring their accuracy and evolution in time (Chevallier, 2021, Deng et al., 2022).

Several studies seem to indicate that the range of horizontal resolution encompassed by this work (individual cell size at the Equator increasing from $417 \text{ km} \times 211 \text{ km}$ to $157 \text{ km} \times 78 \text{ km}$) might be the resolutions that provide the least benefit to atmospheric transport and inverse models. Studies such as Agustí-Panareda et al., 2019 suggest that improvement in the representation of atmospheric CO_2 variability would be more significant at even higher resolutions than those studied in this work, and that the relationship between the two is not linear. The quality of the representation of this synoptic variability can have a strong impact on the data assimilation process of inversions.

At even higher resolutions, physical processes that were previously parametrized at the sub-grid scale can start to be directly resolved. Low clouds, for example, are generated by 10 m to 100 m wide turbulent updrafts and have a strong impact on surface temperatures and atmospheric transport. While this scale may still seem far away, significant work needs to be invested now into improving the modeling of these processes, at the risk otherwise of seeing our understanding of them outpaced by the speed of technical progress (Schneider et al., 2017). Another assumption for the parametrization of our atmospheric inversions is that the errors of our assimilated observations are not correlated in space. Our satellite observations are actually 10-second averages, which correspond to a distance of 67.5 km along the satellite track. As resolution increases, if bins are made smaller so that super-observations do not spread widely across multiple grid cells, error correlations will increase because retrieval errors are significantly correlated at fine scales (Baker et al., 2022): the initial assumption will not hold true anymore. If bins are kept as they are and the observation operator is adapted to account for the contribution from several pixels, we will not benefit from many details of the CO₂ column variability revealed by the retrievals despite their error correlations (Zheng et al., 2020). Finally, the current averaging of the input meteorological variables over three hours may be too coarse for high resolutions, but reducing it also implies enlarging the files and the time needed by the model to read them.

As global inverse models reach new and yet unexplored spatial resolutions, all of these aspects must be critically evaluated: what worked for yesterday's models cannot be assumed to still work tomorrow.



Bibliography

- Acosta, Mario C. et al. (Apr. 19, 2024). “The computational and energy cost of simulation and storage for climate science: lessons from CMIP6”. In: *Geoscientific Model Development* 17.8, pp. 3081–3098. DOI: [10.5194/gmd-17-3081-2024](https://doi.org/10.5194/gmd-17-3081-2024).
- Agustí-Panareda, Anna et al. (June 4, 2019). “Modelling CO₂ weather – why horizontal resolution matters”. In: *Atmospheric Chemistry and Physics* 19.11, pp. 7347–7376. DOI: [10.5194/acp-19-7347-2019](https://doi.org/10.5194/acp-19-7347-2019).
- Andrew, Robbie M. (June 29, 2020). “A comparison of estimates of global carbon dioxide emissions from fossil carbon sources”. In: *Earth System Science Data* 12.2, pp. 1437–1465. DOI: [10.5194/essd-12-1437-2020](https://doi.org/10.5194/essd-12-1437-2020).
- Baker, David F. et al. (Jan. 26, 2022). “A new exponentially decaying error correlation model for assimilating OCO-2 column-average CO₂ data using a length scale computed from airborne lidar measurements”. In: *Geoscientific Model Development* 15.2, pp. 649–668. DOI: [10.5194/gmd-15-649-2022](https://doi.org/10.5194/gmd-15-649-2022).
- Balaji, Venkatramani et al. (Jan. 2, 2017). “CPMIP: measurements of real computational performance of Earth system models in CMIP6”. In: *Geoscientific Model Development* 10.1, pp. 19–34. DOI: [10.5194/gmd-10-19-2017](https://doi.org/10.5194/gmd-10-19-2017).
- Ballantyne, A. P. et al. (Aug. 2012). “Increase in observed net carbon dioxide uptake by land and oceans during the past 50 years”. In: *Nature* 488.7409, pp. 70–72. DOI: [10.1038/nature11299](https://doi.org/10.1038/nature11299).
- Basu, S. et al. (Sept. 3, 2013). “Global CO₂ fluxes estimated from GOSAT retrievals of total column CO₂”. In: *Atmospheric Chemistry and Physics* 13.17, pp. 8695–8717. DOI: [10.5194/acp-13-8695-2013](https://doi.org/10.5194/acp-13-8695-2013).
- Basu, Sourish et al. (May 24, 2018). “The impact of transport model differences on CO₂ surface flux estimates from OCO-2 retrievals of column average CO₂”. In: *Atmospheric Chemistry and Physics* 18.10, pp. 7189–7215. DOI: [10.5194/acp-18-7189-2018](https://doi.org/10.5194/acp-18-7189-2018).
- Bauer, Peter et al. (Feb. 2021). “The digital revolution of Earth-system science”. In: *Nature Computational Science* 1.2, pp. 104–113. DOI: [10.1038/s43588-021-00023-0](https://doi.org/10.1038/s43588-021-00023-0).
- Berchet, Antoine et al. (Aug. 26, 2021). “The Community Inversion Framework v1.0: a unified system for atmospheric inversion studies”. In: *Geoscientific Model Development* 14.8, pp. 5331–5354. DOI: [10.5194/gmd-14-5331-2021](https://doi.org/10.5194/gmd-14-5331-2021).
- Boesch, Hartmut et al. (Feb. 2011). “Global Characterization of CO₂ Column Retrievals from Shortwave-Infrared Satellite Observations of the Orbiting Carbon Observatory-2 Mission”. In: *Remote Sensing* 3.2, pp. 270–304. DOI: [10.3390/rs3020270](https://doi.org/10.3390/rs3020270).
- Bouallègue, Zied Ben et al. (June 6, 2024). “The Rise of Data-Driven Weather Forecasting: A First Statistical Assessment of Machine Learning–Based Weather Forecasts in an Operational-

- Like Context”. In: *Bulletin of the American Meteorological Society* 105.6, E864–E883. DOI: [10.1175/BAMS-D-23-0162.1](https://doi.org/10.1175/BAMS-D-23-0162.1).
- Bousquet, Philippe et al. (Nov. 17, 2000). “Regional Changes in Carbon Dioxide Fluxes of Land and Oceans Since 1980”. In: *Science* 290.5495, pp. 1342–1346. DOI: [10.1126/science.290.5495.1342](https://doi.org/10.1126/science.290.5495.1342).
- Bréon, F. M. et al. (Feb. 18, 2015). “An attempt at estimating Paris area CO₂ emissions from atmospheric concentration measurements”. In: *Atmospheric Chemistry and Physics* 15.4, pp. 1707–1724. DOI: [10.5194/acp-15-1707-2015](https://doi.org/10.5194/acp-15-1707-2015).
- Brioude, J. et al. (2011). “Top-down estimate of anthropogenic emission inventories and their interannual variability in Houston using a mesoscale inverse modeling technique”. In: *Journal of Geophysical Research: Atmospheres* 116 (D20). DOI: [10.1029/2011JD016215](https://doi.org/10.1029/2011JD016215).
- Byrne, Brendan et al. (Mar. 7, 2023). “National CO₂ budgets (2015–2020) inferred from atmospheric CO₂ observations in support of the global stocktake”. In: *Earth System Science Data* 15.2, pp. 963–1004. DOI: [10.5194/essd-15-963-2023](https://doi.org/10.5194/essd-15-963-2023).
- Chandra, Naveen et al. (July 18, 2022). “Estimated regional CO₂ flux and uncertainty based on an ensemble of atmospheric CO₂ inversions”. In: *Atmospheric Chemistry and Physics* 22.14, pp. 9215–9243. DOI: [10.5194/acp-22-9215-2022](https://doi.org/10.5194/acp-22-9215-2022).
- Chatterjee, Abhishek et al. (2012). “Toward reliable ensemble Kalman filter estimates of CO₂ fluxes”. In: *Journal of Geophysical Research: Atmospheres* 117 (D22). DOI: [10.1029/2012JD018176](https://doi.org/10.1029/2012JD018176).
- Cheruy, Frédérique et al. (2020). “Improved Near-Surface Continental Climate in IPSL-CM6A-LR by Combined Evolutions of Atmospheric and Land Surface Physics”. In: *Journal of Advances in Modeling Earth Systems* 12.10, e2019MS002005. DOI: [10.1029/2019MS002005](https://doi.org/10.1029/2019MS002005).
- Chevallier, F. (June 7, 2013). “On the parallelization of atmospheric inversions of CO₂ surface fluxes within a variational framework”. In: *Geoscientific Model Development* 6.3, pp. 783–790. DOI: [10.5194/gmd-6-783-2013](https://doi.org/10.5194/gmd-6-783-2013).
- Chevallier, F. et al. (2005). “Inferring CO₂ sources and sinks from satellite observations: Method and application to TOVS data”. In: *Journal of Geophysical Research: Atmospheres* 110 (D24). DOI: [10.1029/2005JD006390](https://doi.org/10.1029/2005JD006390).
- Chevallier, F. et al. (2010). “CO₂ surface fluxes at grid point scale estimated from a global 21 year reanalysis of atmospheric measurements”. In: *Journal of Geophysical Research: Atmospheres* 115 (D21). DOI: [10.1029/2010JD013887](https://doi.org/10.1029/2010JD013887).
- Chevallier, Frédéric (2021). “Fluxes of Carbon Dioxide From Managed Ecosystems Estimated by National Inventories Compared to Atmospheric Inverse Modeling”. In: *Geophysical Research Letters* 48.15, e2021GL093565. DOI: [10.1029/2021GL093565](https://doi.org/10.1029/2021GL093565).
- Chevallier, Frédéric, François-Marie Bréon, and Peter J. Rayner (May 8, 2007). “Contribution of the Orbiting Carbon Observatory to the estimation of CO₂ sources and sinks: Theoretical study in a variational data assimilation framework”. In: *Journal of Geophysical Research* 112 (D9), p. D09307. DOI: [10.1029/2006JD007375](https://doi.org/10.1029/2006JD007375).
- Chevallier, Frédéric et al. (July 2006). “On the assignment of prior errors in Bayesian inversions of CO₂ surface fluxes”. In: *Geophysical Research Letters* 33.13, 2006GL026496. DOI: [10.1029/2006GL026496](https://doi.org/10.1029/2006GL026496).
- Chevallier, Frédéric et al. (Nov. 26, 2019). “Objective evaluation of surface- and satellite-driven carbon dioxide atmospheric inversions”. In: *Atmospheric Chemistry and Physics* 19.22, pp. 14233–14251. DOI: [10.5194/acp-19-14233-2019](https://doi.org/10.5194/acp-19-14233-2019).
- Chevallier, Frédéric et al. (Mar. 16, 2023). “Toward High-Resolution Global Atmospheric Inverse Modeling Using Graphics Accelerators”. In: *Geophysical Research Letters* 50.5, e2022GL102135. DOI: [10.1029/2022GL102135](https://doi.org/10.1029/2022GL102135).
- Chmielowiec, Philip et al. (Sept. 7, 2023). *UXARRAY/uxarray: v2023.09.0*. DOI: [10.5281/zenodo.8326146](https://doi.org/10.5281/zenodo.8326146).

- Ciais, P. et al. (Nov. 1, 2010). “Atmospheric inversions for estimating CO₂ fluxes: methods and perspectives”. In: *Climatic Change* 103.1, pp. 69–92. DOI: [10.1007/s10584-010-9909-3](https://doi.org/10.1007/s10584-010-9909-3).
- Crisp, David et al. (Jan. 5, 2017). “The on-orbit performance of the Orbiting Carbon Observatory-2 (OCO-2) instrument and its radiometrically calibrated products”. In: *Atmospheric Measurement Techniques* 10.1, pp. 59–81. DOI: [10.5194/amt-10-59-2017](https://doi.org/10.5194/amt-10-59-2017).
- Dai, Shifeng et al. (Sept. 13, 2022). “Metalliferous Coals of Cretaceous Age: A Review”. In: *Minerals* 12.9, p. 1154. DOI: [10.3390/min12091154](https://doi.org/10.3390/min12091154).
- Deng, Zhu et al. (Apr. 11, 2022). “Comparing national greenhouse gas budgets reported in UNFCCC inventories against atmospheric inversions”. In: *Earth System Science Data* 14.4, pp. 1639–1675. DOI: [10.5194/essd-14-1639-2022](https://doi.org/10.5194/essd-14-1639-2022).
- Dogniaux, Matthieu (Dec. 13, 2021). “Suivi de la concentration atmosphérique de CO₂ par satellite : performances et sensibilités des prochains concepts d’observation dans le proche infrarouge”. These de doctorat. Institut polytechnique de Paris.
- Dubos, T. et al. (Oct. 7, 2015). “DYNAMICO-1.0, an icosahedral hydrostatic dynamical core designed for consistency and versatility”. In: *Geoscientific Model Development* 8.10, pp. 3131–3150. DOI: [10.5194/gmd-8-3131-2015](https://doi.org/10.5194/gmd-8-3131-2015).
- Dubos, Thomas and Marine Tort (Oct. 1, 2014). “Equations of Atmospheric Motion in Non-Eulerian Vertical Coordinates: Vector-Invariant Form and Quasi-Hamiltonian Formulation”. In: *Monthly Weather Review* 142.10, pp. 3860–3880. DOI: [10.1175/MWR-D-14-00069.1](https://doi.org/10.1175/MWR-D-14-00069.1).
- Dumont Le Brazidec, Joffrey et al. (Mar. 5, 2024). “Deep learning applied to CO₂ power plant emissions quantification using simulated satellite images”. In: *Geoscientific Model Development* 17.5, pp. 1995–2014. DOI: [10.5194/gmd-17-1995-2024](https://doi.org/10.5194/gmd-17-1995-2024).
- Emanuel, Kerry A. (Nov. 1, 1991). “A Scheme for Representing Cumulus Convection in Large-Scale Models”. In: *Journal of the Atmospheric Sciences* 48.21, pp. 2313–2329. DOI: [10.1175/1520-0469\(1991\)048<2313:ASFRCC>2.0.CO;2](https://doi.org/10.1175/1520-0469(1991)048<2313:ASFRCC>2.0.CO;2).
- Engwirda, Darren (Dec. 15, 2018). “Generalised primal-dual grids for unstructured co-volume schemes”. In: *Journal of Computational Physics* 375, pp. 155–176. DOI: [10.1016/j.jcp.2018.07.025](https://doi.org/10.1016/j.jcp.2018.07.025).
- Erb, Karl-Heinz et al. (Oct. 2013). “Bias in the attribution of forest carbon sinks”. In: *Nature Climate Change* 3.10, pp. 854–856. DOI: [10.1038/nclimate2004](https://doi.org/10.1038/nclimate2004).
- Feng, L. et al. (Feb. 4, 2016). “Estimates of European uptake of CO₂ inferred from GOSAT X_{CO₂} retrievals: sensitivity to measurement bias inside and outside Europe”. In: *Atmospheric Chemistry and Physics* 16.3, pp. 1289–1302. DOI: [10.5194/acp-16-1289-2016](https://doi.org/10.5194/acp-16-1289-2016).
- Fisher, M. (1998). “Minimization algorithms for variational data assimilation”. In: *Proc. of ECMWF Seminar on Recent Development in Numerical Methods for Atmospheric Modelling (7-11 September 1998, Reading, UK)*. ECMWF.
- Folberth, G. A. et al. (June 21, 2006). “Interactive chemistry in the Laboratoire de Météorologie Dynamique general circulation model: model description and impact analysis of biogenic hydrocarbons on tropospheric chemistry”. In: *Atmospheric Chemistry and Physics* 6.8, pp. 2273–2319. DOI: [10.5194/acp-6-2273-2006](https://doi.org/10.5194/acp-6-2273-2006).
- Frey, Matthias et al. (Mar. 11, 2019). “Building the Collaborative Carbon Column Observing Network (COCCON): long-term stability and ensemble performance of the EM27/SUN Fourier transform spectrometer”. In: *Atmospheric Measurement Techniques* 12.3, pp. 1513–1530. DOI: [10.5194/amt-12-1513-2019](https://doi.org/10.5194/amt-12-1513-2019).
- Friedlingstein, Pierre et al. (Nov. 11, 2022). “Global Carbon Budget 2022”. In: *Earth System Science Data* 14.11, pp. 4811–4900. DOI: [10.5194/essd-14-4811-2022](https://doi.org/10.5194/essd-14-4811-2022).
- Fuhrer, Oliver et al. (May 2, 2018). “Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0”. In: *Geoscientific Model Development* 11.4, pp. 1665–1681. DOI: [10.5194/gmd-11-1665-2018](https://doi.org/10.5194/gmd-11-1665-2018).

- Gasser, Thomas et al. (Aug. 13, 2020). “Historical CO₂ emissions from land use and land cover change and their uncertainty”. In: *Biogeosciences* 17.15, pp. 4075–4101. DOI: [10.5194/bg-17-4075-2020](https://doi.org/10.5194/bg-17-4075-2020).
- Gassmann, Almut (2013). “A global hexagonal C-grid non-hydrostatic dynamical core (ICON-IAP) designed for energetic consistency”. In: *Quarterly Journal of the Royal Meteorological Society* 139.670, pp. 152–175. DOI: [10.1002/qj.1960](https://doi.org/10.1002/qj.1960).
- Gilbert, Jean Charles and Claude Lemaréchal (Aug. 1, 1989). “Some numerical experiments with variable-storage quasi-Newton algorithms”. In: *Mathematical Programming* 45.1, pp. 407–435. DOI: [10.1007/BF01589113](https://doi.org/10.1007/BF01589113).
- Giorgetta, Marco A. et al. (Sept. 16, 2022). “The ICON-A model for direct QBO simulations on GPUs (version icon-cscs:baf28a514)”. In: *Geoscientific Model Development* 15.18, pp. 6985–7016. DOI: [10.5194/gmd-15-6985-2022](https://doi.org/10.5194/gmd-15-6985-2022).
- Hansis, Eberhard, Steven J. Davis, and Julia Pongratz (2015). “Relevance of methodological choices for accounting of land use change carbon fluxes”. In: *Global Biogeochemical Cycles* 29.8, pp. 1230–1246. DOI: [10.1002/2014GB004997](https://doi.org/10.1002/2014GB004997).
- Hauck, Judith et al. (2020). “Consistency and Challenges in the Ocean Carbon Sink Estimate for the Global Carbon Budget”. In: *Frontiers in Marine Science* 7.
- Hauglustaine, D. A. et al. (2004). “Interactive chemistry in the Laboratoire de Météorologie Dynamique general circulation model: Description and background tropospheric chemistry evaluation”. In: *Journal of Geophysical Research: Atmospheres* 109 (D4). DOI: [10.1029/2003JD003957](https://doi.org/10.1029/2003JD003957).
- Hedelius, Jacob K. et al. (2017). “Emissions and topographic effects on column CO₂ variations, with a focus on the Southern California Megacity”. In: *Journal of Geophysical Research: Atmospheres* 122.13, pp. 7200–7215. DOI: [10.1002/2017JD026455](https://doi.org/10.1002/2017JD026455).
- Heikes, Ross P., David A. Randall, and Celal S. Konor (Dec. 1, 2013). “Optimized Icosahedral Grids: Performance of Finite-Difference Operators and Multigrid Solver”. In: *Monthly Weather Review* 141.12, pp. 4450–4469. DOI: [10.1175/MWR-D-12-00236.1](https://doi.org/10.1175/MWR-D-12-00236.1).
- Hollingsworth, A. et al. (Aug. 1, 2008). “TOWARD A MONITORING AND FORECASTING SYSTEM FOR ATMOSPHERIC COMPOSITION: The GEMS Project”. In: *Bulletin of the American Meteorological Society* 89.8, pp. 1147–1164. DOI: [10.1175/2008BAMS2355.1](https://doi.org/10.1175/2008BAMS2355.1).
- Hong, Xinhua et al. (2022). “Retrieval of Global Carbon Dioxide From TanSat Satellite and Comprehensive Validation With TCCON Measurements and Satellite Observations”. In: *IEEE Transactions on Geoscience and Remote Sensing* 60, pp. 1–16. DOI: [10.1109/TGRS.2021.3066623](https://doi.org/10.1109/TGRS.2021.3066623).
- Houghton, R. A. and Alexander A. Nassikas (2017). “Global and regional fluxes of carbon from land use and land cover change 1850–2015”. In: *Global Biogeochemical Cycles* 31.3, pp. 456–472. DOI: [10.1002/2016GB005546](https://doi.org/10.1002/2016GB005546).
- Hourdin, Frederic (2015). “Représentation du transport direct et inverse dans les modèles globaux de climat et étude des couplages entre composition et dynamique atmosphérique sur Titan.” HDR defense.
- Hourdin, Frédéric et al. (Dec. 1, 2006). “The LMDZ4 general circulation model: climate performance and sensitivity to parametrized physics with emphasis on tropical convection”. In: *Climate Dynamics* 27.7, pp. 787–813. DOI: [10.1007/s00382-006-0158-0](https://doi.org/10.1007/s00382-006-0158-0).
- Hourdin, Frédéric et al. (May 1, 2013a). “Impact of the LMDZ atmospheric grid configuration on the climate and sensitivity of the IPSL-CM5A coupled model”. In: *Climate Dynamics* 40.9, pp. 2167–2192. DOI: [10.1007/s00382-012-1411-3](https://doi.org/10.1007/s00382-012-1411-3).
- Hourdin, Frédéric et al. (May 1, 2013b). “LMDZ5B: the atmospheric component of the IPSL climate model with revisited parameterizations for clouds and convection”. In: *Climate Dynamics* 40.9, pp. 2193–2222. DOI: [10.1007/s00382-012-1343-y](https://doi.org/10.1007/s00382-012-1343-y).

- Hourdin, Frédéric et al. (2020). “LMDZ6A: The Atmospheric Component of the IPSL Climate Model With Improved and Better Tuned Physics”. In: *Journal of Advances in Modeling Earth Systems* 12.7, e2019MS001892. DOI: [10.1029/2019MS001892](https://doi.org/10.1029/2019MS001892).
- Houweling, S. et al. (Nov. 8, 2005). “Evidence of systematic errors in SCIAMACHY-observed CO₂ due to aerosols”. In: *Atmospheric Chemistry and Physics* 5.11, pp. 3003–3013. DOI: [10.5194/acp-5-3003-2005](https://doi.org/10.5194/acp-5-3003-2005).
- Houweling, S. et al. (2015). “An intercomparison of inverse models for estimating sources and sinks of CO₂ using GOSAT measurements”. In: *Journal of Geophysical Research: Atmospheres* 120.10, pp. 5253–5266. DOI: [10.1002/2014JD022962](https://doi.org/10.1002/2014JD022962).
- Intergovernmental Panel On Climate Change (July 6, 2023). *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. 1st ed. Cambridge University Press. DOI: [10.1017/9781009157896](https://doi.org/10.1017/9781009157896).
- Janssens-Maenhout, G. et al. (Sept. 2, 2020). “Toward an Operational Anthropogenic CO₂ Emissions Monitoring and Verification Support Capacity”. In: *Bulletin of the American Meteorological Society* 101.8, E1439–E1451. DOI: [10.1175/BAMS-D-19-0017.1](https://doi.org/10.1175/BAMS-D-19-0017.1).
- Jin, Zhe et al. (Mar. 1, 2023). “Constraint of satellite CO₂ retrieval on the global carbon cycle from a Chinese atmospheric inversion system”. In: *Science China Earth Sciences* 66.3, pp. 609–618. DOI: [10.1007/s11430-022-1036-7](https://doi.org/10.1007/s11430-022-1036-7).
- Karion, Anna et al. (Nov. 1, 2010). “AirCore: An Innovative Atmospheric Sampling System”. In: *Journal of Atmospheric and Oceanic Technology* 27.11, pp. 1839–1853. DOI: [10.1175/2010JTECHA1448.1](https://doi.org/10.1175/2010JTECHA1448.1).
- Kasahara, AKIRA (Jan. 1, 1977). “Computational Aspects of Numerical Models for Weather Prediction and Climate Simulation”. In: *Methods in Computational Physics: Advances in Research and Applications*. Ed. by JULIUS Chang. Vol. 17. General Circulation Models of the Atmosphere. Elsevier, pp. 1–66. DOI: [10.1016/B978-0-12-460817-7.50006-9](https://doi.org/10.1016/B978-0-12-460817-7.50006-9).
- Keely, William R. et al. (Nov. 29, 2023). “A nonlinear data-driven approach to bias correction of XCO₂ for NASA’s OCO-2 ACOS version 10”. In: *Atmospheric Measurement Techniques* 16.23, pp. 5725–5748. DOI: [10.5194/amt-16-5725-2023](https://doi.org/10.5194/amt-16-5725-2023).
- Koch, Alexander et al. (Mar. 1, 2019). “Earth system impacts of the European arrival and Great Dying in the Americas after 1492”. In: *Quaternary Science Reviews* 207, pp. 13–36. DOI: [10.1016/j.quascirev.2018.12.004](https://doi.org/10.1016/j.quascirev.2018.12.004).
- Kong, Yawen et al. (Aug. 25, 2022). “Global and regional carbon budget for 2015–2020 inferred from OCO-2 based on an ensemble Kalman filter coupled with GEOS-Chem”. In: *Atmospheric Chemistry and Physics* 22.16, pp. 10769–10788. DOI: [10.5194/acp-22-10769-2022](https://doi.org/10.5194/acp-22-10769-2022).
- Kravitz, B. et al. (Oct. 27, 2015). “The Geoengineering Model Intercomparison Project Phase 6 (GeoMIP6): simulation design and preliminary results”. In: *Geoscientific Model Development* 8.10, pp. 3379–3392. DOI: [10.5194/gmd-8-3379-2015](https://doi.org/10.5194/gmd-8-3379-2015).
- Krinner, G. et al. (2005). “A dynamic global vegetation model for studies of the coupled atmosphere-biosphere system”. In: *Global Biogeochemical Cycles* 19.1. DOI: [10.1029/2003GB002199](https://doi.org/10.1029/2003GB002199).
- Kritsikis, Evaggelos et al. (Jan. 30, 2017). “Conservative interpolation between general spherical meshes”. In: *Geoscientific Model Development* 10.1, pp. 425–431. DOI: [10.5194/gmd-10-425-2017](https://doi.org/10.5194/gmd-10-425-2017).
- Lacis, Andrew A. et al. (Oct. 15, 2010). “Atmospheric CO₂: Principal Control Knob Governing Earth’s Temperature”. In: *Science* 330.6002, pp. 356–359. DOI: [10.1126/science.1190653](https://doi.org/10.1126/science.1190653).
- Lan, Xin, Pieter Tans, and Thoning Kirk (June 2024a). *Trends of CH₄ and N₂O growth*. URL: https://gml.noaa.gov/ccgg/trends/gl_data.html (visited on 06/18/2024).
- (June 2024b). *Trends of CO₂ growth*. URL: https://gml.noaa.gov/ccgg/trends/gl_data.html (visited on 06/18/2024).

- Leiserson, Charles E. et al. (June 5, 2020). “There’s plenty of room at the Top: What will drive computer performance after Moore’s law?” In: *Science* 368.6495, eaam9744. DOI: [10.1126/science.aam9744](https://doi.org/10.1126/science.aam9744).
- Lian, Jinghui et al. (Jan. 25, 2024). *Development and deployment of a mid-cost CO₂ sensor monitoring network to support atmospheric inverse modeling for quantifying urban CO₂ emissions in Paris*. DOI: [10.5194/egusphere-2024-125](https://doi.org/10.5194/egusphere-2024-125).
- Lin, Xin et al. (July 6, 2018). “Simulating CH₄ and CO₂ over South and East Asia using the zoomed chemistry transport model LMDz-INCA”. In: *Atmospheric Chemistry and Physics* 18.13, pp. 9475–9497. DOI: [10.5194/acp-18-9475-2018](https://doi.org/10.5194/acp-18-9475-2018).
- Liu, Junjie et al. (Feb. 10, 2021). “Carbon Monitoring System Flux Net Biosphere Exchange 2020 (CMS-Flux NBE 2020)”. In: *Earth System Science Data* 13.2, pp. 299–330. DOI: [10.5194/essd-13-299-2021](https://doi.org/10.5194/essd-13-299-2021).
- Lüthi, Dieter et al. (May 2008). “High-resolution carbon dioxide concentration record 650,000–800,000 years before present”. In: *Nature* 453.7193, pp. 379–382. DOI: [10.1038/nature06949](https://doi.org/10.1038/nature06949).
- Machida, T. et al. (Oct. 1, 2008). “Worldwide Measurements of Atmospheric CO₂ and Other Trace Gas Species Using Commercial Airlines”. In: *Journal of Atmospheric and Oceanic Technology* 25.10, pp. 1744–1754. DOI: [10.1175/2008JTECHA1082.1](https://doi.org/10.1175/2008JTECHA1082.1).
- Malardel, S. et al. (2016). “A new grid for the IFS”. In: DOI: [10.21957/ZWDU9U5I](https://doi.org/10.21957/ZWDU9U5I).
- Manabe, Syukuro and Kirk Bryan (July 1, 1969). “Climate Calculations with a Combined Ocean-Atmosphere Model”. In: *Journal of the Atmospheric Sciences* 26.4, pp. 786–789. DOI: [10.1175/1520-0469\(1969\)026<0786:CCWACO>2.0.CO;2](https://doi.org/10.1175/1520-0469(1969)026<0786:CCWACO>2.0.CO;2).
- McGovern, Amy et al. (Nov. 1, 2019). “Making the Black Box More Transparent: Understanding the Physical Implications of Machine Learning”. In: *Bulletin of the American Meteorological Society* 100.11, pp. 2175–2199. DOI: [10.1175/BAMS-D-18-0195.1](https://doi.org/10.1175/BAMS-D-18-0195.1).
- Membrive, Olivier et al. (June 12, 2017). “AirCore-HR: a high-resolution column sampling to enhance the vertical description of CH₄ and CO₂”. In: *Atmospheric Measurement Techniques* 10.6, pp. 2163–2181. DOI: [10.5194/amt-10-2163-2017](https://doi.org/10.5194/amt-10-2163-2017).
- Mlawer, Eli J. et al. (1997). “Radiative transfer for inhomogeneous atmospheres: RRTM, a validated correlated-k model for the longwave”. In: *Journal of Geophysical Research: Atmospheres* 102 (D14), pp. 16663–16682. DOI: [10.1029/97JD00237](https://doi.org/10.1029/97JD00237).
- Monteil, Guillaume et al. (Oct. 26, 2020). “The regional European atmospheric transport inversion comparison, EUROCOM: first results on European-wide terrestrial carbon fluxes for the period 2006–2015”. In: *Atmospheric Chemistry and Physics* 20.20, pp. 12063–12091. DOI: [10.5194/acp-20-12063-2020](https://doi.org/10.5194/acp-20-12063-2020).
- Navarro, T., G. Schubert, and S. Lebonnois (July 2018). “Atmospheric mountain wave generation on Venus and its influence on the solid planet’s rotation rate”. In: *Nature Geoscience* 11.7, pp. 487–491. DOI: [10.1038/s41561-018-0157-x](https://doi.org/10.1038/s41561-018-0157-x).
- Niwa, Yosuke et al. (Mar. 17, 2017). “A 4D-Var inversion system based on the icosahedral grid model (NICAM-TM 4D-Var v1.0) – Part 1: Offline forward and adjoint transport models”. In: *Geoscientific Model Development* 10.3, pp. 1157–1174. DOI: [10.5194/gmd-10-1157-2017](https://doi.org/10.5194/gmd-10-1157-2017).
- Niwa, Yosuke et al. (Dec. 2022). “Toward a long-term atmospheric CO₂ inversion for elucidating natural carbon fluxes: technical notes of NISMON-CO₂ v2021.1”. In: *Progress in Earth and Planetary Science* 9.1, pp. 1–19. DOI: [10.1186/s40645-022-00502-6](https://doi.org/10.1186/s40645-022-00502-6).
- O’Brien, D. M. and P. J. Rayner (2002). “Global observations of the carbon budget, 2, CO₂ column from differential absorption of reflected sunlight in the 1.61 μm band of CO₂”. In: *Journal of Geophysical Research: Atmospheres* 107 (D18), ACH 6–1–ACH 6–16. DOI: [10.1029/2001JD000617](https://doi.org/10.1029/2001JD000617).
- Petersen, A. K. et al. (2008). “First ground-based FTIR observations of the seasonal variation of carbon monoxide in the tropics”. In: *Geophysical Research Letters* 35.3. DOI: [10.1029/2007GL031393](https://doi.org/10.1029/2007GL031393).

- Pfeil, B. et al. (Apr. 4, 2013). “A uniform, quality controlled Surface Ocean CO₂ Atlas (SOCAT)”. In: *Earth System Science Data* 5.1, pp. 125–143. DOI: [10.5194/essd-5-125-2013](https://doi.org/10.5194/essd-5-125-2013).
- Pison, I. et al. (July 29, 2009). “Multi-species inversion of CH₄, CO and H₂ emissions from surface measurements”. In: *Atmospheric Chemistry and Physics* 9.14, pp. 5281–5297. DOI: [10.5194/acp-9-5281-2009](https://doi.org/10.5194/acp-9-5281-2009).
- Pisso, Ignacio et al. (Dec. 2, 2019). “The Lagrangian particle dispersion model FLEXPART version 10.4”. In: *Geoscientific Model Development* 12.12, pp. 4955–4997. DOI: [10.5194/gmd-12-4955-2019](https://doi.org/10.5194/gmd-12-4955-2019).
- Pottier, Alizée et al. (July 15, 2017). “Unraveling the martian water cycle with high-resolution global climate simulations”. In: *Icarus* 291, pp. 82–106. DOI: [10.1016/j.icarus.2017.02.016](https://doi.org/10.1016/j.icarus.2017.02.016).
- Remaud, Marine et al. (2018). “On the impact of recent developments of the LMDz atmospheric general circulation model on the simulation of CO₂ transport”. In: *Geosci. Model Dev.*, p. 25.
- Richter, Jadwiga H. et al. (2017). “Stratospheric Dynamical Response and Ozone Feedbacks in the Presence of SO₂ Injections”. In: *Journal of Geophysical Research: Atmospheres* 122.23, pp. 12, 557–12, 573. DOI: [10.1002/2017JD026912](https://doi.org/10.1002/2017JD026912).
- Rödenbeck, C. et al. (Oct. 8, 2018). “History of El Niño impacts on the global carbon cycle 1957–2017: a quantification from atmospheric CO₂ data”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 373.1760, p. 20170303. DOI: [10.1098/rstb.2017.0303](https://doi.org/10.1098/rstb.2017.0303).
- Rosnay, P. de et al. (2002). “Impact of a physically based soil water flow and soil-plant interaction representation for modeling large-scale land surface processes”. In: *Journal of Geophysical Research: Atmospheres* 107 (D11), ACL 3–1–ACL 3–19. DOI: [10.1029/2001JD000634](https://doi.org/10.1029/2001JD000634).
- Schank, Jeffrey and Charles Twardy (2009). “Mathematical Models”. In: *The Cambridge History of Science: Volume 6: The Modern Biological and Earth Sciences*. Ed. by John V. Pickstone and Peter J. Bowler. Vol. 6. The Cambridge History of Science. Cambridge: Cambridge University Press, pp. 416–431. DOI: [10.1017/CHOL9780521572019.023](https://doi.org/10.1017/CHOL9780521572019.023).
- Schär, Christoph et al. (May 1, 2020). “Kilometer-Scale Climate Models: Prospects and Challenges”. In: *Bulletin of the American Meteorological Society* 101.5, E567–E587. DOI: [10.1175/BAMS-D-18-0167.1](https://doi.org/10.1175/BAMS-D-18-0167.1).
- Schmidt, Gavin A. et al. (Oct. 16, 2010). “Attribution of the present-day total greenhouse effect”. In: *Journal of Geophysical Research* 115 (D20), p. D20106. DOI: [10.1029/2010JD014287](https://doi.org/10.1029/2010JD014287).
- Schneider, Tapio et al. (Jan. 2017). “Climate goals and computing the future of clouds”. In: *Nature Climate Change* 7.1, pp. 3–5. DOI: [10.1038/nclimate3190](https://doi.org/10.1038/nclimate3190).
- Schuh, Andrew E. et al. (May 8, 2019). “Quantifying the Impact of Atmospheric Transport Uncertainty on CO₂ Surface Flux Estimates”. In: *Global Biogeochemical Cycles* 33.4, pp. 484–500. DOI: [10.1029/2018GB006086](https://doi.org/10.1029/2018GB006086).
- Schulthess, Thomas C. et al. (Jan. 2019). “Reflecting on the Goal and Baseline for Exascale Computing: A Roadmap Based on Weather and Climate Simulations”. In: *Computing in Science & Engineering* 21.1, pp. 30–41. DOI: [10.1109/MCSE.2018.2888788](https://doi.org/10.1109/MCSE.2018.2888788).
- Schurer, Andrew P., Simon F. B. Tett, and Gabriele C. Hegerl (Feb. 2014). “Small influence of solar variability on climate over the past millennium”. In: *Nature Geoscience* 7.2, pp. 104–108. DOI: [10.1038/ngeo2040](https://doi.org/10.1038/ngeo2040).
- Shalf, John, Sudip Dosanjh, and John Morrison (2011). “Exascale Computing Technology Challenges”. In: *High Performance Computing for Computational Science – VECPAR 2010*. Ed. by José M. Laginha M. Palma et al. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 1–25. DOI: [10.1007/978-3-642-19328-6_1](https://doi.org/10.1007/978-3-642-19328-6_1).

- Shi, Yong et al. (Feb. 1, 2000). “Three-dimensional Nonhydrostatic Numerical Simulation for the PBL of Open-pit Mine”. In: *Boundary-Layer Meteorology* 94.2, pp. 197–224. DOI: [10.1023/A:1002404106737](https://doi.org/10.1023/A:1002404106737).
- Skamarock, William C. et al. (Sept. 1, 2012). “A Multiscale Nonhydrostatic Atmospheric Model Using Centroidal Voronoi Tessellations and C-Grid Staggering”. In: *Monthly Weather Review* 140.9, pp. 3090–3105. DOI: [10.1175/MWR-D-11-00215.1](https://doi.org/10.1175/MWR-D-11-00215.1).
- Sommer, Philipp S. (Aug. 22, 2017). “The psyplot interactive visualization framework”. In: *Journal of Open Source Software* 2.16, p. 363. DOI: [10.21105/joss.00363](https://doi.org/10.21105/joss.00363).
- Staniforth, Andrew and John Thuburn (Jan. 1, 2012). “Horizontal grids for global weather and climate prediction models: a review”. In: *Quarterly Journal of the Royal Meteorological Society* 138.662, pp. 1–26. DOI: [10.1002/qj.958](https://doi.org/10.1002/qj.958).
- Stohl, A. et al. (Sept. 21, 2005). “Technical note: The Lagrangian particle dispersion model FLEXPART version 6.2”. In: *Atmospheric Chemistry and Physics* 5.9, pp. 2461–2474. DOI: [10.5194/acp-5-2461-2005](https://doi.org/10.5194/acp-5-2461-2005).
- Stohl, A. et al. (Mar. 3, 2009). “An analytical inversion method for determining regional and global emissions of greenhouse gases: Sensitivity studies and application to halocarbons”. In: *Atmospheric Chemistry and Physics* 9.5, pp. 1597–1620. DOI: [10.5194/acp-9-1597-2009](https://doi.org/10.5194/acp-9-1597-2009).
- Szopa, Sophie et al. (May 1, 2013). “Aerosol and ozone changes as forcing for climate evolution between 1850 and 2100”. In: *Climate Dynamics* 40.9, pp. 2223–2250. DOI: [10.1007/s00382-012-1408-y](https://doi.org/10.1007/s00382-012-1408-y).
- Thanwerdas, Joël (Dec. 15, 2021). “Estimation des émissions et puits de méthane par inversion atmosphérique multi-contraintes du transport et de la chimie atmosphérique à l’aide d’un ensemble d’observations”. These de doctorat. université Paris-Saclay.
- Thompson, R. L. et al. (Feb. 17, 2014). “Nitrous oxide emissions 1999 to 2009 from a global atmospheric inversion”. In: *Atmospheric Chemistry and Physics* 14.4, pp. 1801–1817. DOI: [10.5194/acp-14-1801-2014](https://doi.org/10.5194/acp-14-1801-2014).
- Thuburn, J., C. J. Cotter, and T. Dubos (May 20, 2014). “A mimetic, semi-implicit, forward-in-time, finite volume shallow water model: comparison of hexagonal-icosahedral and cubed-sphere grids”. In: *Geoscientific Model Development* 7.3, pp. 909–929. DOI: [10.5194/gmd-7-909-2014](https://doi.org/10.5194/gmd-7-909-2014).
- Thuburn, John (Sept. 1, 1997). “A PV-Based Shallow-Water Model on a Hexagonal-Icosahedral Grid”. In: *Monthly Weather Review* 125.9, pp. 2328–2347. DOI: [10.1175/1520-0493\(1997\)125<2328:APBSWM>2.0.CO;2](https://doi.org/10.1175/1520-0493(1997)125<2328:APBSWM>2.0.CO;2).
- Tilmes, Simone et al. (Apr. 8, 2022). “Stratospheric ozone response to sulfate aerosol and solar dimming climate interventions based on the G6 Geoengineering Model Intercomparison Project (GeoMIP) simulations”. In: *Atmospheric Chemistry and Physics* 22.7, pp. 4557–4579. DOI: [10.5194/acp-22-4557-2022](https://doi.org/10.5194/acp-22-4557-2022).
- Tolk, L. F. et al. (Nov. 17, 2008). “Modelling representation errors of atmospheric CO₂ mixing ratios at a regional scale”. In: *Atmospheric Chemistry and Physics* 8.22, pp. 6587–6596. DOI: [10.5194/acp-8-6587-2008](https://doi.org/10.5194/acp-8-6587-2008).
- Tort, Marine and Thomas Dubos (July 1, 2014). “Usual Approximations to the Equations of Atmospheric Motion: A Variational Perspective”. In: *Journal of the Atmospheric Sciences* 71.7, pp. 2452–2466. DOI: [10.1175/JAS-D-13-0339.1](https://doi.org/10.1175/JAS-D-13-0339.1).
- Ullrich, Paul A. et al. (Dec. 6, 2017). “DCMIP2016: a review of non-hydrostatic dynamical core design and intercomparison of participating models”. In: *Geoscientific Model Development* 10.12, pp. 4477–4509. DOI: [10.5194/gmd-10-4477-2017](https://doi.org/10.5194/gmd-10-4477-2017).
- Van Der Laan-Luijkx, Ingrid T. et al. (July 18, 2017). “The CarbonTracker Data Assimilation Shell (CTDAS) v1.0: implementation and global carbon balance 2001–2015”. In: *Geoscientific Model Development* 10.7, pp. 2785–2800. DOI: [10.5194/gmd-10-2785-2017](https://doi.org/10.5194/gmd-10-2785-2017).

- Van Leer, Bram (Mar. 1, 1977). “Towards the ultimate conservative difference scheme. IV. A new approach to numerical convection”. In: *Journal of Computational Physics* 23.3, pp. 276–299. DOI: [10.1016/0021-9991\(77\)90095-X](https://doi.org/10.1016/0021-9991(77)90095-X).
- Walko, Robert L. and Roni Avissar (Nov. 1, 2008). “The Ocean–Land–Atmosphere Model (OLAM). Part II: Formulation and Tests of the Nonhydrostatic Dynamic Core”. In: *Monthly Weather Review* 136.11, pp. 4045–4062. DOI: [10.1175/2008MWR2523.1](https://doi.org/10.1175/2008MWR2523.1).
- Warneke, T. et al. (Aug. 5, 2005). “Shipborne solar absorption measurements of CO₂, CH₄, N₂O and CO and comparison with SCIAMACHY WFM-DOAS retrievals”. In: *Atmospheric Chemistry and Physics* 5.8, pp. 2029–2034. DOI: [10.5194/acp-5-2029-2005](https://doi.org/10.5194/acp-5-2029-2005).
- Wild, Martin et al. (June 2013). “The global energy balance from a surface perspective”. In: *Climate Dynamics* 40.11, pp. 3107–3134. DOI: [10.1007/s00382-012-1569-8](https://doi.org/10.1007/s00382-012-1569-8).
- Williamson, David (July 1, 2007). “The Evolution of Dynamical Cores for Global Atmospheric Models”. In: *J. Royal Met. Soc. Japan* 85B, pp. 241–269. DOI: [10.2151/jmsj.85B.241](https://doi.org/10.2151/jmsj.85B.241).
- Wu, Lin et al. (June 24, 2016). “What would dense atmospheric observation networks bring to the quantification of city CO₂ emissions?” In: *Atmospheric Chemistry and Physics* 16.12, pp. 7743–7771. DOI: [10.5194/acp-16-7743-2016](https://doi.org/10.5194/acp-16-7743-2016).
- Yamada, T. (Jan. 1983). “Simulations of Nocturnal Drainage Flows by a q^2 l Turbulence Closure Model”. In: *Journal of the Atmospheric Sciences* 40.1, pp. 91–106. DOI: [10.1175/1520-0469\(1983\)040<0091:SONDFB>2.0.CO;2](https://doi.org/10.1175/1520-0469(1983)040<0091:SONDFB>2.0.CO;2).
- Yashiro, Hisashi et al. (2016). “Performance Analysis and Optimization of Nonhydrostatic ICosahedral Atmospheric Model (NICAM) on the K Computer and TSUBAME2.5”. In: *Proceedings of the Platform for Advanced Scientific Computing Conference. PASC '16*. New York, NY, USA: Association for Computing Machinery, pp. 1–8. DOI: [10.1145/2929908.2929911](https://doi.org/10.1145/2929908.2929911).
- Yepes-Arbós, Xavier et al. (Jan. 18, 2022). “Evaluation and optimisation of the I/O scalability for the next generation of Earth system models: IFS CY43R3 and XIOS 2.0 integration as a case study”. In: *Geoscientific Model Development* 15.2, pp. 379–394. DOI: [10.5194/gmd-15-379-2022](https://doi.org/10.5194/gmd-15-379-2022).
- Zängl, Günther et al. (2015). “The ICON (ICOsahedral Non-hydrostatic) modelling framework of DWD and MPI-M: Description of the non-hydrostatic dynamical core”. In: *Quarterly Journal of the Royal Meteorological Society* 141.687, pp. 563–579. DOI: [10.1002/qj.2378](https://doi.org/10.1002/qj.2378).
- Zheng, Bo et al. (July 21, 2020). “Observing carbon dioxide emissions over China’s cities and industrial areas with the Orbiting Carbon Observatory-2”. In: *Atmospheric Chemistry and Physics* 20.14, pp. 8501–8510. DOI: [10.5194/acp-20-8501-2020](https://doi.org/10.5194/acp-20-8501-2020).

Appendices

Acronyms and common notations

Table A.1: Acronyms and abbreviations used in this report

Acronym	Meaning
ACOS	Atmospheric Carbon Observations from Space
AI	Artificial Intelligence
ASYPD	Actual Simulated Year Per Day
BIM	Budget Imbalance
CAMS	Copernicus Atmosphere Monitoring Service
CCRT	Centre de calcul recherche et technologie
CFL	Courant-Friedrichs-Lewy
CHSY	Core Hours per Simulated Year
CIF	Community Inversion Framework
CO2M	Carbon Dioxide Monitoring
COCCON	Collaborative Carbon Column Observing Network
CPU	Central Processing Unit
CMEMS	Copernicus Marine Environment Monitoring Service
CMIP	Coupled Model Intercomparison Project
CONGRAD	Conjugate Gradient Algorithm
CTE	Carbon-Tracker Europe
DCMIP	Dynamical Core Model Intercomparison Project
DGVM	Dynamic Global Vegetation Model
ECMWF	European Centre for Medium-Range Weather Forecasts
ERA5	ECMWF Re-Analysis 5

FLEXPART	FLEXible PARTicle dispersion model
GCM	General Circulation Model
GCP	Global Carbon Project
GEMS	Global Earth-system Monitoring using Satellite and in-situ data
GFED	Global Fire Emission Database
GMD	Geoscientific Model Development Journal
GOSAT	Greenhouse Gases Observing Satellite
GOSAT-GW	Global Observing SATellite for Greenhouse gases and Water cycle
GPU	Graphics Processing Unit
GridFED	Gridded Fossil Emissions Dataset
HR	High resolution
ICO	Icosahedral
IFS	Integrated Forecasting System
INCA	Interaction between Chemistry and Aerosol
IPCC	Intergovernmental Panel on Climate Change
IPSL	Institut Pierre-Simon Laplace
LAM	Limited Area Model
LMD	Laboratoire de Météorologie Dynamique
LMDZ	Laboratoire de Météorologie Dynamique Zoom
LSCE	Laboratory for Sciences of Climate and Environment
LR	Low resolution
MIP	Model Intercomparison Project
ML	Machine Learning
MPI	Message Passing Interface
MVS	Monitoring and Verification Support
NASA	National Aeronautics and Space Administration
NEMO	Nucleus for European Modelling of the Ocean
NetCDF	Network Common Data Form
NICAM	Nonhydrostatic ICosahedral Atmospheric Model
NILU	Norwegian Institute for Air Research
NOAA	National Ocean and Atmosphere Administration
NOAA/GML	National Oceanic and Atmospheric Administration Global Monitoring Laboratory
NPtP	Normalized peak-to-peak amplitude
NSD	Normalized standard deviation

ObsPack	Observation Package
OCO-2	Orbiting Carbon Observatory 2
OCO-3	Orbiting Carbon Observatory 3
OpenMP	Open Multiprocessing
ORCHIDEE	Organising Carbon and Hydrology In Dynamic Ecosystems
pdf	Probability density function
REG	Regular
RRTM	Rapid Radiative Transfer Model
SATINV	Inverse modeling for atmospheric and satellite measurements
SYPD	Simulated Year Per Day
TCCON	Total Carbon Column Observing Network
TGCC	Very Large Computing Center
UGRID	Unstructured Grid
UNFCCC	United Nations Framework Convention on Climate Change
WMO	World Meteorological Organizatio
XCH ₄	Total dry air column of CH ₄
XCO ₂	Total dry air column of CO ₂
XIOS	XML Input/Output Server
<hr/>	
a.s.l.	above sea level
e.g.	<i>exempli gratia</i> = for example
et al.	<i>et alii</i> = and contributors
<hr/>	

Table A.2: Chemical & mathematical symbols and physical units

Symbols	Meaning
C	Carbon
CH ₄	Methane
CO	Carbon monoxide
CO ₂	Carbon dioxide
H ₂	Molecular hydrogen
H ₂ O	Water vapour
N ₂ O	Dinitrogen monoxide
O ₂	Molecular oxygen

nm	nanometre 1 nm = 10^{-9} m
km	kilometre 1 km = 10^3 m
Pg	petagram 1 Pg = 10^{15} g
t	ton 1 t = 10^6 g
Gt	gigaton 1 Gt = 10^{15} g
ppm	parts per million 1 ppm = 10^{-6}
$\mathbb{E}[\cdot]$	Expected value
$p(\mathbf{x} \mathbf{y})$	probability density of the variable \mathbf{x} knowing \mathbf{y}
\mathcal{H}	Observation operator
\mathbf{H}	Jacobian matrix of the observation operator
\mathbf{x}	State vector
\mathbf{x}^b	State vector (prior)
\mathbf{x}^t	State vector (truth)
\mathbf{x}^a	State vector (posterior)
ϵ^b	Prior error
\mathbf{B}	Prior error covariance matrix
\mathbf{A}	Posterior error covariance matrix
\mathbf{y}^o	Observation vector
ϵ^o	Observation error
\mathbf{R}	Observation error covariance matrix
$J(\mathbf{x})$	Cost function

List of communications

Scientific and technical challenges of increasing horizontal resolution in atmospheric CO₂ inversion systems

2022, *EGU General Assembly 2022, Vienna (Austria)* (on-site presentation)

Zoé Lloret¹, Frédéric Chevallier¹, Anne Cozic¹.

The gradual densification of CO₂ observation networks and CO₂ observation systems around the Earth, particularly from space, has increased the observational information available for data assimilation and atmospheric inverse modeling to all spatial scales. In particular, it makes it possible to infer surface fluxes of CO₂ over increasingly small regions. This densification must be accompanied by a corresponding increase in the horizontal resolution of the transport models in which the observations are assimilated or which are inverted. In the latter application, the timescales involved extend over weeks, months or even years, and controlling computational speed despite increasing resolution is particularly critical. This challenge can be met by adapting transport models to new high-performance computing architectures and their new paradigms (multicore processors or accelerators based on graphics processing units). It deeply affects the structure of the codes, in particular the geometry of their mesh and the management of their inputs-outputs. In this study, we redesign the offline transport model of the Laboratoire de Météorologie Dynamique (LMDz) Global Atmospheric General Circulation Model used in the Copernicus Atmosphere Monitoring Service inversion system (<https://atmosphere.copernicus.eu/>) in order to test such solutions. First, we use a new dynamic core associated with an icosahedral-hexagonal spherical mesh, called DYNAMICO. DYNAMICO has a much better scalability than the current Cartesian grid of LMDz, while being efficiently vectorizable. Second, we use the parallel and asynchronous input-output management system called XIOS. XIOS helps damp performance losses associated with disk reads and writes. The technical performances of the new version will be presented in the case of a regular mesh of 16,000 hexagons on the sphere, equivalent to a global resolution of about 180 km, and with 79 vertical layers, by comparison to the regular Cartesian grid. The scientific assessment is based on a large set of CO₂ observations from the ground, from airplanes and from surface remote sensing reference sites. Particular attention is paid to the skill at high latitudes where the new grid avoids the singularity of the previous version at the pole, but at the cost of a coarser resolution.

¹Laboratoire des Sciences du Climat et de l'Environnement, CEA-CNRS-UVSQ, Gif-sur-Yvette, France

Exploiting satellite data with Dispersion, a global atmospheric inversion model at increased resolution

2023, IWGGMS-19 Workshop, Paris (France) (on-site presentation)

Zoé Lloret¹, Frédéric Chevallier¹, Anne Cozic¹, Sakina Takache¹.

Satellite data offer a high resolution view of greenhouse gas (GHG) column-average concentrations in the atmosphere, down to a few square kilometers now. Incorporating them into global-scale atmospheric inversion systems and effectively utilizing their high resolution information without unreasonably extending the computing time of atmospheric inversions, remains a central challenge, for instance for the atmospheric inversions of the Copernicus Atmosphere Monitoring Service Using the global model called Dispersion, which is based on the LMDz General Circulation Model (GCM) from the Laboratoire de Météorologie Dynamique our work addresses this challenge using two different approaches. The first one consists in replacing the latitude-longitude regular grid by a hexagonal mesh based on the Dynamico advection model which was developed for the Earth System Model of Institut Pierre-Simon-Laplace. For the same resolution at the Equator, this grid reduces the number of global cells by about one fourth by decreasing the effective resolution at high latitudes due to the absence of polar singularities. Using the full LMDz GCM to simulate CO₂ transport and compared to independent atmospheric measurements, we find a very similar performance for the new dynamical core compared to the previous one for a similar resolution at the equator. However, the new dynamical core involves a profound change in the parallelization of the code, which has slowed down its insertion in Dispersion. The second approach is an increase in resolution of the regular latitude-longitude grid after porting the Dispersion code on Graphics Processing Units (GPUs) and optimizing the volume of input data files. It is technically lighter than the first approach and allowed us to successively increase the model resolution from 3.75° longitude by 1.80° latitude and 39 vertical layers, to 2.5°x1.5°x79 layers, and then to 1.4°x0.7°x79 layers. We will synthesize the inversion results for the assimilation of NASA's OCO-2 retrievals at those resolutions, in terms of surface fluxes and fit to atmospheric observations. We will finally discuss the added value of the grid refinement.

¹Laboratoire des Sciences du Climat et de l'Environnement, CEA-CNRS-UVSQ, Gif-sur-Yvette, France

Refining the Global Picture: the Impact of Increased Resolution on CO₂ Atmospheric Inversions using OCO-2 XCO₂ retrievals

2024, *TransCom-2024 Meeting, Boulder (USA)* (remote presentation)

Zoé Lloret¹, Frédéric Chevallier¹, Anne Cozic¹.

The threat posed by the increasing concentration of carbon dioxide (CO₂) in the atmosphere motivates a detailed and precise estimation of CO₂ emissions and absorptions over the globe. This study refines the spatial resolution of the CAMS/LSCE inversion system, achieving a global resolution of 0.7° latitude and 1.4° longitude, or three times as many grid boxes as the current operational setup. In a two-year inversion assimilating the midday clear-sky retrievals of the column-average dry-air mole fraction of carbon dioxide (XCO₂) from NASA's second Orbiting Carbon Observatory (OCO-2), the elevated resolution demonstrates an improvement in the representation of atmospheric CO₂, particularly at the synoptic time scale, as validated against independent surface measurements. Vertical profiles of the CO₂ concentration differ slightly above 22 km between resolutions compared to AirCore profiles, and highlight differences in the vertical distribution of CO₂ between resolutions. However, this disparity is not evident for XCO₂, as evaluated against independent reference ground-based observations. Global and regional estimates of natural fluxes for 2015-2016 are similar between the two resolutions, but with North America exhibiting a higher natural sink at high-resolution for 2016. Overall, both inversions seem to yield reasonable estimates of global and regional natural carbon fluxes. The increase in calculation time is less than the increase in the number of operations and in the volume of input data, revealing greater efficiency of the code executed on a Graphics Processing Unit. This allows us to make this higher resolution the new standard for the CAMS/LSCE system.

¹Laboratoire des Sciences du Climat et de l'Environnement, CEA-CNRS-UVSQ, Gif-sur-Yvette, France

