



HAL
open science

**The role of greenhouse gases in past climatic variations :
an approach based on accurate chronologies of deep
polar ice cores**

Jai Chowdhry Beeman

► **To cite this version:**

Jai Chowdhry Beeman. The role of greenhouse gases in past climatic variations : an approach based on accurate chronologies of deep polar ice cores. Climatology. Université Grenoble Alpes, 2019. English. NNT : 2019GREAU023 . tel-02450783

HAL Id: tel-02450783

<https://theses.hal.science/tel-02450783>

Submitted on 23 Jan 2020

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

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

THÈSE

Pour obtenir le grade de

DOCTEUR DE LA COMMUNAUTE UNIVERSITE GRENOBLE ALPES

Spécialité : Sciences de la Terre et de l'Univers et de l'Environnement

Arrêté ministériel : 25 mai 2016

Présentée par

Jai Chowdhry Beeman

Thèse dirigée par **Frédéric Parrenin**, Chargé de recherche, CNRS, et codirigée par **Amaelle Landais**, Directeur de recherche, CNRS et **Emmanuel Witrant**, Maître de conférences, Université Grenoble Alpes

préparée au sein de l'**Institut des Géosciences de l'Environnement**

dans l'**École Doctorale Terre, Univers, Environnement**

Le rôle des gaz à effet de serre dans les variations climatiques passées : une approche basée sur des chronologies précises des forages polaires profonds.

The role of greenhouse gases in past climatic variations: an approach based on accurate chronologies of deep polar ice cores

Thèse soutenue publiquement le **21 Octobre 2019**
devant le jury composé de :

Monsieur Frédéric Parrenin

Chargé de recherche, CNRS Délégation Alpes, Directeur de thèse

Madame Dorte Dahl-Jensen

Professeur, Université de Copenhague, Rapporteur

Monsieur Eric Wolff

Professeur, Université de Cambridge, Rapporteur

Monsieur Shaun Marcott

Professeur Assistant, Université du Wisconsin à Madison, Examineur

Monsieur Eric Blayo

Professeur, Université Grenoble Alpes, Président



The role of greenhouse gases in past climatic variations: an
approach based on accurate chronologies of deep polar ice
cores

Jai Chowdhry Beeman

Abstract/Résumé

Deep polar ice cores contain records of both past climate and trapped air that reflects past atmospheric compositions, notably of greenhouse gases. This record allows us to investigate the role of greenhouse gases in climate variations over eight glacial-interglacial cycles. The ice core record, like all paleoclimate records, contains uncertainties associated both with the relationships between proxies and climate variables, and with the chronologies of the records contained in the ice and trapped air bubbles. In this thesis, we develop a framework, based on Bayesian inverse modeling and the evaluation of complex probability densities, to accurately treat uncertainty in the ice core paleoclimate record. Using this framework, we develop two studies, the first about Antarctic Temperature and CO₂ during the last deglaciation, and the second developing a Bayesian synchronization method for ice cores. In the first study, we use inverse modeling to identify the probabilities of piecewise linear fits to CO₂ and a stack of Antarctic Temperature records from five ice cores, along with the individual temperature records from each core, over the last deglacial warming, known as Termination 1. Using the nodes, or change points in the piecewise linear fits accepted during the stochastic sampling of the posterior probability density, we discuss the timings of millennial-scale changes in trend in the series, and calculate the phasings between coherent changes. We find that the phasing between Antarctic Temperature and CO₂ likely varied, though the response times remain within a range of 500 years from synchrony, both between events during the deglaciation and across the individual ice core records. This result indicates both regional-scale complexity and modulations or variations in the mechanisms linking Antarctic temperature and CO₂ across the deglaciation. In the second study, we develop a Bayesian method to synchronize ice cores using corresponding time series in the IceChrono inverse chronological model. Tests show that this method is able to accurately synchronize CH₄ series, and is capable of including external chronological observations and prior information about the glaciological characteristics at the coring site. The method is continuous and objective, bringing a new degree of accuracy and precision to the use of synchronization in ice core chronologies.

Les forages polaires profonds contiennent des enregistrements des conditions climatiques du passé et de l'air piégé qui témoignent des compositions atmosphériques du passé, notamment des gaz à effet de serre. Cette archive nous permet de décrypter le rôle des gaz à effet de serre dans les variations climatiques pendant huit cycles glaciaire-interglaciaires, soit l'équivalent de plus de 800 000 ans. Les carottes de glace, comme toute archive paléoclimatique, sont caractérisées par des incertitudes liées aux processus qui traduisent les variables climatiques en proxy, ainsi que par des incertitudes dues aux chronologies de la glace et des bulles d'air piégées. Nous développons un cadre méthodologique, basé sur la modélisation inverse dite Bayésienne et l'évaluation de fonctions complexes de densité de probabilité, pour traiter les incertitudes liées aux enregistrements paléoclimatiques des carottes de glace de manière précise. Nous proposons deux études dans ce cadre. Pour la première étude, nous identifions les probabilités de localisation des points de changement de pente de l'enregistrement du CO_2 dans la carotte de WAIS Divide et d'un stack d'enregistrements de paléotempérature à partir de cinq carottes Antarctiques avec des fonctions linéaires par morceaux. Nous identifions aussi les probabilités pour chaque enregistrement individuel de température. Cela nous permet d'examiner les changements de pente à l'échelle millénaire dans chacune des séries, et de calculer les déphasages entre les changements cohérents. Nous trouvons que le déphasage entre la température en Antarctique et le CO_2 a probablement varié (en restant inférieur, généralement, à 500 ans) lors de la déglaciation. L'âge des changements de température varie probablement entre les sites de carottage aussi. Ce résultat indique que les mécanismes qui reliaient la température en Antarctique et le CO_2 lors de la déglaciation pouvaient être différents temporellement et spatialement. Dans la deuxième étude nous développons une méthode Bayésienne pour la synchronisation des carottes de glace dans le modèle inverse chronologique IceChrono. Nos simulations indiquent que cette méthode est capable de synchroniser des séries de CH_4 avec précision, tout en prenant en compte des observations chronologiques externes et de l'information à priori sur les caractéristiques glaciologiques aux sites de forage. La méthode est continue et objective, apportant de la précision à la synchronisation des carottes de glace.

Dedication

To Geeta Chowdhry

Acknowledgements

First, I would like to thank Frédéric Parrenin, Emmanuel Witrant and Amaelle Landais, my thesis advisors. Their combined expertise made this thesis possible, and they have been insightful, motivating and always approachable throughout the three years. It has been a real pleasure to work with each of them.

I would also like to thank all of the collaborators and co-authors on the studies presented in this thesis, in particular Léa Gest, who developed and co-led the study on Antarctic Temperature and CO₂; Jinhwa Shin, who invited me to work with her on a study on CO₂ during the MIS 6 period and who kindly proofread several of the chapters of my thesis and shared near-fatally spicy food to help me during moments of stress; and Dominique Raynaud, the godfather of the CO₂ study and a limitless source of research inspiration. Mai Winstrup provided invaluable advice as a member of my thesis committee. Thanks as well to Delphine Six, Christian Vincent, Olivier Laarman et al. for the many opportunities to help out with glacier fieldwork in the Alps.

I am grateful to a very long list of family and friends, my father Mark Beeman for his optimism and constant encouragement, Priya, Tim, Sara and Isha for providing a second home not too far from Grenoble, and Julien, Laura and Gabi for making a home in Revel that I hope to delay leaving for as long as possible. I thank many people for the adventures spent together: Romain, in particular for a bike trip through Switzerland that restored my sanity at the end of the third year, Julien, Maria, Gauthier, Jonathan, Claudio, Jean Carlos, Astrid, Diego, Hans, Sophie, Anil, Nate, Amber, Jordi, Cinthya and Nico; Pranav; Sarah, Peter and Cynthia; Sara, Ron and Teo, Sol, Brooke, Paige, Debanjan ... and this list is very far from being complete.

Finally, a huge thanks to family in India and the US for always providing solid support and encouragement.

Contents

Foreword	6
1 The state of the art in ice core chronology	9
2 Paleoclimate chronologies and Bayesian Methods	52
3 The Last Deglaciation : Antarctic Temperature and CO ₂	80
4 Methane Synchronization	95
5 Conclusions and Perspectives	124
A Supplementary Material	134
B CO ₂ during MIS 6	156

Foreword

Rapid, anthropogenic climate change is set to mark the coming century, with disproportionate impacts on developing societies and vulnerable populations. Observations and modeling have rapidly advanced our understanding of the climate system, but as we prepare to combat and adapt to climate change, many important scientific questions remain to be assessed. One of these: how do the mechanisms of the earth's climate operate on the long time scales of thousands to hundreds of thousands of years? In the coming centuries and millennia, what might we expect?

Records of paleoclimate allow us to begin to investigate these questions, which are difficult to answer using contemporary measurements and modeling studies. Unique among these archives, ice cores record past atmospheric composition and proxies of past climate, allowing us to investigate the greenhouse gas-climate relationship that is key to current climate change on much longer timescales.

Paleoclimate records are, of course, imperfect observations. Measurements of chemical species are only accurate within a range of uncertainty. The transfer functions relating proxies to climate variables are by nature approximate. Finally, the age-depth relationships that we use to date paleoclimate records are based on uncertain and limited observations. Thus, paleoclimate records contain two dimensions of uncertainty – age uncertainty, and uncertainty with respect to the climate variables themselves.

This thesis focuses on accurately characterizing these two sources of uncertainty. In order to do so, we apply Bayesian inversion – the use of observations and their uncertainties to determine the range of uncertainty within which we can understand unknown quantities. Bayesian inversion is a fitting framework for paleoclimatology – the science of determining the unknown climate of the past from the observable properties of the earth's layered surface in the present.

The first chapter of this thesis introduces ice core chronology. In this chapter, we discuss each of the sources of chronological information available in ice cores, and their respective limitations and uncertainties.

The second chapter focuses on Bayesian inversion. This chapter begins with an introduction to Bayesian statistics. We then introduce the IceChrono model, which develops ice core chronologies in a Bayesian framework. Finally, we introduce a few stochastic solution methods, adapted to solving complex Bayesian inverse problems, that we use in the later chapters of this thesis. This chapter is more mathematical than the first, as it introduces in considerable detail the methods used in the two next chapters, which are articles describing original research.

In the third chapter, an article published in *Climate of the Past* and co-led with Léa Gest, we apply an inverse framework to determine the timings of millennial-scale changes in CO₂ and Antarctic (regional and continental) temperature over the course of the last deglaciation, from roughly 20,000 to 10,000 years ago. This problem is perhaps less simple than it seems – we need to model the structure of the changes, in this case, using piecewise linear change; and then determine the probabilities of when they occur, using an inverse framework. All the while, we take into account measurement and chronological uncertainty with respect to the time series. We arrive at a probabilistic, accurate assessment of the timings of changes and phase lags between the two series. This estimate gives us a more accurate observation of how these two key variables of the carbon-climate cycle operated on the millennial timescales of the glacial termination.

In the fourth chapter, an article in preparation for submission to *Geoscientific Model Development*, we develop a method to automate the synchronization of paleoclimate records. Many time series measured on ice cores are known to be similar, and matching them is a powerful tool for building detailed chronologies. This is usually done by eye, meaning that the match is discrete, and uncertainties are difficult to estimate. Our method is continuous and Bayesian, making the estimation of uncertainties inherent to the calculation.

The similarities in the last two studies are numerous. In both studies, we use residuals between a model and observations to define a cost function – though in the first study, the model is much simpler. In both studies, this cost function is used to develop a posterior probability density function, in the Bayesian sense. We do not assume, in either study, that the posterior will take a particular shape, but rather apply stochastic methods that allow us to characterize complex, multimodal probability density functions.

In the final chapter, we bring the two studies together in the context of this Bayesian framework. We outline the similarities and differences between the method design for the two studies and the limitations of the methodologies, before proposing some concluding

remarks about the broader significance of the two works.

We begin, as such, to propose an approach for tackling complex state-of-knowledge problems in paleoclimate. This approach will surely see further development, as more powerful stochastic sampling algorithms allow us to better characterize probability density functions, as computational resources increase, and, most importantly, as more paleoclimate archives are collected in the field and measured. We hope that further scientific investigation will continue to improve upon the limitations detailed in our studies, and that these developments are only a beginning.

Chapter 1

The state of the art in ice core chronology

The state of the art in ice core chronology*

Jai Chowdhry Beeman Frédéric Parrenin Volodya Lipenkov
Dominique Raynaud Amaelle Landais Emmanuel Witrant

Introduction

Deep polar ice cores contain a continuous record of atmospheric composition and proxy representations of Earth’s climate currently reaching 800,000 years into the past. To access this valuable archive of the climate system, we must precisely date the air and ice contained in the ice cores. The methods developed in order to do so integrate the disciplines of chemistry, physics, mathematics and geology to extract as much information as possible from chronological clues in the air and ice.

An ice core begins as snow on the surface of an ice sheet, ice cap, or glacier. As more snow accumulates, the underlying layers begin to densify into firn. When these layers become sufficiently dense and firn becomes ice, the air that circulates through the overlying layers is locked in to air bubbles, which are integrated into the ice matrix as clathrates deeper in the core. As a result, the air at a given depth in a core is younger than the surrounding ice. The ice layers themselves become thinner as we move deeper in the core, as a function of ice flow conditions at the site where the core is drilled. Together, these processes—accumulation, snow densification, air lock-in and ice flow—are referred to as sedimentation. Models of the sedimentation process are an important component in dating ice cores.

The snow accumulated each year on a glacial surface is referred to as an annual layer. Where snow accumulation is high enough so that annual layers are not broken down by wind, they can be identified and counted using seasonal cycles recorded in dust, water isotopes, soluble compounds, and visual stratigraphy. Layer counting is the most consequential method for the dating of Greenland ice cores, where annual layers tend to be

*A version of this chapter was written for a monograph in preparation by Volodya Lipenkov, and will be translated into Russian. The chapter is inspired by Parrenin [2014], which provides a comprehensive introduction to ice core dating in French.

very well preserved, and has also been applied to cores close to the Antarctic coast, where accumulation is stronger than in the Antarctic interior.

In low-accumulation ice cores, which record longer time periods, annual layers become very thin or even unrecognizable at depth. As such, other methods have been developed to date ice cores on these longer time scales. The isotopic ratio between ^{18}O and ^{16}O in air; the fractionation ratio between O_2 and N_2 , and Total Air Content (TAC) have been tuned to the periodicities of Earth's orbit to produce age models. Since Earth's orbital parameters can be calculated well into the distant past, this method can be used on the full depth of even the deepest ice cores; however, the physical relationships between insolation and its proxies are not perfectly known, resulting in uncertainty on the scale of a few thousand years.

Significant events in Earth's history often leave indicators or markers in ice core layers, which we also refer to as dated horizons. These are another important tool in dating ice cores. Volcanic events can be identified and linked to known eruptions by their chemical and dielectric properties. Rapid variations in temperature during the last glaciation, known as Dansgaard-Oeschger events, are well-dated by annual layer counting in Greenland ice cores and by Uranium-Thorium dating in speleothems, and can be identified and linked to these archives in Antarctic cores. Concentrations of Beryllium-10, which is influenced by the Earth's magnetic field and by solar wind, can be measured in ice cores, allowing us to identify links between cores and with other dated archives. This is particularly useful during major reversals in the Earth's magnetic field, which can be dated radiometrically in other paleoclimate archives. Carbon-14 spectrometry can be applied to air trapped in Antarctic ice cores, but consumes a large amount of CO_2 , limiting its application to cases where large amounts of ice are available (see Petrenko et al. [2009, 2016] for example).

Finally, ice cores can be linked to each other and to other paleoclimate archives in both the ice and gas phases. Methane concentrations, which are assumed to be globally consistent, are commonly used to synchronize ice cores in the air phase, and have even been used to synchronize ice cores to speleothem isotopic records; the ratio of $\delta^{18}\text{O}$ of atmospheric O_2 can be used to synchronize ice cores in the air phase as well. Volcanic events, even when not dated absolutely, provide a powerful tool for relative dating between ice cores in the ice phase, and can be used to link ice cores to speleothems as well.

The methods used to date ice cores each contain sources of uncertainty. Annual layer counting, for example, is very precise in terms of resolution, but counting errors tend to

accumulate as we move away from a layer whose age is known, and annual layers become unresolvable in older and lower-accumulation ice cores. Tuning to orbital parameters is accurate on very long time scales, but significant uncertainty is introduced since the physical relationships between proxies of orbital parameters and the parameters themselves are not well known. Models of the sedimentation process, on the other hand, are based on well-known principles of flow, but require knowledge of external parameters like accumulation and basal melt rate. The identification of volcanic events and changes in Earth's magnetic field and solar activity allow us to date individual layers, but the ages of the ice and gas in between these layers are left unknown.

We can mitigate some of this uncertainty by combining multiple sources of chronological information. The methods developed to do so are federative: that is, they combine dating information from multiple dating methods and often multiple ice cores. They are often formulated as inverse problems, and involve estimating the probabilities of the variables that create an age model using the chronological information available about the age of the core. Any of the dating methods mentioned in this introduction can be included as a source of information in a federative dating experiment.

Each of the dating methods described in this introduction will be treated in its own section in this chapter. In the sections, we make an attempt to include well-known studies as examples to guide the reader through the method, describing how it has been used in dating ice cores, how the associated uncertainties are evaluated, and the scientific advances that the method has allowed. We leave the reader with an introduction to federative dating methods, which are arguably the state of the art in ice core dating, and with perspectives of what is to come in dating ice cores.

A note to the reader: the convention of counting time in years or kiloyears before 1950 or before present, with the present referring to 1950 AD (ka B1950, ka BP, or simply ka), is commonly used in ice core and paleoclimate studies. We will as a rule use this convention in this chapter, though alternatives, such as years before 2000 AD, are also common.

The sedimentation process

Polar ice cores record the past in two forms: ice and air. The ice begins as fresh snow, which is deposited on the surface of an ice sheet. We quantify this deposition as accumulation, or a. As the snow densifies under its own weight into firn and then ice, the air that moves through the porous snow becomes trapped in bubbles in the denser ice matrix.




Figure 1: Schematic (not to scale) of the firn column at WAIS Divide. A darker background color indicates a denser ice matrix. Pores are indicated in white. From Battle et al. [2011].

Sowers et al. [1992] identify four distinct zones that characterize the transition from snow to firn and then ice. The first of these, in the top layer of snow and firn, is the convective zone. This zone is porous enough for air to be transported via convection: that is, the movement of molecules induced by the difference in temperature between the surface and the bottom of the layer. Convection consists of two components: advection, the transfer of mass via fluid currents, and diffusion, the motion of individual molecules.

Firn becomes less porous in the next layer, the diffusive zone. Here, advection (and therefore convection) no longer occurs, and molecules are transported by diffusion alone. Below the diffusive zone, in the non-diffusive or lock-in zone, pores in the firn are almost closed, such that even diffusion is negligible. The depth in the firn at which this occurs is known as the lock-in depth, or LID. The lock-in depth varies in time and is site-dependent: at the site of the NEEM core in Greenland, it varies on the order of 50 meters, while at interior Antarctic sites like Vostok, this depth can be on the order of 100 meters, and the surrounding firn can be many decades to several thousand years older than the air it traps. At the close-off depth, located even deeper than the lock-in depth, all pores become definitely closed. However, little additional air is trapped in pores between the lock-in depth and the close-off depth, so we tend to use the lock-in depth to calculate the age of the air in a pore. We quantify the chronological difference between the ice and air in an ice core in two ways. Δage refers to the difference in age between ice and air at an equivalent depth. Analogously, Δdepth refers to the difference in depth between ice and air with equivalent ages. Both of these variables are related to the LID.

Finally, as ice is weighted with newly accumulated snow, it flows. The layers of ice that began as surface snow are transported deeper into the ice sheet, and begin to move horizontally along flow lines as well. Consequently, the layers representing a given period in an ice core become thinner at depth. Along the vertical line of an ice core, we can quantify the thinning of layers with a thinning function τ . If we have estimates of the accumulation, the thinning function, and the lock-in depth along the depth of an ice core, we can construct a chronology based on sedimentation modeling. When creating an ice core chronology, modeling and proxies are used to effectively gain information about these three variables.

Lock-in depth is perhaps the most difficult of the three variables to estimate, as it involves both the gas transport processes and the densification mechanisms that occur in the firn column. Models based on the principles of mechanics are used to calculate the evolution of the profile of density in the firn—a classical formulation is that of Herron

and Langway [1980]. In these models, lock-in depth is parameterized in one of two ways: the depth where the density of the firn column becomes sufficiently high; or the depth at which the ratio of closed to total porosity becomes sufficiently high [Goujon et al., 2003].

These parameterizations are based on modern-day observations of the properties of firn at lock-in depth. Firn modeling is treated in much greater depth by Salamatin et al. [2009], for example.

In addition to firn densification modeling, isotopes of Nitrogen (^{15}N and ^{14}N) and Argon (^{40}Ar and ^{36}Ar) can be used as indicators of lock-in depth. [Sowers et al., 1992, Severinghaus et al., 2003] Nitrogen and Argon concentrations in the atmosphere have not undergone significant changes on the timescales relevant to dating ice cores. However, their isotopic ratios as recorded in ice cores are not constant. These variations result from fractionation processes—processes that separate different isotopes—that occur in the firn. The processes that separate Nitrogen and Argon isotopes can indeed be related to the lock-in depth at the time of pore closure.

The first and most important of these fractionation processes, gravitational fractionation, occurs due to the mass difference between isotopes. Heavier isotopes fall to the bottom of the diffusive zone and become trapped as pores in the firn begin to close. The greater the height of the diffusive column, the more these heavier isotopes are preferred deeper in the core. We can describe the gravitational fractionation ratio, δg , as

$$\delta g = \left[e^{\frac{\Delta m g z}{RT}} - 1 \right] \cdot 1000 \quad (1)$$

where Δm is the mass difference (in atomic mass units, or g mol⁻¹) between the two isotopes, g is the gravitational acceleration (in m s⁻²), z is the height of the diffusive column (in m), T is temperature (in K), and R is the ideal gas constant (expressed in kg m² s⁻² K⁻¹ mol⁻¹). Note that gravitational fractionation is stronger for greater differences in isotopic mass: e.g. it is stronger for Argon than for Nitrogen isotopes.

Gravitational fractionation is dependent on the height of the diffusive column, rather than the height of the firn column from the lock-in depth to the surface. This means that in order to estimate the full LID using isotopic ratios, we must also make an assumption about the height of the convective zone. The observed present-day height of the convective zone is variable by site, from near-zero at EDC to more than 10m at Vostok [Landais et al., 2006]. We do not, however, know the height of the past convective zone, introducing a source of uncertainty in the LID estimation that is difficult to quantify.

The second type of fractionation process, thermal fractionation, is related to the tem-

perature gradient in the firn column. It can be written as:

$$\delta_t = \left[\left(\frac{T}{T_0} \right)^c - 1 \right] \cdot 1000 \quad (2)$$

where T and T_0 are the temperatures at the bottom of the diffusive column and at the surface, respectively, and c , the thermal diffusion coefficient, is related to temperature as well.

These relationships can be used in addition to temperature proxies, like Deuterium isotopes, to deduce the lock-in depth associated with a given air depth in the ice core. However, the lock-in depth estimated by these relationships is not always in agreement with the lock-in depth calculated by firn densification modeling, particularly in the case of low-accumulation interior Antarctica in the glacial period [Landais et al., 2006]. Several explanations have been proposed, including corrections of accumulation rates [Landais et al., 2006] and continued dispersive mixing in the lock-in zone driven by pressure variations [Buizert and Severinghaus, 2016]; a modeling correction for low-accumulation Antarctic sites, including a relationship between activation energy and temperature, and a parameterization of the influence of impurities on densification, has been proposed by Bréant et al. [2017].

In addition to gravitational and thermal fractionation, Kawamura et al. [2013] describe a kinetic fractionation mechanism in the convective zone, where fractionation was previously thought negligible. Very heavy gases are slower to diffuse than lighter gases, and are thus partially depleted lower in the firn. The authors propose krypton and xenon, which are slow enough to diffuse to be significantly impacted by kinetic fractionation, as proxy tracers of this effect.

Estimating lock-in depth is the first step to modeling the age of the air in an ice core. In order to date the air and the surrounding ice using a model of the sedimentation process, we need to additionally estimate the accumulation, which is time-dependent, and the thinning function.

Fortunately, there are a few ways to do this. Accumulation is most often initially estimated using relationships with water isotope records, like δD or $\delta^{18}\text{O}$, which are known as paleothermometers because of their strong relationships with temperature. The saturation vapor pressure, which is an exponential function of temperature, is assumed to control precipitation. However, the relationships between water isotopes and temperature are imperfect, and this contributes to the complexity of the relationships between isotopic proxies and accumulation. The uncertainty in calculating past accumulation in turn

motivates some of the methods described in later sections, in particular inverse modeling.

Estimating the thinning function is simple for ice cores which preserve annual layers, in which case it suffices to measure the thickness of a given layer and derive the thinning function at that layer using an accumulation estimate. Since the ice is already dated by the annual layer count, the only methodological uncertainties associated with sedimentation modeling are in the estimate of the air age. Many Greenland ice cores, and some Antarctic cores, are at least partially dated by annual layer counting, as we will see in the next section.

Where annual layers are not retained in the ice, we can model the flow of ice at the location of the core using dynamic ice flow models. These models vary in complexity, and often the characteristics of the site at which an ice core is drilled allow us to use a simpler model. At the site of a geographic dome, such as Dome C or Dome Fuji, we can simplify the calculation of horizontal ice flux. As such, one-dimensional, vertical flow models with approximations of flux and basal conditions have been applied (see Parrenin et al. [2007b] for example).

Cores situated downslope of geographic domes require more detailed modeling of horizontal ice flow. In the case of the Vostok ice core, a 2.5 dimensional ice flow model, which supposes that ice flows in the direction of the greatest slope, has been applied [Parrenin et al., 2004]. 3-dimensional models, while costly in terms of computation, can be used to tune or verify simplified models. [Ritz et al., 2001], for example, simulated the Antarctic ice sheet using a 3-dimensional model over the last 420,000 years, and the results of this experiment were used to constrain the Vostok simulation performed by [Parrenin et al., 2004]. Note that ice flow modeling is described in much greater detail in chapters 4 and 5.

Once the accumulation α , the lock-in depth LID , and the thinning function τ have been constructed, we can make age models for the air and ice in a core. At depth z , the age of the ice X can be written:

$$X(z) = \int_0^z \frac{D(z')}{\alpha(z')\tau(z')} dz' \quad (3)$$

Then, the age of the air contained at a given depth, ψ , is given by:

$$\psi(z) = X(z - \Delta d(z)) \quad (4)$$

where Δd is Δ depth (the difference in depth between synchronous air and ice layers). Finally, we can calculate Δd from lock-in depth LID by integrating the following

relationship:

$$\int_{z-\Delta D(z)}^z \frac{D(z')}{\tau(z')} dz' = \int_0^{LID(z)D^{firn}(z)} \frac{1}{\tau(z^{ie})} dz^{ie} \quad (5)$$

where D^{firn} , which we assume we know from densification modeling, is the average relative density of firn when the air at depth z was at the lock-in depth. We call z^{ie} the ice equivalent depth, defined as:

$$z^{ie} = \int_0^z D(z') dz' \quad (6)$$

These four equations allow us to construct a chronology from the information we have about the sedimentation processes in an ice core. Additionally, they form the forward model for many inverse simulations: that is, they allow us to deduce the most probable scenarios for the three input variables α , LID , and τ when we have some additional information about the age of the ice or air. In the following sections, we explore the methods that provide us with this additional chronological information. We return to inverse dating methods, which allow us to combine this information, in the final section.

It should be noted that further improvements can certainly be made to models of the sedimentation process. Glacier flow modeling is a geoscientific discipline in its own right. The relationships between accumulation and isotopic ratios should be further refined. Finally, continued investigation of the processes of firn-layer transport will surely improve both model-based and proxy-based estimation of lock-in depth.

Annual Layer Counting

When sufficient snow falls on the surface of an ice sheet, and this snow is left relatively undisturbed over the course of a year, distinct annual layers form as the snow densifies into firn and then ice. The identification of these annual layers in the ice has allowed high accumulation ice cores in Greenland and near the Antarctic coast to be dated at very high resolution.

Layer counting is made possible by the presence of tracers in the ice that are seasonally cyclical. Early on, the presence of annual-resolution visual stratigraphy—comparable to varves in lake sediment cores or annual tree rings, was recognized in ice cores [Benson, 1962], as the largest of these layers are evident to the naked eye. The visual stratigraphy present in ice cores is related to annual cycles in the content of impurities in the ice, which determine its color. Though visual stratigraphy is a methodologically simple, powerful

dating tool, it was difficult to apply in practice until the advent of digital photography.

Seasonal cycles in isotopic composition were also used to date ice cores early on (see, for example, Hammer et al. [1978]). The water isotope ratios $\delta^{18}\text{O}$ and δD contain seasonal cycles that are, in large part, related to seasonal differences in isotopic fractionation during moisture transport. Air masses moving from the tropics toward the poles tend to be cooled more in winter than in summer, and more precipitation occurs along the transport pathway as a result. This rainout leads to a depletion in heavier isotopes and, generally, lower values of the isotopic ratios in winter than in summer [Epstein and Sharp, 1959].

Since water vapor can diffuse through the firn column, annual layer counting using water isotopes is limited. This effect is less important for ions trapped in the ice matrix, notably Na^+ , Ca^{2+} , H_2O_2 , NH_4^+ , NO_3^- and SO_4^{2-} , which can also vary in concentration on a sub-annual timescale, tracking seasonal components of the biogeochemical system. The concentrations of Na^+ , Ca^{2+} , and SO_4^{2-} , for example, are affected by the influx of sea salt aerosols to the site of an ice core. SO_4^{2-} can also contain a biogenic signal [Curran et al., 1998]. Insoluble dust concentrations in ice cores show a significant seasonal component as well, which is dominated by seasonal precipitation patterns at the land source [Hammer et al., 1978].

It is important to note that the atmospheric transport mechanisms hypothesized to produce the annual cycles observed in water isotopes and impurities, and the timing of the annual peaks and lows in these tracers, can vary from one paleoclimate period to another [Andersen et al., 2006], and indeed from one site to another as well. Thus, even when we still observe annual cycles, identifying individual seasons can prove ambiguous.

The recognition of annual layers has become considerably more accurate as new methods to measure visual stratigraphy, water isotopes, and impurities have been developed. Hammer [1980] identified Electrical Conductivity Measurement (ECM) as an integrative source of information about the ions contained in an ice core; ECM is still commonly applied as a component of annual layer counting studies. Using digital photographic processing techniques, very high-resolution line scans of ice cores can be taken on site [Nielsen, 2005] and processed to make depth series of grey-tone intensity [Winstrup et al., 2012] which show identifiable annual cycles. This method makes the simple concept of visual stratigraphic counting easier to apply.

Measurements of water isotopes and impurities in ice cores were transformed, in terms of resolution, by the advent of Continuous Flow Analysis (CFA). By continuously melting a core, measurements in the liquid phase and of many components in the gas phase can be

made at increasingly high resolution (e.g. Bigler et al. [2011]), limited by the quantities of water or air required by the experimental setup, rather than by the work hours required to cut individual ice core samples. Of particular interest to annual layer counting studies, CFA setups have been developed for both water-soluble ions [Röthlisberger et al., 2000] and water isotopes [Gkinis et al., 2011]. These very high resolution measurements make the identification of an annual cycle considerably easier.

Compared to low-accumulation Antarctic ice cores like those from Dome C, Dome Fuji and Vostok, Greenland ice cores are generally higher-accumulation records, with more easily identifiable annual layers reaching considerable depths. The Greenland Ice Core Chronology 2005 (GICC05; Svensson et al. [2008]), the most up-to-date chronology of Greenland ice cores, reaches 60ka at annual resolution using layer counting (and 123 ka using glaciological modeling), with uncertainty estimated at under 1% in the Holocene and under 5% during the last glaciation. Antarctic cores, on the other hand, cannot be left out: the West Antarctic Ice Sheet Divide Core (WAIS Divide core; WD) chronology (WD2014) is based on annual layer counting up to 31ka [Sigl et al., 2016] and significant methodological advances were produced during the dating of the WAIS Divide core as well.

Several integrative dating scenarios for multiple Greenland ice cores have been developed that take multiple sources of layer-counting information into account. The GISP2 scenario of Meese et al. [1997] was used by many paleoclimate studies; the most current chronology for Greenland ice cores is GICC05, which uses high-resolution data series from three Greenland ice cores—DYE-3, GRIP, and NorthGRIP—to perform a multi-parameter, continuous count of annual layers over the last 60 ka. The multi-core, multi-parameter approach holds two significant advantages over a single-core, single-parameter dating scenario. First, the seasonal tracers: water isotopes, soluble ions, ECM measurements, and visual stratigraphy, can complement each other. The use of multiple tracers is of particular value when the seasonal cycle in a single tracer is not evident for a given paleoclimate period. Second, the measurements made in the core with the highest resolution can be used to limit uncertainty at a given point, using stratigraphic links (volcanic events, for example) to match the other cores to the least uncertain time scale.

An example of layer counting from GICC05, within Greenland Interstadial 14 (beginning at $54\,220 \pm 1150$ yr BP, ending at $49\,280 \pm 1015$ yr BP) is shown in Figure 2. During this period, layer counting was performed using visual stratigraphy grey scale, ECM, Conductivity, and Na^+ measurements.

A few aspects of the practical implementation annual layer counting are evident in the figure that may not be obvious when the subject is treated in a purely theoretical way. Annual cycles themselves can sometimes be ambiguous: some series show peaks on a sub-annual time scale, while others seem to be missing some annual peaks and troughs. While layer counting is by far the highest-resolution method used to date ice cores, the ambiguity of annual layers introduces some uncertainty into the dating scenario. A less-apparent layer might be left uncounted, or an extra layer may be counted when in fact the ice stratum was simply abnormally thick.

These counting errors accumulate as we move further from ice strata with known ages. In the GICC05 dating scenario, to address the cumulative nature of counting error, layers thought to be present but deemed to be uncertain were counted as $\frac{1}{2} \pm \frac{1}{2}$ years (shown as dashed lines), with the uncertainty term summing between absolutely dated points.

In addition, the ambiguity in identifying annual layers can add a degree of subjectivity to the method: two scientists counting annual layers in the same core may, and likely will, identify a different amount of layers, and this difference will increase deeper in the core. The GICC05 time scale made an effort to reduce human error by combining the layer counting results of multiple authors for the final age scenario.

Several recent studies have made significant strides toward introducing automation in the counting of annual layers, and replacing subjectivity with reproducibility [Rasmussen et al., 2002, Smith et al., 2009, McGwire et al., 2011]. The study of Winstrup et al. [2012] made the largest advances in counting annual layers automatically, and this method, detailed in Chapter 2, has been used in several major ice core chronologies.

While automated methods reduce the subjectivity of layer counting, human input is still essential to verify the results. Indeed, in practical applications, some strongly ambiguous layer counting scenarios still arise. In these cases, manual counting is often the best-performing method.

Annual layer counting is highly precise in dating the duration of periods at depth, but because of the cumulative nature of errors (more layers are missed in total as the count progresses) it can be uncertain in terms of absolute dating at depth. Additional dating information, like the inclusion of dated horizons, is thus often used (as is the case in GICC05, for example) to constrain the absolute dating uncertainty when annual layer counting is the primary method used to date a core.

In Antarctica, the chronology of the high-accumulation WAIS Divide ice core (WD2014; Sigl et al. [2016]; Buizert et al. [2015]) is based on annual layer counting from 0 to 31 kyr

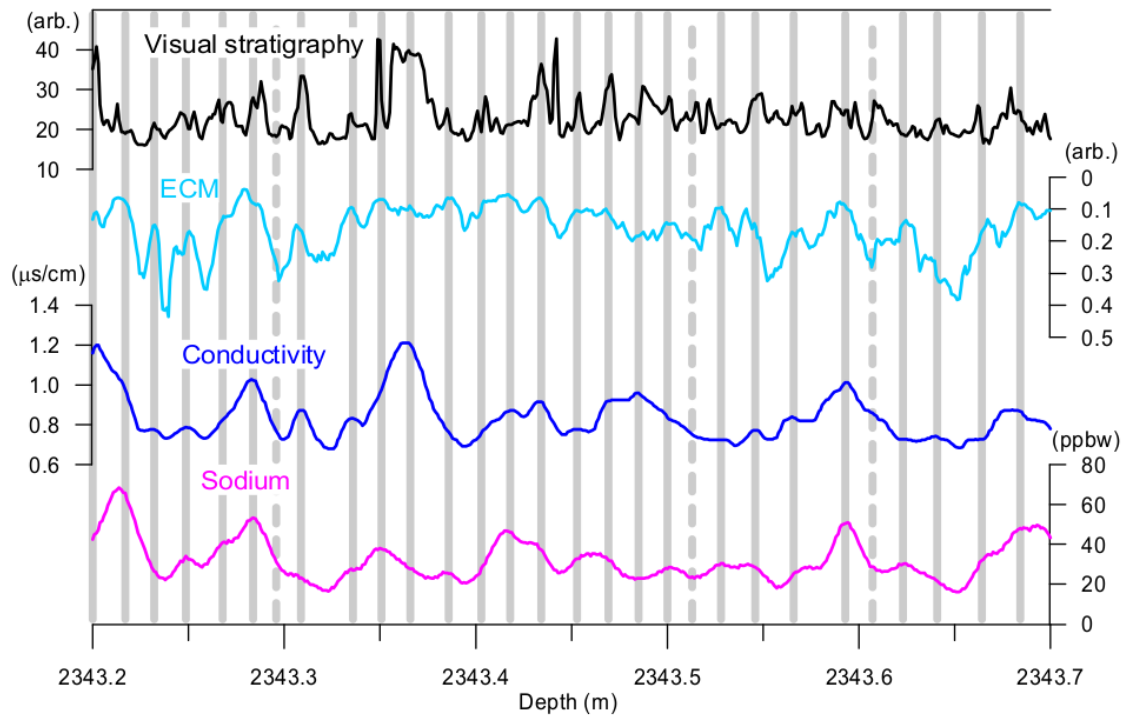


Figure 2: Example of annual layer counting from GICC05, within Greenland Interstadial 14 (54-49 kyr BP). The records are visual stratigraphy grey scale, ECM, conductivity, and Na^+ concentration. (Uncertain) Annual layers are indicated by (dashed) grey vertical bars. Units of the grey scale and the ECM profiles are arbitrary but comparable to those in Fig. 2. The counting of this section is mostly based on the conductivity and Na^+ profiles because the other records are known to have multiple peaks within an annual layer during milder climatic periods. Figure from Svensson et al. [2008].

BP. For the WD2014 chronology, the multi-parameter approach is used as well, with CFA and discrete measurements of water isotopes and major ions used along ECM and di-electric profiling measurements to count layers.

In contrast to the GICC05 chronology, which was developed using purely manual layer counting, in the layer-counted portion of the WD2014 chronology, the algorithms of Winstrup et al. [2012] or McGwire et al. [2011] were applied whenever possible. In some sections where annual layers became particularly difficult to identify, notably in a zone of brittle ice at 577–1300 m depth, manual counting was used extensively. Uncertainty was evaluated by comparison between methods and with other records.

The WD2014 chronology is a major stepping stone to the future of Antarctic ice core dating. It is by far the most precise chronology of a deep Antarctic ice core up to 31kyr BP, thanks to the application of layer counting. The data from the WAIS divide core itself, and the synchronization of other cores to WD2014, will allow for unprecedented detail in studies of Antarctic and global paleoclimate, particularly during the last deglaciation.

Orbital Tuning

Annual layer counting, as seen in the previous section, is largely based on the idea that certain indicators in ice cores respond to seasonal cycles in the climate system. Orbital tuning is based on a similar principle, but on much longer time scales: that proxies in ice cores can track the Earth’s orbital shape, and that the dating scenarios for these cores can be improved by matching the proxies to their respective orbital parameter targets. Where annual layer counting allows relatively new ice to be dated with great precision, orbital tuning allows ice to be dated deep into the past.

The shape of Earth’s orbit and rotation is altered as the gravitational pull from nearby planetary bodies itself changes in time. Milankovitch [1941] theorized that the major periodicities of these changes in shape, which cause insolation on the planetary surface to vary in time, could explain long-term variations in Earth’s climate. This theory has been confirmed by innumerable paleoclimate records, and the orbital and rotational frequency variations in Earth’s climate are now widely known as Milankovitch cycles.

Earth’s orbit can be calculated well into the past and future using fundamental physical principles. The insolation solution of Laskar et al. [2004], for example, numerically integrates a model of the solar system with nine major planetary bodies and the moon, and is precise enough to be used for calibrations of geologic records on the scale of 50 Ma (million of years); that is, nearly two orders of magnitude greater than the oldest ice core

record. We often speak of an “orbital clock”—that is, the variations in orbit and rotation provide us with a time reference with which to compare paleoclimatic and geologic records. However, as is the case for annual layer counting, we do not directly observe these orbital parameters in ice cores, but rather use proxy records which contain sources of uncertainty.

The isotopic ratio between ^{18}O and ^{16}O concentrations of foraminifera was first compared to orbital parameters in the context of dating deep marine sediment cores (see, for example, Hays et al. [1976], Imbrie et al. [1984]). In ice cores, the ratio between these isotopes in trapped gas, $\delta^{18}\text{O}_{atm}$ (or atmospheric $\delta^{18}\text{O}$, not to be confused with water/ice $\delta^{18}\text{O}$), is well-correlated with precession. Several early studies made note of this relationship, with the earliest dating scenarios tuning ice core $\delta^{18}\text{O}_{atm}$ to sea water $\delta^{18}\text{O}$ measured in ocean cores [Sowers et al., 1992, Bender et al., 1994]. [Waelbroeck et al., 1995] compared the Vostok δD record, which is closely related to $\delta^{18}\text{O}_{atm}$, to insolation filtered at the major periodicities of obliquity(41kyr) and precession (23kyr)¹.

Theoretically, $\delta^{18}\text{O}_{atm}$ is related to insolation via processes in the climate cycle. Since these processes affect the gas concentration when air is trapped, $\delta^{18}\text{O}_{atm}$ tuning is used to date the air contained in a core.

The chronological application of $\delta^{18}\text{O}_{atm}$ was first proposed for the deepened Vostok ice core by Jouzel et al. [2002], who found strong correlations between $\delta^{18}\text{O}_{atm}$ in the Vostok core and both 65°N maximum insolation and its filtered precessional component, when the core was placed on a previously developed chronology based on glaciological flow modeling. Similarly, Petit et al. [1999] took a strong correlation between $\delta^{18}\text{O}_{atm}$ and precession as a confirmation of their glaciological chronology. Neither of these studies, though, applied orbital tuning to create a dating scenario. The use of this method to uncover chronological clues was left to later studies.

Shackleton [2000] continued where these studies left off, developing a true orbital chronology using $\delta^{18}\text{O}_{atm}$ for Vostok by comparing its precession and obliquity components to the same components of a tuning target. Dating using orbital tuning was soon integrated with other methods as well: Parrenin et al. [2001] refined this chronology by using visual tie points between $\delta^{18}\text{O}_{atm}$ and insolation as control points for a glaciological

¹In addition to local insolation, a few of the parameters of Earth’s orbit and rotation have periodicities that make them of interest to studies of the climate of the Quaternary period. Eccentricity, which describes the degree to which Earth’s orbit is elliptical, is strongly periodic at 412 kyr, and weakly periodic at several small peaks between 95 and 131 kyr. Obliquity, the inclination of Earth’s rotational axis, shows a strong periodicity at 41kyr, and weaker periodicities at 39.6 and 53.7 kyr. Axial and apsidal precession, the changes in orientation of earth’s rotational axis and its orbital plane, show strong periodicities at 19 and 23 kyr.

model; Dreyfus et al. [2007] and Parrenin et al. [2007b] similarly used $\delta^{18}\text{O}_{atm}$ tie points, along with additional dated horizons, in the EDC3 chronology.

The relationship between $\delta^{18}\text{O}_{atm}$ and precession, while useful for dating the oldest ice contained in ice cores, is limited in precision by several sources of uncertainty. $\delta^{18}\text{O}_{atm}$ does not respond directly to insolation, but rather tracks responses to the insolation cycle within the climate cycle. Monsoon precipitation systems in the subtropics, which can lag insolation forcing by thousands of years, can impact $\delta^{18}\text{O}_{atm}$. A component of $\delta^{18}\text{O}_{atm}$ related to the 100-kyr cycles in ice sheet volume that occurred over the last million years, which do not agree in amplitude with any orbital parameter, is also well-known. The biosphere also influences $\delta^{18}\text{O}_{atm}$ by way of photosynthesis and respiration, a phenomenon known as the Dole Effect. The lag between $\delta^{18}\text{O}_{atm}$ and insolation cycles should vary with these mechanisms, but since the exact mechanisms are difficult to model, each of the four studies mentioned in the previous paragraph estimated a lag of 6 kyr between the precession component of insolation and $\delta^{18}\text{O}_{atm}$. Indeed, the lag itself between insolation and $\delta^{18}\text{O}_{atm}$ is known to vary [Jouzel et al., 2002], adding an important source of uncertainty to the method.

Using a $\delta^{18}\text{O}_{atm}$ -tuned chronology for the Vostok ice core, Bender [2002] observed that the ratio between O_2 and N_2 trapped in the air contained in the core ($\delta\text{O}_2/\text{N}_2$) seemed to be related with the on-site (78°S) summer insolation. More importantly, unlike the relationship between $\delta^{18}\text{O}_{atm}$ and precession, this relationship did not seem to show a lag. Measurements of $\delta\text{O}_2/\text{N}_2$ taken at Dome Fuji and EDC indeed showed the same phase agreement with local summer insolation. This localized relationship would minimize the influence of global climate on the signal, thus significantly reducing the uncertainty present in $\delta^{18}\text{O}_{atm}$ tuning.

A third orbital tuning method would be developed soon after: Kawamura et al. [2004] and Raynaud et al. [2007] observed that the Total Air Content (TAC) in the Dome Fuji and EDC ice cores also varied in phase with local summer insolation. Interestingly, different frequencies were observed to dominate the spectra of the two proxies: the strongest component in $\delta\text{O}_2/\text{N}_2$ is at the frequency of precession (23kyr); the strongest component in the TAC signal is at the frequency of obliquity (41kyr).

$\delta\text{O}_2/\text{N}_2$ and TAC record insolation via local, small-scale mechanisms, in clear contrast to $\delta^{18}\text{O}_{atm}$. Both the amount of air trapped in bubbles in the core (which determines TAC) and the size of the molecules trapped in the bubbles (which determines the O_2/N_2 ratio) are controlled by the process of pore close-off. The total amount of air trapped in a bubble

in an ice core depends on the porosity of the ice, with a weaker dependence on temperature and pressure at the time of close-off via the ideal gas law [Raynaud et al., 2007]. In the case of $\delta\text{O}_2/\text{N}_2$, smaller O_2 molecules continue to effuse near the lock-in depth, while larger N_2 molecules become trapped, depleting the $\delta\text{O}_2/\text{N}_2$ ratio [Severinghaus and Battle, 2006]. The physical properties of snow are hypothesized to determine the lock-in depth; these properties can be significantly affected by the snow surface energy balance, which can be largely driven by insolation [Libois et al., 2014]. Since $\delta\text{O}_2/\text{N}_2$ and TAC are related to snow processes, both tuning methods are used to date the ice phase of a core.

Orbital tuning using these localized proxies has led to advances in our understanding of Antarctic paleoclimate. When [Kawamura et al., 2007] adjusted the Dome Fuji, Vostok and EDC time scales by tuning $\delta\text{O}_2/\text{N}_2$ to local insolation, the better dating resolution deep in the cores allowed them to investigate relationships between Northern Hemisphere and Antarctic climate with unprecedented accuracy. In particular, carbon dioxide concentrations in the cores and Antarctic temperature as recorded in $\delta^{18}\text{O}_{ice}$ were found to lag Northern Hemisphere insolation by a few millenia during the last four deglaciations.

While the results of dating scenarios using $\delta\text{O}_2/\text{N}_2$ and TAC are promising, the hypotheses used to support matching these series to orbital parameters require additional confirmation. Since they are hypothesized to vary via similar mechanisms, they should produce dating scenarios that are in good agreement. Indeed, Lipenkov et al. [2011] showed that this was the case, within a range of about 1kyr, for the Vostok ice core. Of course, this result only confirms the methods with respect to each other, and that sources of error common to the two methods could additionally exist.

$\delta\text{O}_2/\text{N}_2$ and TAC series are also affected by several processes not related to insolation. A few examples are the impact of the accumulation rate via pressure differences and the change in depth of the non-diffusive zone (Hutterli et al. [2009]; Witrant et al. [2012]); microstructural changes in the firn that can occur well below the surface [Hörhold et al., 2009] and dust content [Freitag et al., 2013]. Bazin et al. [2016], in an extensive test of the phasing between orbital forcing, $\delta^{18}\text{O}_{atm}$ and $\delta\text{O}_2/\text{N}_2$, using new, high-resolution measurements made on the EDC ice core in addition to measurements from Vostok, indicated the presence of a 100kyr spectral peak in the $\delta\text{O}_2/\text{N}_2$ records that is not present in the summer insolation series. Together, these studies considerably complicate the hypothesis that $\delta\text{O}_2/\text{N}_2$ and TAC are purely local signals. Additionally, corrections for gas loss, which is exacerbated if cores are not stored at low enough temperature, are necessary for $\delta\text{O}_2/\text{N}_2$ records.

Models of the processes that link these three proxies to insolation will help clarify exactly how orbital tuning should be performed for the best possible accuracy. Fujita et al. [2009] have performed an extensive initial modeling study to this effect. It is also possible to automate the tuning process itself, in order to reduce the subjectivity inherent in visually picking tie points. Both Raynaud et al. [2007] and Lipenkov et al. [2011] used a method based on Continuous Wavelet Transform analysis to calculate time delays between the tuning proxy and the insolation target. Figure 3 shows the result of this method when applied to $\delta\text{O}_2/\text{N}_2$ and TAC data from the Vostok ice core, and two respective Integrated Summer Insolation (ISI) targets. The result is an estimate of time delay from the two series, which can then be subtracted from the base chronology to give a new, tuned chronology

Future developments in orbital tuning will likely include further advances in process-based modeling of pore closure and the automation of the tuning process; indeed, we will likely soon see combinations of the two. Thanks to new measurements of $\delta\text{O}_2/\text{N}_2$ and $\delta^{18}\text{O}_{atm}$ made on the EDC core [Bazin et al., 2016], we can investigate and apply all three orbital tuning methods well into the past.

Dated Horizons

Although the ice and air contained in ice cores are not easily dated using radiometric methods, some layers can be dated by relating them to strata in other, previously dated paleoclimate archives, like speleothems. Volcanic materials, the rapid changes in northern hemisphere climate during the last deglaciation known as Dansgaard-Oeschger events, and variations in solar activity and of earth’s magnetic field can all be used to provide individual age points for a core.

Ice cores can often be synchronized with radiometrically dated paleoclimate archives using volcanic events². Major volcanic eruptions eject large amounts of volcanic material

²Radiometric dating methods are standard in the broader field of geochronology. These methods rely on the decay of radioactive isotopes with lifetimes long enough to cover the timescale of the record for which an age is desired. Using exponential relationships known as decay equations, we can model the ratios of radioactive parent isotopes to their decayed counterparts, known as daughter isotopes. The values of the constants in the decay equations, however, are not known. To estimate the values of these constants, reference materials must be chosen in well-constrained sedimentary sequences: that is, we require measurements of parent and daughter isotopes in a material whose age is already well-known. The standard reference material for $^{40}\text{Ar}/^{39}\text{Ar}$ dating, for example, is the Fish Canyon sanidine, whose age is constrained by matching stratigraphic succession in Morocco to oceanic sediments, which are then dated

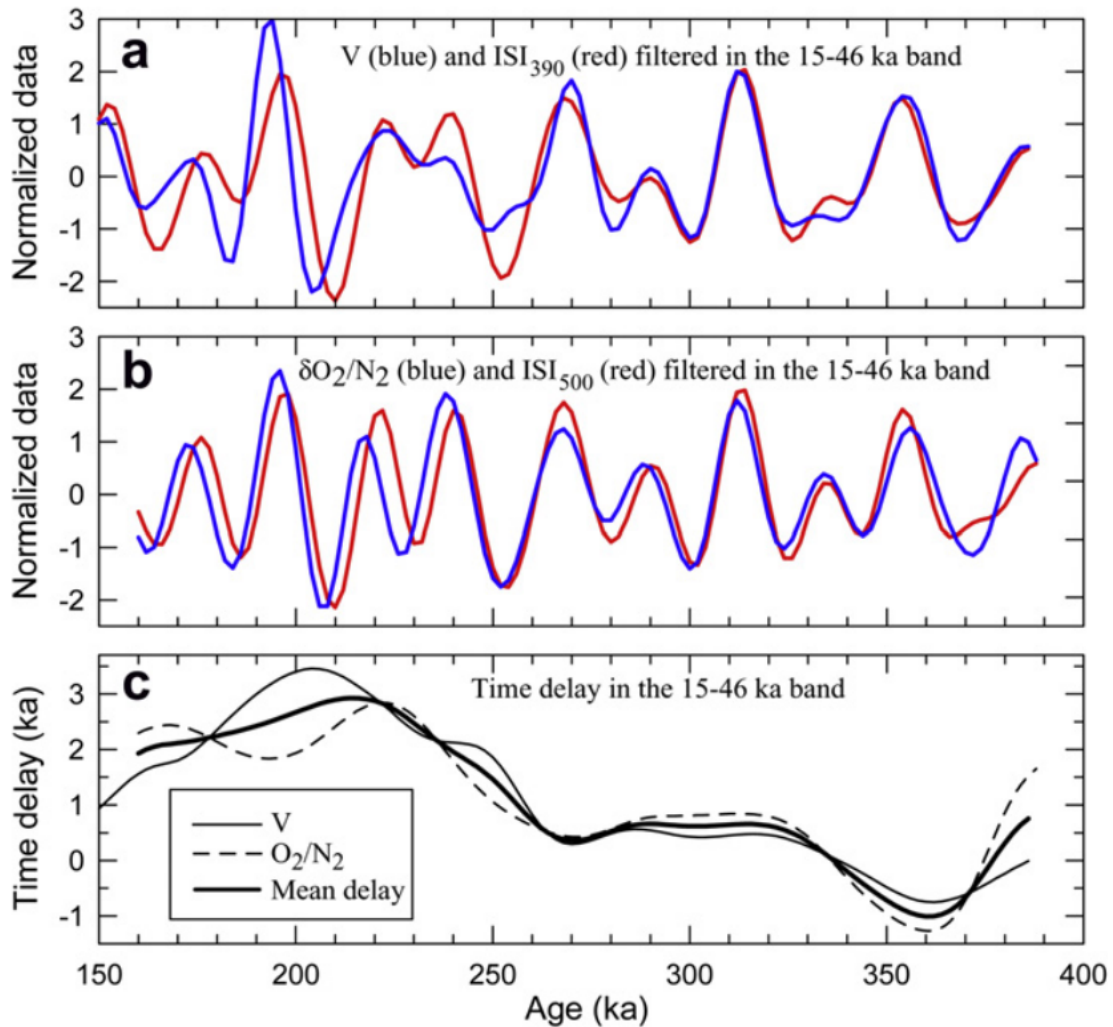


Figure 3: An example of automated orbital tuning, based on the wavelet transformation, from the Vostok ice core: Time series of TAC (V in the figure), $\delta O_2/N_2$, and their Integrated Summer Insolation (ISI) targets; calculated time delays. The mean time delay is shown by the thick curve. All time series are filtered in the 15-46 ka band. Figure from Lipenkov et al. [2011].

into the high troposphere and even the stratosphere. This volcanic material includes glass, minerals and aerosols, and is deposited onto the surface as tephra, which are visible in many paleoclimate records including ice cores [Gao et al., 2006]. Near the source volcano, the material embedded in the tuffs, sedimented igneous rocks made of deposits of volcanic ash, associated with the tephra is often abundant enough to be dated radiometrically. The use of tephra and the corresponding tuff layers to date paleoclimate records is known as tephrostratigraphy. The major and trace element composition of the same tephra horizon observed in two different paleoclimate records should be coherent, allowing us to relate records to each other.

The sulfate aerosols produced during the eruption process are rapidly oxidized in the atmosphere to form sulfuric acid, which is in turn deposited with precipitation on the surface of ice sheets close enough to the volcano [Gao et al., 2006]. The characteristic acidity of this sulfuric acid-rich snow is retained as the snow densifies into firn, and finally into the ice which composes a core. These acid peaks can be detected in ECM and Di-electric profiling (DEP) records, preserving the ice for additional use (in addition to the possibility of directly measuring sulfate). Some acid peaks, in the age range of the historical record or of massive, well-known eruptions, can be traced to known volcanic events using the geochemical signature of the tephra [Narcisi et al., 2006]. In addition to serving as dated horizons, tephra can also be used for the synchronization of ice cores: the GICC05 dating scenario, for example, uses five tephra corresponding to volcanic events around the North Atlantic to align Greenland ice cores to a single layer counted chronology [Svensson et al., 2008].

The eruption of the Toba supervolcano in Sumatra around 74,000 years ago is perhaps the most powerful known volcanic event of the Quaternary period, and deposited volcanic material in both hemispheres. This eruption is dated within a small range of uncertainty using astronomically calibrated $^{40}\text{Ar}/^{39}\text{Ar}$ [Storey et al., 2012]. Zielinski et al. [1996] (1994) first proposed a tephra layer corresponding to the Toba eruption in the GISP2 Greenland ice core, though at the time it had been dated to 74 ka BP. Svensson et al. [2013], taking the updated radiometric dating information into account, identify a series of

by orbital tuning. The age assigned to this layer is 28.201 ± 0.046 Ma [Kuiper et al., 2008]. The parallel between the orbital methods used for these older geologic records and the methods used to date ice cores is too apparent to leave unmentioned. Note, for example, that the error in the age assigned to the Fish Canyon sanidine is on the same order of magnitude as the uncertainties when dating ice cores using $\delta^{18}\text{O}$! However, since decay relationships are exponential, we can estimate the age of materials younger than the reference material with much greater precision.

nine peaks corresponding to this eruption in the layer-counted NGRIP core in Greenland and the EDML core in Antarctica. The Toba eruption can thus serve as a valuable interhemispheric reference point for dating the ice phase.

Narcisi et al. [2006] identify another major tephra layer in the EDC and Dome Fuji ice cores, at 1265 and 1362 m depths, respectively, corresponding to an eruption of the Mt. Berlin Volcano in West Antarctica. This eruption has since been dated to 92.1 ka, with a ± 0.9 ka range of uncertainty, making it a useful, precise constraint for chronologies built on orbital tuning or ice flow modeling.

The rapid variations in Greenland climate during the last deglaciation known as Dansgaard-Oeschger events can also be used to assign ages to ice core layers. These events, though first identified, in the Greenland isotope record, are also well-recorded in the global methane record (supporting the bipolar seesaw hypothesis). They are resolved in other paleoclimate archives as well, notably in the speleothem isotopic record. The $\delta^{18}\text{O}$ record from the Chinese Hulu cave speleothem, which is dated using U-Th radiometry, clearly resolves the sequence of Dansgaard-Oeschger events from the last deglaciation. These events can thus be assigned ages as well.

Recently, Buizert et al. [2015] synchronized the Antarctic WAIS divide core to the Hulu cave chronology using the methane record, expanding on the observation that the Hulu cave $\delta^{18}\text{O}$ and WAIS divide methane series both recorded the Dansgaard-Oeschger events to hypothesize that they recorded the same climatic signals in general. As we will see in more detail in the next section, this introduces some uncertainty into the WAIS divide gas chronology, but perhaps more importantly allows for an estimation of Δdepth entirely independent of firn modeling, and hence free of the associated uncertainties.

Ice core strata can also be assigned absolute ages using Beryllium-10 (^{10}Be), which is among the class of isotopes created by interactions with cosmic radiation in the upper atmosphere. These isotopes are collectively known as cosmogenic isotopes or cosmogenic nuclides. The variations of Beryllium-10 in particular are affected almost exclusively by nucleosynthesis in the upper atmosphere, which is modulated by variations in solar activity and of the earth's magnetic field. These variations are also present in cosmogenic isotopes affected by other processes, such as carbon-14 (^{14}C), whose atmospheric concentration is additionally linked to the terrestrial carbon cycle. ^{10}Be can be precisely measured in ice cores, and contributes several sources of dating information. Cores can be synchronized using ^{10}Be ; dated magnetic and solar anomalies can be identified as ^{10}Be peaks; and major variations in ^{10}Be can even be linked to variations in ^{14}C in tree ring archives.

The Laschamp magnetic event, for example, is a marked, very brief reversal in the Earth’s magnetic field dated radiometrically to $40,650 \pm 950$ yr BP [Singer et al., 2009]. During the reversal, the earth’s magnetic field weakened considerably, allowing for greater interactions with cosmic radiation in the upper atmosphere. This event can thus be identified as a ^{10}Be peak in the Vostok ice core [Yiou et al., 1997]; in Greenland [Yiou et al., 1997]; and in the EDC core [Raisbeck et al., 2006].

Raisbeck et al. [2006] identify a series of sharp ^{10}Be peaks is found between 3,160 m and 3,170m depth in the EDC core to correspond with the Matuyama-Brunhes magnetic reversal, the latest total reversal of Earth’s magnetic field.³ This provides an important age marker for the oldest part of the core, confirming that the oldest ice drilled at EDC was indeed more than 800 ka old.

In general, improvements in dating external geochronologic markers will help better constrain ice core chronologies as well. In turn, the precise dating of ice cores can also verify the radiometrically calculated ages of the dated horizons identified within. When dealing with geochronologic markers, of course, we must be cautious to not apply circular reasoning: that is, we cannot use ice cores to constrain the ages of dated horizons which are used to date the ice core in the first place!

Synchronization of Ice Cores

The information available to date individual ice cores, as we have seen in the previous sections, can be sparse. This is most true for the low-accumulation cores taken in the Antarctic interior, like the Vostok and EDC cores, which are also the oldest available records. To alleviate some of the uncertainty in dating these records, they can be synchronized with other ice cores, for which higher-resolution age information, like annual layer

³At the time of Raisbeck et al. [2006]’s study, the Matuyama-Brunhes reversal was placed at 776 ± 12 ka BP by radiometric dating, with an orbitally tuned estimate of 778 ± 2 ka BP. The latest estimate of the Matuyama-Brunhes boundary age, from [Mark et al., 2017], uses ^{40}Ar dating of volcanic tuffs, including material from a much older eruption of the same Toba volcano whose 74 ka eruption has been identified in both Greenland and Antarctic ice cores—in combination with oceanic sediment records where these tuffs lie on either side of the reversal, to constrain this age to 783.4 ± 0.6 ka BP. This new, extremely well-constrained age will undoubtedly be used in future chronologies for the EDC core, and for older ice cores if and when they are drilled. The ^{10}Be peaks identified by Raisbeck et al. [2006] should also be reconsidered in future chronologies. The presence of a ^{10}Be flux peak approximately 20 ka BP before the Matuyama-Brunhes boundary in the EDC record is not coherent with measurements made on ocean cores [Simon et al., 2016]. As such, an additional correction for environmental sources of ^{10}Be should be considered for future chronologies.

counts, are available. Cores can also be synchronized to other paleoclimate records, like speleothems or deep ocean sediment cores, which can add valuable dating information as well.

The records used to synchronize two cores should have a few key characteristics. These records should be global in nature or at least reasonably common to the two sites. In practice, we look for records for which the influence of on-site conditions like local temperature and precipitation or ice densification on the records should be as minimal as possible. In the case of gas-phase synchronization, the gases should be globally well-mixed: that is, not presenting large spatial variations in the atmosphere, which relies on them having an atmospheric lifetime longer than the timescale of global transport. On the other hand, if this atmospheric lifetime is too long, the absence of rapid variations makes synchronization less precise.

In the ice phase, the records that best fit these characteristics are volcanic events, which are marked in ice cores by peaks in acidity, and dust content. In the gas phase, methane in particular fits the profile of a gas that is well-mixed but varies on short enough timescales to allow for precise synchronization [Blunier et al., 2007].

Volcanic tephtras which can be identified and traced to a volcanic eruption can be assigned absolute ages, as discussed in the previous section. Many tephtras present in ice cores, however, cannot be traced to a volcano, and thus cannot be dated absolutely. These layers, if they can be identified in multiple cores, can be used to minimize disagreements in age models. The peaks in ECM, DEP, or sulfate measurements that come from volcanic eruptions are often orders of magnitude stronger than background noise, and tephtra deposits usually only last for one or two years. This allows the cores to be synchronized with relative certainty. Cores can similarly be synchronized using ^{10}Be measurements, even when the peaks themselves cannot be identified (see, for example, Raisbeck et al. [2006]). On the other hand, a degree of uncertainty is introduced because of the possibility of matching two unrelated tephtra layers. This is complicated by the fact that the same volcanic event in two ice cores does not necessarily produce time series of peaks of the same magnitude, because of differences in transport and sedimentation processes.

Figure 4 shows an example of volcanic synchronization between the EDC and Vostok ice cores. Tie points, shown as red diamonds, are selected at major peaks in the ECM, DEP and sulfate records. The synchronization is confirmed by the match in Deuterium records shown in the last panel. Note, though, that these peaks occur with varying frequency, and that matching smaller peaks in these records would be considerably more difficult. In




Figure 4: Figure 4. An example of volcanic synchronization, between the EDC and Vostok Ice cores, during the period 65–80 kyr BP, centered on the likely age of the Toba eruptive event. All records have been placed on the EDC3 age scale. (Top) Black: ECM record from the VK-5G ice core. Blue: ECM record from the EDC99 ice core. Purple: DEP record from the EDC99 ice core. Green: sulfate records from the EDC99 ice core. Red diamonds: volcanic tie points. For simplicity, the y-axes are not plotted. (Bottom) Blue: Deuterium record from the EDC ice core. Green: Deuterium record from the Vostok ice core. For simplicity, the y-axes are not plotted. Figure from Parrenin et al. [2012].

addition, the smaller peaks in acidity may represent non-volcanic, local sources.

Another limitation in volcanic synchronization comes from atmospheric circulation. Aerosols are blocked by the circulation barrier at the equator, so that only the strongest volcanic eruptions are present in cores from both hemispheres [Gao et al., 2006]. As such, volcanic peak matching is most commonly used to synchronize cores from the same ice sheet.

Volcanic synchronization has been applied to most of the major ice cores: important examples include the synchronization of the Epica Dronning Maud Land (EDML) core [Severi et al., 2007] and the Talos Dome ice core (TALDICE) to the EDC3 timescale [Severi et al., 2012]; the synchronization of the EDC and Vostok cores [Parrenin et al., 2012]; and the synchronization of the Dome Fuji and Dome C cores [Fujita et al., 2015]. Volcanic synchronization will likely be the principal method used to synchronize Antarctic ice cores with the layer-counted portion of the WD2014 time scale.

Just as volcanic synchronization can be used to align ice core age scales in the solid phase, CH₄ synchronization can be used to align air age scales. The lifetime of methane in the atmosphere is on the order of ten years. Methane concentrations can thus accurately track changes in terrestrial sources or sinks, which themselves respond to rapid climate change, without presenting major variations related to pole-to-pole atmospheric transport, whose timescale is on the order of a year [Blunier et al., 2007]. A gradient between methane concentrations exists between northern hemisphere and southern hemisphere ice cores, as a result of the differences in sources and sinks between the two hemispheres. This gradient can vary between climate intervals, and is largest during warm periods due to stronger northern hemisphere sources [Dällenbach et al., 2000]. However, this gradient is small when compared to the concentrations themselves. The application by Blunier et al. [2007] of CH₄ synchronization of the EDML record to a composite Greenland record on the GICC05 timescale is a good example of the utility of this tool for interhemispheric core synchronization.

A few sources of uncertainty exist for methane synchronization (and synchronization in general in the gas phase). The first source comes from the method itself: outside of very major changes, the match between the two series can be ambiguous. The second source is due to the process by which air is trapped in the ice core. Because air continues to diffuse into pores during the close-off process, the age of the air at a given depth in an ice core is not unique, but rather a distribution of ages. Since the close-off processes depend on local snow characteristics, these air age distributions are different between cores [Köhler,

2010]. This means that matching gas concentrations is limited in terms of precision to the time-scale of the close-off process, which can be on the order of decades to centuries.

Methane synchronization is also an important component of the WD2014 time scale [Buizert et al., 2015]. WD2014, however, is not synchronized to methane from other ice cores, making the example a variation of the classic method of Blunier et al. [2007]. For the section of the core older than 31ka, where annual layers become difficult to count, the air in the WAIS divide core is dated by synchronizing methane measurements to $\delta^{18}\text{O}_{ice}$ measurements from the NGRIP core in Greenland, [NGRIP Community Members, 2004] which is on the layer-counted GICC05 timescale, and additionally to the Hulu Cave speleothem $\delta^{18}\text{O}$ record (Wang et al. [2001] among many others). The synchronization of CH_4 in the WAIS divide core to both NGRIP and Hulu Cave $\delta^{18}\text{O}$ poses advantages as well in terms of minimizing dating uncertainty. First, NGRIP $\delta^{18}\text{O}_{ice}$ is dated directly by layer counting, without adding the uncertainty involved in estimating Δage for the NGRIP methane record. Second, including the Hulu Cave record, which is dated radiometrically (and thus is dated precisely in terms of absolute age), helps to compensate for the cumulative error inherent in layer counting.

However, this synchronization makes the assumption that the rapid changes in $\delta^{18}\text{O}$ recorded in the Hulu Cave speleothem and in the NGRIP core are associated with the same climatic events as the rapid changes in methane recorded by the WAIS divide core. This introduces some uncertainty due to the different sources of information in the records—indeed, $\delta^{18}\text{O}$ can record regional-scale variations. It should be emphasized that speleothem isotopic records themselves often differ, and synchronization to two different speleothems may give different results (e.g. Fleitmann et al. [2009]).

$\delta^{18}\text{O}_{atm}$ has been used to synchronize records in the gas phase as well, particularly in earlier studies [Blunier, 2001]. Its lifetime in the atmosphere is considerably longer than that of methane, often making quick transitions more difficult to identify and match. However, it shows less of an atmospheric gradient (i.e. its interhemispheric mixing time is much shorter), and can be useful to clarify synchronization when the match between methane time series is ambiguous [Capron et al., 2010]. This is often the case when millennial-scale variability is not visible, as is the case for full glacial periods like Marine Isotopic Stage (MIS) 2 and late MIS 6.

Extier et al. [2018] develop an improved approach to synchronize ice cores with speleothems by matching $\delta^{18}\text{O}_{atm}$ in ice cores to speleothem calcite $\delta^{18}\text{O}$. These two isotopic records present shared orbital-scale and millennial-scale features, which may be related

to the position of the Intertropical Convergence Zone (ICTZ). Speleothems can often be absolutely dated, and the age uncertainty of the match between speleothem calcite $\delta^{18}\text{O}$ and $\delta^{18}\text{O}_{atm}$ is assessed to be much lower than that between $\delta^{18}\text{O}_{atm}$ and the insolation targets used for orbital tuning.

A few research challenges could make significant improvements to the synchronization of ice cores. One of these is the mathematical automation the synchronization of records, which is treated in this thesis. In the method we develop, we focus on providing an objective assessment of the methodological uncertainty of the match. In the gas phase, this uncertainty can be difficult to accurately quantify because of the age uncertainty related with the air lock-in process, and modeling studies focused on this process will certainly improve the quantification of the uncertainty involved in synchronizing cores in the gas phase.

Bayesian dating models: a federative approach

To most comprehensively date an ice core, the dating methods described in this chapter should be used in conjunction with each other. For the earliest ice core chronologies, this was often done implicitly: a dated horizon was used to check the result of an ice flow modeling simulation, for example. Afterwards, if the simulation did not match the dated layer, the parameters of the ice flow model would be adjusted until an agreement was obtained. This result could then be checked against an orbital tuning chronology, or links with another core, and so on. However, as we consider more cores and dating methods, the manual adjustment and verification of chronologies becomes more and more time-consuming. This motivates a formalized, mathematical approach.

Including multiple sources of information in an ice core chronology falls under the mathematical domain of inverse problems. In the first section of this chapter, we used a few equations to describe the evolution of an ice core over time. We can call this set of equations the forward model. Starting with some inputs, we can use the model to arrive at a state in the future: that is, we run the model in the forward direction in time. In practice, starting from accumulation, lock-in depth, and the thinning function, we can follow ice horizons as they are embedded deeper in the core.

Though we have some prior knowledge of the accumulation, lock-in depth, and thinning function at a given site, we certainly do not know them exactly. We also know the ages of the ice and gas at certain depths. The age of the ice might be identified by layer counting, or at a peak in acidity corresponding to a known volcanic eruption. The age of an interval

in the gas phase might be well-known during a Dansgaard-Oeschger event. We may know that two layers in two ice cores are synchronous, without knowing their absolute age. But these pieces of information, which tell us where a layer of ice or air with a given age has arrived at depth, are not inputs to the forward model. They are, in fact, better described by the outputs!

To improve our chronology, we want to effectively run our model in the inverse direction in time: starting from the dated layers in the ice core, we want to arrive at the most probable scenarios of accumulation, lock-in depth, and the thinning function, in order to construct a chronology for the full core. Unlike the forward model, the inverse problem by nature does not have a unique solution. Thus, we work with probabilities. Rather than a single, absolute chronology, we would like to have an idea of the probability, given the available information, of any given chronology.

Bayesian inference allows us to infer the probability of a dating scenario from the information that we have available: the chronological information for the ice cores we would like to date, and our prior knowledge of the forward model parameters. We provide a more mathematical description of the Bayesian formulation of inverse problems in the specific context of ice core dating in Chapter 2. For a more in-depth treatment of inverse modeling in general, we refer the reader to Tarantola [2005].

Let us group all the forward model parameters—accumulation, lock-in depth, and the thinning function, for example—together, as a forward model vector m in the model space M . We have some prior knowledge about these parameters. Let us say, for example, that we have an idea of accumulation throughout the core from the δ_D isotopic record, that we have estimated lock-in depth using Nitrogen isotopes, and that model simulations have been run that give us an idea of the thinning function at the site. We can formally describe this prior knowledge as a probability distribution ρ_M , which specifies how likely a given model vector is based on our prior ideas of the individual parameters.

Let us also group all the chronological information together as an observation vector d . Then, we can also identify a probability distribution ρ_D , which tells us how likely a given model vector is based on the chronological information.

Using Bayes' theorem, we can calculate the posterior probability distribution of the model vector, σ_M , as follows [Tarantola, 2005]:

$$\sigma_M = k\rho_D(m) \cdot \rho_M(M) \tag{7}$$

σ_M , in practice, gives us the probability of any model vector, and the associated ice

core chronology, given all the prior and observational information that we have.

This framework allows us to set up the inverse problem. First, we must define the variables included in the model vector and the associated prior distributions. Second, we must specify the observational data used to define the probability distribution of the data. Third, we must develop a mathematical method to solve the inverse problem. In the last decade, each three of these aspects have seen significant development with respect to the inverse problem of dating ice cores.

The first probabilistic ice core dating experiments using a formal inverse modeling framework were made by Parrenin et al. [2004] for the Vostok core, Parrenin et al. [2007b] and Parrenin et al. [2007a] for the EDC and Dome Fuji cores. These experiments did not use the three parameters we describe as the base of a chronology above, but rather inverted the control parameters of a flow model, an accumulation model, and a firn densification model. This approach greatly reduces the number of individual number of parameters to be inverted, as we do not have to find the values of accumulation or thinning at each point in the core for the resolution we want to invert, but rather a few constants from which we solve for these values.

In the setup of an inverse problem, the “goodness of fit” of a proposed set of model parameters to the available data and prior parameter information must be quantified. This quantification is referred to as a cost function. The cost function is often formulated as the sum of residuals, that is, of the errors in age, related to the data and the prior when a certain set of model parameters is proposed. It should additionally take the autocorrelation of these residuals, and the uncertainty associated with modeling and measurements, into account. Essentially, the cost function measures how far a model is from the information available.

In the studies above, the Metropolis-Hastings algorithm, a stochastic search requiring thousands of proposed scenarios for each parameter to be evaluated, was used to minimize the cost function. A benefit of stochastic searches is that they are global in nature: that is, any set of parameters can potentially be proposed. Additionally, the Metropolis-Hastings algorithm characterizes the posterior probability distribution of the parameters evaluated, making it a Bayesian inversion method. If the parameter set is small, the Metropolis-Hastings algorithm can be applied without requiring a prohibitive amount of computational resources.

However, reducing the parameters of an ice core dating scenario to a few control parameters related to accumulation, flow, and firn densification has a few drawbacks. The

flow, accumulation, and firn densification models make assumptions that are not perfect. The uncertainties contributed to the dating scenario by these assumptions can be hard to quantify.

Lemieux-Dudon et al. [2010] developed an approach to bypass this uncertainty in the Datice ice core dating model. Rather than calculating the accumulation, thinning function and lock-in depth, these variables, which are the base of any chronology (as discussed in the first section) could be inverted directly. Information gained from modeling each of the variables could still be used: rather placing these models within the inversion, they were applied to develop the prior probability distribution of each variable.

The drawback of this method is that the number of parameters to be inverted grows substantially, depending on the resolution at which we date the ice core. As such, a Metropolis-Hastings stochastic search was no longer computationally feasible for Lemieux-Dudon et al. [2010]. Rather, the authors proposed an analytical gradient of the cost function which defines the posterior probability distribution, greatly reducing computation time. However, the computation of posterior uncertainty required some assumptions to be made about the shape of the probability distribution, which are not necessary when using a stochastic search.

The Datice model was used by Veres et al. [2013] and Bazin et al. [2013] to create the Antarctic Ice Core Chronology 2012 (AICC2012), one of the most current chronologies for deep Antarctic ice cores. The AICC2012 experiment was a major milestone in Bayesian dating. In addition, it is a federative dating experiment: that is, it dates and uses information from and links between four Antarctic ice cores: European Project for Ice Coring in Antarctica (EPICA) Dome C (EDC), Vostok, Talos Dome (TALDICE), and EPICA Dronning Maud Land (EDML), and the North Greenland Ice Core Project (NGRIP) core from Greenland. AICC2012 synthesized several, if not all, of the sources of dating information presented in this chapter. At the time of writing, AICC2012 is still a reference chronology for the four Antarctic cores involved.

A technical advance in the field of Bayesian and federative dating was made by Parrenin et al. [2015], who developed the IceChrono model. Based on similar principles as Datice, IceChrono is built to be more concise (reducing the number of code lines by a factor of ten), and is coded to be modular: that is, it is flexible in terms of the optimization method used to solve the inverse problem, rather than being constrained to gradient-based optimization. This opens up several possibilities in terms of considering new information.

With federative dating models like IceChrono, we can provide not only an optimal

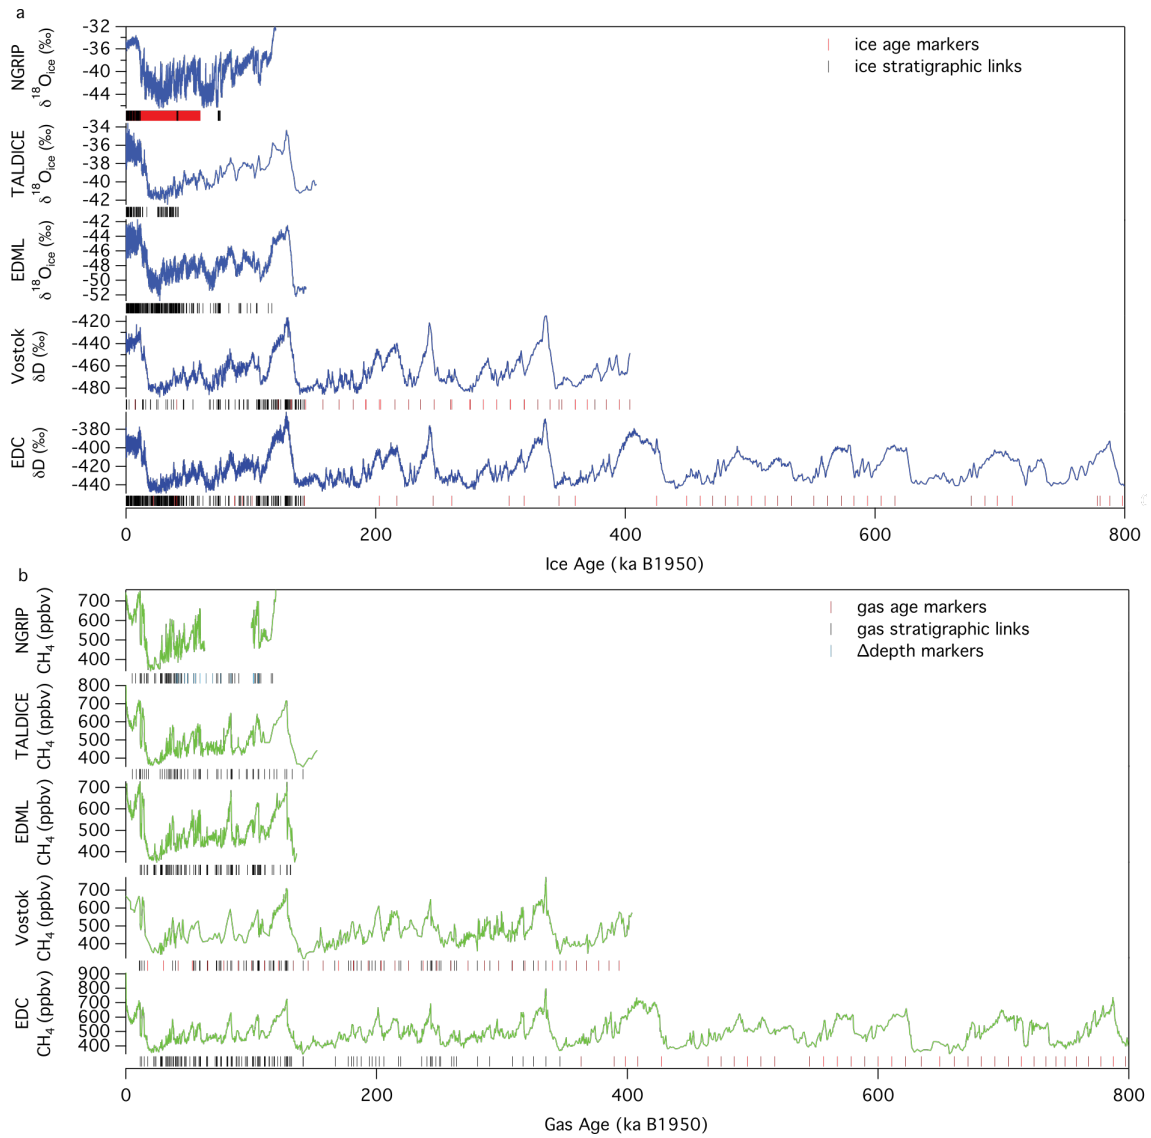


Figure 5: (a) Water stable isotope records and (b) Methane records of the cores dated in the AICC2012 dating experiment on the AICC2012 age scale. Stratigraphic links and age marker positions are displayed under each core. From Bazin et al. (2013).

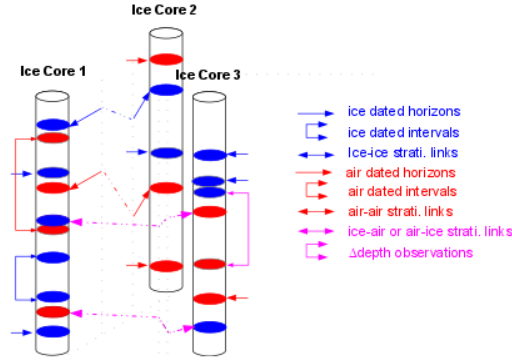


Figure 6: The sources of chronological information included in the IceChrono 1.0 dating model [Parrenin et al., 2015].

chronology for a set of ice cores, but a formalized statistical estimate of the uncertainty associated with that chronology as well. Federative models make improvements both in reproducibility and in work-hours over manual integrative dating.

In this thesis, we tackle several challenges in federative dating. We attempt to improve uncertainty estimates by using stochastic optimization methods, like Metropolis-Hastings simulations, to relax the assumption of a Gaussian posterior distribution. Further advances will follow developments in applied mathematics and computation, which will reduce the computational intensity of stochastic optimization.

Perspectives and Motivation

The problem of dating ice cores is far from being fully solved. Multiple useful sources of information have been uncovered, which bear clues about the age of the ice and air in a core. However, many uncertainties still remain, and identifying and accurately representing the sources of these uncertainties is perhaps the most salient challenge in ice core dating.

As we strive to accurately represent and minimize uncertainty, dating ice cores becomes an even more distinctly interdisciplinary problem. Improving firn modeling, for example, and better understanding the mechanisms that link earth’s orbital parameters to their tracers in ice cores, are problems strongly rooted in physics. The state of the art in annual layer counting has been reached thanks to both gas and isotope chemistry and probabilistic mathematics. The identification of dated horizons like the Matuyama-Brunhes boundary will be made more accurate by advances in geochronology; indeed, the

development of absolute dating methods for ice cores themselves is underway (e.g. Buizert et al. [2014]). The automation of manually-performed dating methods like layer counting, synchronization and orbital tuning will allow us to quantify the uncertainties inherent to the dating methods themselves; and indeed will allow these methods to be incorporated formally into the federative dating inverse problem.

This last point is addressed in this thesis. Over three years of research, I have worked with a team of three advisors and numerous colleagues to develop an automated method for synchronization, potentially applicable to layer counting and orbital tuning as well, that is incorporated in the federative dating inverse problem. This method is incorporated into a new version of IceChrono, now called PaleoChrono, as it is capable of treating records other than ice cores as well. The method, its application to methane synchronization, volcanic synchronization and orbital tuning are all treated in the following chapters. I present, as well, an example of a study that uses Bayesian principles to accurately represent chronological uncertainty about the timings of changes in CO₂ and Antarctic Temperature during the last deglaciation.

Of course, some uncertainty in dating will always remain. We cannot observe the passage of time directly in an ice core, but must rely on the often imperfect tracers it leaves behind. But as we learn to use these tracers more and more effectively, we approach a more temporally accurate understanding of what is perhaps our best record of climatic mechanisms in the past.

References

- Katrine K Andersen, Anders Svensson, Sigfus J Johnsen, Sune O Rasmussen, Matthias Bigler, Regine Röthlisberger, Urs Ruth, Marie-Louise Siggaard-Andersen, Jørgen Peder Steffensen, Dorthe Dahl-Jensen, et al. The Greenland ice core chronology 2005, 15–42 ka. part 1: constructing the time scale. *Quaternary Science Reviews*, 25(23-24): 3246–3257, 2006.
- MO Battle, JP Severinghaus, ED Sofen, D Plotkin, AJ Orsi, M Aydin, SA Montzka, T Sowers, and PP Tans. Controls on the movement and composition of firn air at the west Antarctic ice sheet Divide. *Atmospheric Chemistry and Physics*, 11(21):11007–11021, 2011.
- Lucie Bazin, Amaëlle Landais, Bénédicte Lemieux-Dudon, Habib Toye Mahamadou Kele, D Veres, Frédéric Parrenin, Patricia Martinerie, Catherine Ritz, Emilie Capron,

- Vladimir Lipenkov, et al. An optimized multi-proxy, multi-site Antarctic ice and gas orbital chronology (AICC2012): 120-800 ka. *Climate of the Past*, 9(4):1715–1731, 2013.
- Lucie Bazin, Amaelle Landais, Emilie Capron, Valérie Masson-Delmotte, Catherine Ritz, Ghislain Picard, Jean Jouzel, Marie Dumont, Markus Leuenberger, and Frédéric Prié. Phase relationships between orbital forcing and the composition of air trapped in Antarctic ice cores. *Climate of the Past*, 12(3):729–748, 2016.
- Michael Bender, Todd Sowers, Mary-Lynn Dickson, Joseph Orchardo, Pieter Grootes, Paul A Mayewski, and Debra A Meese. Climate correlations between Greenland and Antarctica during the past 100,000 years. *Nature*, 372(6507):663, 1994.
- Michael L Bender. Orbital tuning chronology for the Vostok climate record supported by trapped gas composition. *Earth and Planetary Science Letters*, 204(1-2):275–289, 2002.
- Carl S Benson. Stratigraphic studies in the snow and firn of the greenland ice sheet. Technical report, COLD REGIONS RESEARCH AND ENGINEERING LAB HANOVER NH, 1962.
- Matthias Bigler, Anders Svensson, Ernesto Kettner, Paul Vallelonga, Maibritt E Nielsen, and Jørgen Peder Steffensen. Optimization of high-resolution continuous flow analysis for transient climate signals in ice cores. *Environmental science & technology*, 45(10):4483–4489, 2011.
- T. Blunier. Timing of millennial-scale climate change in Antarctica and Greenland during the last glacial period. *Science*, 291(5501):109–112, jan 2001. doi: 10.1126/science.291.5501.109.
- T. Blunier, R. Spahni, J.-M. Barnola, J. Chappellaz, L. Loulergue, and J. Schwander. Synchronization of ice core records via atmospheric gases. *Climate of the Past*, 3(2):325–330, jun 2007. doi: 10.5194/cp-3-325-2007.
- Camille Bréant, Patricia Martinerie, Anaïs Orsi, Laurent Arnaud, and Amaëlle Landais. Modelling firn thickness evolution during the last deglaciation: constraints on sensitivity to temperature and impurities. *Climate of the Past*, 13(7):833–853, 2017.
- C Buizert, KM Cuffey, JP Severinghaus, D Baggenstos, TJ Fudge, EJ Steig, BR Markle, M Winstrup, Rachael H Rhodes, EJ Brook, et al. The WAIS Divide deep ice core WD2014 chronology—part 1: Methane synchronization (68–31 ka bp) and the gas age–ice age difference. *Climate of the Past*, 11(2):153–173, 2015.

- Christo Buizert and Jeffrey P Severinghaus. Dispersion in deep polar firn driven by synoptic-scale surface pressure variability. *The Cryosphere*, 10(5):2099–2111, 2016.
- Christo Buizert, Daniel Baggenstos, Wei Jiang, Roland Purtschert, Vasilii V Petrenko, Zheng-Tian Lu, Peter Müller, Tanner Kuhl, James Lee, Jeffrey P Severinghaus, et al. Radiometric 81kr dating identifies 120,000-year-old ice at Taylor Glacier, Antarctica. *Proceedings of the National Academy of Sciences*, 111(19):6876–6881, 2014.
- E Capron, A Landais, B Lemieux-Dudon, A Schilt, Valérie Masson-Delmotte, D Buiron, J Chappellaz, Dorthe Dahl-Jensen, S Johnsen, M Leuenberger, et al. Synchronising EDML and NorthGRIP ice cores using $\delta^{18}\text{O}$ of atmospheric oxygen ($\delta^{18}\text{O}_{\text{atm}}$) and CH_4 measurements over MIS 5 (80–123 kyr). *Quaternary Science Reviews*, 29(1-2):222–234, 2010.
- Mark AJ Curran, Tas D Van Ommen, and Vin Morgan. Seasonal characteristics of the major ions in the high-accumulation Dome Summit South ice core, Law Dome, Antarctica. *Annals of Glaciology*, 27:385–390, 1998.
- GB Dreyfus, F Parrenin, B Lemieux-Dudon, G Durand, Valérie Masson-Delmotte, Jean Jouzel, J-M Barnola, L Panno, R Spahni, A Tisserand, et al. Anomalous flow below 2700 m in the EPICA Dome C ice core detected using $\delta^{18}\text{O}$ of atmospheric oxygen measurements. *Climate of the Past*, 3(2):341–353, 2007.
- A. Dällenbach, T. Blunier, J. Flückiger, B. Stauffer, J. Chappellaz, and D. Raynaud. Changes in the atmospheric CH_4 gradient between Greenland and Antarctica during the last glacial and the transition to the Holocene. *Geophysical Research Letters*, 27(7):1005–1008, apr 2000. doi: 10.1029/1999gl010873.
- Samuel Epstein and Robert P Sharp. Oxygen-isotope variations in the Malaspina and Saskatchewan Glaciers. *The Journal of Geology*, 67(1):88–102, 1959.
- Thomas Extier, Amaelle Landais, Camille Bréant, Frédéric Prié, Lucie Bazin, Gabrielle Dreyfus, Didier M Roche, and Markus Leuenberger. On the use of $\delta^{18}\text{O}_{\text{atm}}$ for ice core dating. *Quaternary Science Reviews*, 185:244–257, 2018.
- D. Fleitmann, H. Cheng, S. Badertscher, R. L. Edwards, M. Mudelsee, O. M. Göktürk, A. Fankhauser, R. Pickering, C. C. Raible, A. Matter, J. Kramers, and O. Tüysüz. Timing and climatic impact of Greenland interstadials recorded in stalagmites from northern Turkey. *Geophysical Research Letters*, 36(19), oct 2009. doi: 10.1029/2009gl040050.

- Johannes Freitag, Sepp Kipfstuhl, and Thomas Laepple. Core-scale radioscopic imaging: a new method reveals density-calcium link in Antarctic firn. *J. Glaciol.*, 59(218):1009–1014, 2013.
- S. Fujita, F. Parrenin, M. Severi, H. Motoyama, and E. W. Wolff. Volcanic synchronization of Dome Fuji and dome c Antarctic deep ice cores over the past 216 kyr. *Climate of the Past*, 11(10):1395–1416, oct 2015. doi: 10.5194/cp-11-1395-2015.
- Shuji Fujita, Junichi Okuyama, Akira Hori, and Takeo Hondoh. Metamorphism of stratified firn at Dome Fuji, Antarctica: A mechanism for local insolation modulation of gas transport conditions during bubble close off. *Journal of Geophysical Research: Earth Surface*, 114(F3), 2009.
- Chaochao Gao, Alan Robock, Stephen Self, Jeffrey B Witter, JP Steffenson, Henrik Brink Clausen, Marie-Louise Siggaard-Andersen, Sigfus Johnsen, Paul A Mayewski, and Caspar Ammann. The 1452 or 1453 ad Kuwae eruption signal derived from multiple ice core records: Greatest volcanic sulfate event of the past 700 years. *Journal of Geophysical Research: Atmospheres*, 111(D12), 2006.
- V Gkinis, TJ Popp, T Blunier, M Bigler, S Schüpbach, E Kettner, and SJ Johnsen. Water isotopic ratios from a continuously melted ice core sample. *Atmospheric Measurement Techniques*, 4(11):2531, 2011.
- C Goujon, J-M Barnola, and C Ritz. Modeling the densification of polar firn including heat diffusion: Application to close-off characteristics and gas isotopic fractionation for Antarctica and greenland sites. *Journal of Geophysical Research: Atmospheres*, 108(D24), 2003.
- Claus U Hammer. Acidity of polar ice cores in relation to absolute dating, past volcanism, and radio-echoes. *Journal of Glaciology*, 25(93):359–372, 1980.
- CU Hammer, HB Clausen, Willi Dansgaard, N Gundestrup, SJ Johnsen, and Niels Reeh. Dating of greenland ice cores by flow models, isotopes, volcanic debris, and continental dust. *Journal of Glaciology*, 20(82):3–26, 1978.
- James D Hays, John Imbrie, Nicholas J Shackleton, et al. Variations in the earth’s orbit: pacemaker of the ice ages. *Science*, 194(4270):1121–1132, 1976.
- Michael M Herron and Chester C Langway. Firn densification: an empirical model. *Journal of Glaciology*, 25(93):373–385, 1980.

- Maria W Hörhold, Mary R Albert, and Johannes Freitag. The impact of accumulation rate on anisotropy and air permeability of polar firn at a high-accumulation site. *Journal of Glaciology*, 55(192):625–630, 2009.
- Manuel A Hutterli, Martin Schneebeli, Johannes Freitag, Josef Kipfstuhl, and Regine Röthlisberger. Impact of local insolation on snow metamorphism and ice core records. *Low Temperature Science*, 68(Supplement):223–232, 2009.
- John Imbrie, James D Hays, Douglas G Martinson, Andrew McIntyre, Alan C Mix, Joseph J Morley, Nicklas G Pisias, Warren L Prell, and Nicholas J Shackleton. The orbital theory of Pleistocene climate: support from a revised chronology of the marine d18o record. 1984.
- Jean Jouzel, Georg Hoffmann, Frédéric Parrenin, and Claire Waelbroeck. Atmospheric oxygen 18 and sea-level changes. *Quaternary Science Reviews*, 21(1-3):307–314, 2002.
- K Kawamura, T Nakazawa, S Aoki, Y Fujii, O Watanabe, and JP Severinghaus. Close resemblance between local summer insolation, O₂/N₂ and total air content from the Dome Fuji ice core, Antarctica. In *AGU Fall Meeting Abstracts*, 2004.
- Kenji Kawamura, Frédéric Parrenin, Lorraine Lisiecki, Ryu Uemura, Françoise Vimeux, Jeffrey P Severinghaus, Manuel A Hutterli, Takakiyo Nakazawa, Shuji Aoki, Jean Jouzel, et al. Northern hemisphere forcing of climatic cycles in Antarctica over the past 360,000 years. *Nature*, 448(7156):912, 2007.
- Kenji Kawamura, Jeffrey P Severinghaus, MR Albert, ZR Courville, MA Fahnestock, T Scambos, E Shields, and CA Shuman. Kinetic fractionation of gases by deep air convection in polar firn. *Atmospheric Chemistry and Physics*, 13(21):11141–11155, 2013.
- KF Kuiper, Alan Deino, FJ Hilgen, W Krijgsman, PR Renne, Wijbrans, and JR. Synchronizing rock clocks of earth history. *Science*, 320(5875):500–504, 2008.
- P. Köhler. Rapid changes in ice core gas records – part 1: On the accuracy of methane synchronisation of ice cores. *Climate of the Past Discussions*, 6(4):1453–1471, aug 2010. doi: 10.5194/cpd-6-1453-2010.
- Amaelle Landais, Jean-Marc Barnola, Kenji Kawamura, Nicolas Caillon, M Delmotte, T Van Ommen, G Dreyfus, Jean Jouzel, Valérie Masson-Delmotte, B Minster, et al. Firn-air $\delta^{15}\text{N}$ in modern polar sites and glacial–interglacial ice: a model-data mismatch during glacial periods in Antarctica? *Quaternary science reviews*, 25(1-2):49–62, 2006.

- Jacques Laskar, Philippe Robutel, Frédéric Joutel, Mickael Gastineau, ACM Correia, and Benjamin Levrard. A long-term numerical solution for the insolation quantities of the earth. *Astronomy & Astrophysics*, 428(1):261–285, 2004.
- Bénédicte Lemieux-Dudon, Eric Blayo, Jean-Robert Petit, Claire Waelbroeck, Anders Svensson, Catherine Ritz, Jean-Marc Barnola, Bianca Maria Narcisi, and Frédéric Parrenin. Consistent dating for Antarctic and Greenland ice cores. *Quaternary Science Reviews*, 29(1-2):8–20, 2010.
- Quentin Libois, Ghislain Picard, Laurent Arnaud, Samuel Morin, and Eric Brun. Modeling the impact of snow drift on the decameter-scale variability of snow properties on the Antarctic plateau. *Journal of Geophysical Research: Atmospheres*, 119(20):11–662, 2014.
- V Ya Lipenkov, D Raynaud, MF Loutre, and P Duval. On the potential of coupling air content and O₂/N₂ from trapped air for establishing an ice core chronology tuned on local insolation. *Quaternary Science Reviews*, 30(23-24):3280–3289, 2011.
- Darren F Mark, Paul R Renne, Ross C Dymock, Victoria C Smith, Justin I Simon, Leah E Morgan, Richard A Staff, Ben S Ellis, and Nicholas JG Pearce. High-precision 40Ar/39Ar dating of Pleistocene tuffs and temporal anchoring of the Matuyama-Brunhes boundary. *Quaternary Geochronology*, 39:1–23, 2017.
- Kenneth C McGwire, Kendrick C Taylor, John R Banta, and Joseph R McConnell. Identifying annual peaks in dielectric profiles with a selection curve. *Journal of Glaciology*, 57(204):763–769, 2011.
- DA Meese, AJ Gow, RB Alley, GA Zielinski, PM Grootes, M Ram, KC Taylor, Paul Andrew Mayewski, and JF Bolzan. The Greenland ice sheet project 2 depth-age scale: methods and results. *Journal of Geophysical Research: Oceans*, 102(C12):26411–26423, 1997.
- Milutin Milankovitch. *Kanon der Erdebestrahlung und seine Anwendung auf das Eiszeitenproblem*. Königlich Serbische Akademie, 1941.
- Biancamaria Narcisi, Jean Robert Petit, and Massimo Tiepolo. A volcanic marker (92 ka) for dating deep East Antarctic ice cores. *Quaternary Science Reviews*, 25(21-22):2682–2687, 2006.

- NGRIP Community Members. High-resolution record of northern hemisphere climate extending into the last interglacial period. *Nature*, 431(7005):147–151, sep 2004. doi: 10.1038/nature02805.
- S.W. Nielsen. Registrering af visuel stratigrafi i ngrip iskernen: Konstruktion, dataopsamling og analyse. Technical report, Master’s Thesis, University of Copenhagen, 2005.
- F Parrenin, G Dreyfus, G Durand, S Fujita, O Gagliardini, F Gillet, Jean Jouzel, K Kawamura, N Lhomme, Valérie Masson-Delmotte, et al. Ice flow modelling at EPICA Dome C and Dome Fuji, East Antarctica. *Climate of the Past Discussions*, 3(1):19–61, 2007a.
- F. Parrenin, J.-R. Petit, V. Masson-Delmotte, E. Wolff, I. Basile-Doelsch, J. Jouzel, V. Lipenkov, S. O. Rasmussen, J. Schwander, M. Severi, R. Udisti, D. Veres, and B. M. Vinther. Volcanic synchronisation between the EPICA dome c and Vostok ice cores (Antarctica) 0-145 kyr BP. *Climate of the Past*, 8(3):1031–1045, jun 2012. doi: 10.5194/cp-8-1031-2012.
- Frédéric Parrenin. La datation des archives glaciaires. In Jean-Claude Duplessy and Gilles Ramstein, editors, *Paléoclimatologie: Trouver, dater et interpréter les indices*, volume 1. EDP Sciences, 2014.
- Frédéric Parrenin, Jean Jouzel, Claire Waelbroeck, Catherine Ritz, and Jean-Marc Barnola. Dating the Vostok ice core by an inverse method. *Journal of Geophysical Research: Atmospheres*, 106(D23):31837–31851, 2001.
- Frederic Parrenin, Frederique Remy, Catherine Ritz, Martin J Siegert, and Jean Jouzel. New modeling of the Vostok ice flow line and implication for the glaciological chronology of the Vostok ice core. *Journal of Geophysical Research: Atmospheres*, 109(D20), 2004.
- Frederic Parrenin, J-M Barnola, J Beer, Thomas Blunier, E Castellano, J Chappellaz, G Dreyfus, Hubertus Fischer, S Fujita, Jean Jouzel, et al. The EDC3 chronology for the EPICA Dome C ice core. *Climate of the Past*, 3:485–497, 2007b.
- Frédéric Parrenin, Lucie Bazin, Emilie Capron, Amaelle Landais, Bénédicte Lemieux-Dudon, and Valérie Masson-Delmotte. Icechronol: a probabilistic model to compute a common and optimal chronology for several ice cores. *Geoscientific Model Development*, 8(5):1473–1492, 2015.
- Jean-Robert Petit, Jean Jouzel, Dominique Raynaud, Narcisse I Barkov, J-M Barnola, Isabelle Basile, Michael Bender, J Chappellaz, M Davis, G Delaygue, et al. Climate

- and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. *Nature*, 399(6735):429, 1999.
- Vasilii V Petrenko, Andrew M Smith, Edward J Brook, Dave Lowe, Katja Riedel, Gordon Brailsford, Quan Hua, Hinrich Schaefer, Niels Reeh, Ray F Weiss, et al. 14ch4 measurements in greenland ice: investigating last glacial termination ch4 sources. *Science*, 324(5926):506–508, 2009.
- Vasilii V Petrenko, Jeffrey P Severinghaus, Hinrich Schaefer, Andrew M Smith, Tanner Kuhl, Daniel Baggenstos, Quan Hua, Edward J Brook, Paul Rose, Robb Kulin, et al. Measurements of 14c in ancient ice from taylor glacier, antarctica constrain in situ cosmogenic 14ch4 and 14co production rates. *Geochimica et Cosmochimica Acta*, 177: 62–77, 2016.
- G Raisbeck, F Yiou, Jean Jouzel, and TF Stocker. Direct north-south synchronization of abrupt climate change record in ice cores using Beryllium 10. *Climate of the Past*, 3: 541–547, 2006.
- Sune Olander Rasmussen, Katrine Krogh Andersen, Marie-Louise Siggaard-Andersen, and Henrik B Clausen. Extracting the annual signal from Greenland ice-core chemistry and isotopic records. *Annals of Glaciology*, 35:131–135, 2002.
- Dominique Raynaud, Vladimir Lipenkov, Bénédicte Lemieux-Dudon, Paul Duval, Marie-France Loutre, and Nicolas Lhomme. The local insolation signature of air content in Antarctic ice. a new step toward an absolute dating of ice records. *Earth and Planetary Science Letters*, 261(3-4):337–349, 2007.
- Catherine Ritz, Vincent Rommelaere, and Christophe Dumas. Modeling the evolution of Antarctic ice sheet over the last 420,000 years: Implications for altitude changes in the Vostok region. *Journal of Geophysical Research: Atmospheres*, 106(D23):31943–31964, 2001.
- Regine Röthlisberger, Matthias Bigler, Manuel Hutterli, Stefan Sommer, Bernhard Stauffer, Hans G Junghans, and Dietmar Wagenbach. Technique for continuous high-resolution analysis of trace substances in firn and ice cores. *Environmental Science & Technology*, 34(2):338–342, 2000.
- Andrey N Salamatin, Vladimir Ya Lipenkov, Jean Marc Barnola, Akira Hori, Paul Duval, and Takeo Hondoh. Snow/firn densification in polar ice sheets. *Low Temperature Science*, 68(Supplement):195–222, 2009.

- M. Severi, S. Becagli, E. Castellano, A. Morganti, R. Traversi, R. Udisti, U. Ruth, H. Fischer, P. Huybrechts, E. Wolff, F. Parrenin, P. Kaufmann, F. Lambert, and J. P. Steffensen. Synchronisation of the EDML and EDC ice cores for the last 52 kyr by volcanic signature matching. *Climate of the Past*, 3(3):367–374, jul 2007. doi: 10.5194/cp-3-367-2007.
- M. Severi, R. Udisti, S. Becagli, B. Stenni, and R. Traversi. Volcanic synchronisation of the EPICA-DC and TALDICE ice cores for the last 42 kyr BP. *Climate of the Past*, 8(2):509–517, mar 2012. doi: 10.5194/cp-8-509-2012.
- Jeffrey P Severinghaus and Mark O Battle. Fractionation of gases in polar ice during bubble close-off: New constraints from firn air Ne, Kr and Xe observations. *Earth and Planetary Science Letters*, 244(1-2):474–500, 2006.
- Jeffrey P Severinghaus, Alexi Grachev, Boaz Luz, and Nicolas Caillon. A method for precise measurement of Argon 40/36 and Krypton/Argon ratios in trapped air in polar ice with applications to past firn thickness and abrupt climate change in Greenland and at Siple Dome, Antarctica. *Geochimica et Cosmochimica Acta*, 67(3):325–343, 2003.
- Nicholas J Shackleton. The 100,000-year ice-age cycle identified and found to lag temperature, carbon dioxide, and orbital eccentricity. *Science*, 289(5486):1897–1902, 2000.
- Michael Sigl, Tyler J Fudge, Mai Winstrup, Jihong Cole-Dai, David Ferris, Joseph R McConnell, Ken C Taylor, Kees C Welten, Thomas E Woodruff, Florian Adolphi, et al. The WAIS Divide deep ice core WD2014 chronology—part 2: Annual-layer counting (0–31 ka bp). *Climate of the Past*, 12(3):769–786, 2016.
- Quentin Simon, Nicolas Thouveny, Didier L Bourlès, Jean-Pierre Valet, Franck Bassinot, Lucie Ménabréaz, Valéry Guillou, Sandrine Choy, and Luc Beaufort. Authigenic $^{10}\text{Be}/^{9}\text{Be}$ ratio signatures of the cosmogenic nuclide production linked to geomagnetic dipole moment variation since the Brunhes/Matuyama boundary. *Journal of Geophysical Research: Solid Earth*, 121(11):7716–7741, 2016.
- Brad S Singer, Hervé Guillou, Brian R Jicha, Carlo Laj, Catherine Kissel, Brian L Beard, and Clark M Johnson. $^{40}\text{Ar}/^{39}\text{Ar}$, K–Ar and ^{230}Th – ^{238}U dating of the Laschamp excursion: a radioisotopic tie-point for ice core and climate chronologies. *Earth and Planetary Science Letters*, 286(1-2):80–88, 2009.
- Claire L Smith, Ian J Fairchild, Christoph Spötl, Silvia Frisia, Andrea Borsato, Steven G Moreton, and Peter M Wynn. Chronology building using objective identification of

- annual signals in trace element profiles of stalagmites. *Quaternary Geochronology*, 4(1): 11–21, 2009.
- Todd Sowers, Michael Bender, Dominique Raynaud, and Ye S Korotkevich. $\delta^{15}\text{N}$ of n_2 in air trapped in polar ice: A tracer of gas transport in the firn and a possible constraint on ice age-gas age differences. *Journal of Geophysical Research: Atmospheres*, 97(D14): 15683–15697, 1992.
- Michael Storey, Richard G Roberts, and Mokhtar Saidin. Astronomically calibrated $^{40}\text{Ar}/^{39}\text{Ar}$ age for the Toba supereruption and global synchronization of late Quaternary records. *Proceedings of the National Academy of Sciences*, 109(46):18684–18688, 2012.
- Anders Svensson, Katrine Krogh Andersen, Matthias Bigler, Henrik Brink Clausen, Dorte Dahl-Jensen, SM Davies, Sigfus Johann Johnsen, Raimund Muscheler, F Parrenin, Sune Olander Rasmussen, et al. A 60 000 year Greenland stratigraphic ice core chronology. *Climate of the Past*, 4(1):47–57, 2008.
- Anders Svensson, Matthias Bigler, Thomas Blunier, Henrik Brink Clausen, Dorte Dahl-Jensen, Hubertus Fischer, S Fujita, K Goto-Azuma, SJ Johnsen, K Kawamura, et al. Direct linking of Greenland and Antarctic ice cores at the Toba eruption (74 ka bp). *Climate of the Past*, 9(2):749–766, 2013.
- Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*, volume 89. siam, 2005.
- Daniel Veres, Lucie Bazin, Amaëlle Landais, H Toyé Mahamadou Kele, Bénédicte Lemieux-Dudon, Frédéric Parrenin, Patricia Martinerie, Eric Blayo, Thomas Blunier, Emilie Capron, et al. The Antarctic ice core chronology (AICC 2012): an optimized multi-parameter and multi-site dating approach for the last 120 thousand years. *Climate of the Past*, 9(4):1733–1748, 2013.
- C Waelbroeck, J Jouzel, L Labeyrie, C Lorius, M Labracherie, M Stiévenard, and NI Barkov. A comparison of the Vostok ice deuterium record and series from Southern Ocean core MD 88-770 over the last two glacial-interglacial cycles. *Climate Dynamics*, 12(2):113–123, 1995.
- Yong-Jin Wang, Hai Cheng, R Lawrence Edwards, ZS An, JY Wu, C-C Shen, and Jeffrey A Dorale. A high-resolution absolute-dated late Pleistocene monsoon record from Hulu Cave, China. *science*, 294(5550):2345–2348, 2001.

Chapter 2

Paleoclimate chronologies and Bayesian Methods

An introduction to stochastic Bayesian methods: Markov Chain Monte Carlo and dating paleoclimate archives

Institut des Géosciences de l'Environnement

Jai Chowdhry Beeman

Dating a paleoclimate archive is the first, and arguably most essential step in understanding the information that it contains. Paleoclimate chronologies are never exact, and the chronological information for a given record is always limited and uncertain to some degree. Constructing a chronology for a paleoclimate record can be thought of as an exercise in assessing our state of knowledge about the timing of the record—an ideal chronology should accurately convey what we know about the timing of given record, and honestly assess how certain we are about this knowledge.

As briefly mentioned in the introduction, Bayes' theorem provides a mathematical framework precisely designed for assessing states of knowledge [Tarantola, 2005]. This chapter, designed to provide readers with some background information about inverse methods, presents Bayes' theorem and a small selection of methods for solving Bayesian inverse problems. The description of methods, far from complete, is intended to provide the reader with greater depth with respect to the algorithms we use to treat the inverse problems developed in later chapters. For a pedagogic, in-depth text about inverse problems, we refer the reader to Tarantola [2005], which inspires much of the content of this chapter. Finally, we introduce the IceChrono/PaleoChrono [Parrenin et al., 2015] dating model, which is central to the bulk of this thesis.

The inverse problem and Bayes' theorem

Inverse problems can be thought of as a framework for understanding a physical system that can be described numerically. In the broadest sense, the physical system we study when dating a paleoclimate archive is the Earth. For example, an ice core's physical properties are determined most importantly by the climate system and the carbon cycle; the dynamics and mass balance of ice sheets; the properties and position of the underlying

crust; and volcanic activity. We can use ice cores to better understand any of these attributes of the Earth system. In order to do so, we measure the physical and chemical properties of the core, obtaining observations that we hope to link to periods in the Earth's past.

With this information, we can advance and test our state of knowledge about the Earth system in the past. One way to do so is to develop numerical models. Tarantola [2005] separates modeling into three steps:

- **Parameterization:** Defining a set of model parameters whose values describe the physical behavior of the system.
- **Forward modeling:** Proposing a set of physical laws allowing us to estimate the values of observational parameters from the model parameters.
- **Inverse modeling:** Using real observations to estimate the values of the model parameters, which are not observable.

In practice, these steps often occur simultaneously. It is worth it to bear in mind that inverse modeling is often done without a formalized mathematical framework: that is, appropriate values of the model parameters can often be efficiently deduced with some knowledge of the behavior of the forward model at hand, the physically possible values for the parameters, and some observations.

On the other hand, the exercise of determining model parameters quickly becomes more difficult as the dimension (i.e. the number of parameters) of the model becomes larger. And often, we require greater accuracy in the calculation of the model parameter values: ideally, we would like to obtain a distribution of all sets of acceptable values and their respective probabilities, rather than just one set of acceptable values. Bayes' theorem permits calculating such a probability distribution.

The state of knowledge with respect to a system under study consists of two parts: the observations (the measurements from the paleoclimate record, in our example), and the model of how these records are physically expressed (in our example, this would be a sedimentation or growth model for the record). In most cases, we cannot make exact measurements of the observations, but rather express them with some uncertainty. This uncertainty is best expressed as a probability distribution of the observed parameters, which we call ρ_{data} . Similarly, we know that the model parameters should fall within a physically (and sometimes computationally) acceptable range. We can thus use this range to define a probability distribution that expresses what we know about the model

parameters; which in Bayesian statistics is referred to as a prior or model prior distribution (here ρ_{model} or ρ_{prior}).

If our model can be expressed as a relationship $G(X)$, where X is a set of model parameters, we can then write Bayes' theorem as:

$$\pi_{model|data} = k \frac{\rho_{model}(X) \rho_{data}(G(X))}{\mu(X)} \quad (1)$$

The term on the left side of equation 1 denotes the probability of a model given the observations $\pi_{model|data}$. This term, referred to in Bayesian terminology as the *posterior probability*, is equal to the conjunction of two additional probability densities. The first is the prior probability of the set of model parameters X , $\rho_{model}(X)$. The second is the probability of the observations given the model $G(X)$, $\rho_{data}(G(X))$.¹ k is a normalizing constant, which in theory ensures that the values of $\pi_{model|data}$ integrate to 1. Finally, μ is a probability density expressing the homogeneous state of information. In this thesis, we generally take the homogeneous state of information to be a uniform distribution across the model-data space (though this is not necessarily always the case), and so the expression simplifies to:

$$\pi_{model|data} = k \rho_{model}(X) \rho_{data}(G(X)) \quad (3)$$

It is difficult to know ρ_{data} exactly, since we would need to integrate over all possible model parameter sets to correctly scale the distribution. However, we can define a likelihood function L for the data: that is, a measure how good a model is at predicting the observed data. This likelihood function is by definition proportional to ρ_{data} . Reassigning the scaling to the constant k , we can rewrite equation 3 as:

$$\pi_{model|data} = k \rho_{model}(X) L_{data}(G(X)) \quad (4)$$

In the particular case of the IceChrono/PaleoChrono model, we assume the model prior and probability distribution of the observations to be Gaussian² in shape. Thus, we

¹Note the parallel to a much more common way to write Bayes' theorem:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \quad (2)$$

where P denotes a probability density, and B and A are probabilistic events.

²What is Gaussian? By Gaussian, we refer to probabilities that follow a normal distribution, or

$$f(x|\mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp -\frac{(x - \mu)^2}{2\sigma^2} \quad (5)$$

This distribution for variable or variable vector x is centered on its mean μ , and its width is defined by its standard deviation σ (or equivalently its variance σ^2). Because of the abundance of natural observations

can define a joint likelihood function for the prior and observations, and simply write:

$$\pi_{model|data} = k L(G(X)) \quad (6)$$

It is worth bearing in mind that the simplifications in Equations 4 and 6 are dependent on the characteristics of the problem at hand, and are not general expressions of Bayes' theorem.

IceChrono 1 [Parrenin et al., 2015]

Up until now, we have discussed Bayes' theorem mostly theoretically, with the prior and posterior probability distributions, for example, remaining rather abstract concepts. Now that these concepts have been defined, we can investigate what they look like in practice, in the development of a real inverse problem.

Several chapters of this thesis are based on the development of the PaleoChrono dating model. In its first iteration, PaleoChrono was known as IceChrono 1 [Parrenin et al., 2015], as it was first developed for use on ice cores, and subsequently adapted to be generally applicable to paleoclimate archives.³ Here, we provide a brief description of the original IceChrono model. The adaptations made for paleoclimate archives other than ice cores are rather simple : for example, we treat accumulation or sedimentation, but not thinning or air lock-in depth, which are unique to the ice core archive (though lock-in depth could be applied to the magnetic properties of sediment cores).

IceChrono treats the forward model of ice core dating much as it is discussed in the introduction. Three variables are used in the forward model—the accumulation or sedimentation rate α , the air lock-in depth LID and the thinning function τ . The composite vector of these three variables, X , corresponds to the X in the previous sections that denoted a model parameter vector.

Accumulation, or a , is expressed in $\text{m}\cdot\text{yr}^{-1}$ ice equivalent; Lock-in Depth (LID), or LID , is expressed in m ; and the ice thinning function, or τ , is unitless.

These three glaciological variables allow us to build an ice core chronology for each following the normal distribution, describing a random variable by its mean and standard deviation has become basic scientific practice. However, not all observable phenomena follow normal or even approximately normal distributions.

³Two additional Bayesian dating models, though not used in this thesis, are highly relevant to the topics discussed here. These are the Stratcounter annual layer counting model Winstrup et al. [2012] developed for ice cores, and the HMM-Match synchronization model Lin et al. [2014], developed for oceanic sediment cores.

core k at depth z . First, we calculate the ice age χ_k , by integrating the total accumulated and thinned ice layers above depth z :

$$\chi_k(z_k) = \int_0^{z_k} \frac{D_k(z'_k)}{a_k(z'_k)\tau_k(z'_k)} dz'_k \quad (7)$$

where D_k is the relative density with respect to pure ice at a given depth in the core. Then, we calculate the air age:

$$\Psi_k(z_k) = \chi_k(z_k - \Delta d(z_k)), \quad (8)$$

where $\Delta d(z_k)$, or Δ depth, is the depth difference between the air trapped at a given depth and the ice with the same age, which is located shallower in the core.

$\Delta d(z_k)$ is derived from Lock-in Depth LID using the following relationships:

$$\int_{z_k - \Delta d_k(z_k)}^{z_k} \frac{D(z'_k)}{\tau(z'_k)} \approx \int_0^{LID_k(z_k) D_k^{firm}(z_k)} \frac{1}{\tau_k(z_k^{ie})} dz_k^{ie}, \quad (9)$$

$$z_k^{ie} = \int_0^{z_k} D_k(z'_k) dz'_k \quad (10)$$

where $D_k^{firm}(z_k)$ is the firn density when the ice at depth z_k was at lock-in depth, and z_k^{ie} is the ice equivalent depth, i.e. the depth if the firn density was adjusted to the density of ice.

These integrals can be imagined in the following way: after lock-in, the firn above the air trapped at Lock-In Depth densifies into ice. The new depth difference between ice and air of equivalent age is the Lock-In Depth in Ice Equivalent, or $LIDIE$. The ice matrix is then thinned (described by the thinning function), to arrive at Δ depth. Δ depth can then be used, along with the ice age scale, to find the ice depth with equivalent age to a given air depth. The age difference between air and the ice at the same depth is referred to as Δ age.

The set of three equations described above corresponds to the model G in the Bayesian formulation. It is intuitive that when we proceed to approximate the posterior, at each step we will propose values for X , and then evaluate $G(X)$ —what we will often call the *proposed chronology*.

In the IceChrono 1 model, Gaussian (or normal) prior probability density functions are defined for each of the three glaciological variables, corresponding to ρ_{model} . The first component of each of these prior probability density functions is a mean scenario, calculated according to the best available glaciological knowledge at the site. Accumulation,

for example, can be estimated from isotopic proxies for a given depth. Standard deviations are given for each depth, and a covariance matrix is defined, completing the prior distributions.

Because the original IceChrono model was intended to be solved via a gradient-based least-squares optimization, the prior variables are transformed from Jeffreys variables to Cartesian variables (centered on zero) by treating any changes with a correction function. This is done by working with logarithmic correction functions. For accumulation, the correction function is expressed as :

$$c_k^a = \ln(a_k/a_k^b) \quad (11)$$

Where a_k is the proposed accumulation scenario, and a_k^b is the mean scenario of the prior. The correction functions for LID and τ are expressed, similarly, as

$$c_k^{LID} = \ln(LID_k/LID_k^b) \quad (12)$$

and

$$c_k^\tau = \ln(\tau_k/\tau_k^b). \quad (13)$$

IceChrono 1 allows for eight types of observations, which complement the prior probability distributions to constrain chronologies. These are :

- **Ice dated horizons** : for example, points dated using matches to external archives, magnetic reversals, or known volcanic events.
- **Ice dated intervals** : for example, intervals for which layers are counted, or climatic events for which the duration is explicitly known.
- **Ice-ice stratigraphic links** : Tie points obtained from synchronization in the ice matrix (volcanic synchronization, for example)
- **Air dated horizons** : points dated, again, using matches to external archives (i.e. orbital tuning)
- **Air dated intervals** : events in the air matrix for which the duration is explicitly known.
- **Air-air stratigraphic links** : Tie points obtained from synchronization in the air matrix (CH_4 synchronization, for example)

- **Air-ice / Ice-air stratigraphic links:** these might be obtained by synchronizing rapid changes in climate proxies, for example.
- **Δ depth observations:** Obtained, for example, using $\delta^{15}\text{N}$

In IceChrono 1, it is assumed that each of these discrete observations is independent, and that the uncertainties with respect to each observation are Gaussian. We can then write residual terms with respect to each of the glaciological parameters and each class of observations. These residuals, eventually, construct the probability of the observations ρ_{data} . In order to do so, the sum of the observation residual terms for core k , J_k , is written as:

$$J_k = R^T P^{-1} R \quad (14)$$

$$R = \frac{m_i - o_i}{\sigma_i} \quad (15)$$

Where m_i are the model approximations of the observations, o_i are the observations, and σ_i are their standard deviations. Note that the residual vector R is transformed by a correlation matrix P . Because the variables and observations are normalized in IceChrono 1, correlation matrices are generally used instead of covariance matrices.

The summed residuals for each of the prior terms are written as:

$$J_k^a = (R_k^a)^T (P_k^a)^{-1} (R_k^a)^T \quad (16)$$

$$J_k^{LID} = (R_k^{LID})^T (P_k^{LID})^{-1} (R_k^{LID})^T \quad (17)$$

$$J_k^\tau = (R_k^\tau)^T (P_k^\tau)^{-1} (R_k^\tau)^T \quad (18)$$

where R_k^τ , R_k^{LID} , and R_k^a are the residual vectors, and P_k^a , P_k^{LID} and P_k^τ are the correlation matrices of the three prior variables. The residual vectors are written as:

$$R_k^{LID} = \frac{c_k^{LID}}{\sigma_k^{LID}} \quad (19)$$

$$R_k^a = \frac{c_k^a}{\sigma_k^a} \quad (20)$$

$$R_k^\tau = \frac{c_k^\tau}{\sigma_k^\tau} \quad (21)$$

where σ_k^{LID} , σ_k^a and σ_k^τ are the standard deviations assigned to each of the glaciological variables.

Then, we can write a cost function J

$$J = \sum J_k + \sum J_{k,m} \quad (22)$$

which is composed of each of the summed residual terms J_k related to each individual core k included in a dating scenario:

$$J_k = J_k^a + J_k^{LID} + J_k^\tau + J_k^{ih} + J_k^{ah} + J_k^{ii} + J_k^{ai} + J_k^{\Delta d} \quad (23)$$

where J_k^a , J_k^{LID} , and J_k^τ are the residual terms related to each of the (prior) glaciological variables; and J_k^{ih} , J_k^{ah} , J_k^{ii} , J_k^{ai} and $J_k^{\Delta d}$ are the residual terms related to ice horizons, air horizons, ice intervals with known duration, air intervals with known duration, and Δ depth observations, respectively. Similarly, for each pair of cores k and m , a summed residual term $J_{k,m}$ is defined:

$$J_{k,m} = J_{k,m}^{ii} + J_{k,m}^{aa} + J_{k,m}^{ia} + J_{k,m}^{ai} \quad (24)$$

where $J_{k,m}^{ii}$, $J_{k,m}^{aa}$, $J_{k,m}^{ia}$, and $J_{k,m}^{ai}$ are ice-ice, air-air, ice-air and air-ice stratigraphic links, respectively.

Each of the individual PDF's used to define the residual terms is expressed as an independent multivariate gaussian, meaning that we can express the posterior likelihood can finally be written as

$$L = \exp\left(-\frac{1}{2}J\right) \quad (25)$$

In IceChrono 1, the structure of the model is such that the shape of the posterior probability distribution defined by this likelihood function is also near-Gaussian. Thus, the Levenberg-Marquardt algorithm [Levenberg, 1944, Marquardt, 1963] can be used to converge toward the minimum of the cost function. This is done using the Jacobian of the residual vector. The model Jacobian is also used to estimate the approximate error covariance matrix of the model, thus providing a Gaussian estimate of the posterior probability distribution as well.

This last attribute changes with the development of IceChrono/PaleoChrono discussed in this thesis. When a continuous synchronization is introduced, the associated cost function term can have multiple local minima, which are often surrounded by high cost function

values thus making the posterior probability density multimodal and jagged—i.e. highly non-Gaussian. Thus, the optimization framework of IceChrono/PaleoChrono is revised for the continuous synchronization applications.

Solving inverse problems: evaluating the posterior distribution

Once the inverse problem is set up (in the case of IceChrono, once the likelihood function is defined) we can proceed to the solution step: evaluating the posterior distribution. Several methods exist to evaluate posterior distributions. No method is a black box – instead, the choice of the method will depend on the shape, size and characteristics of the distribution at hand. In some special cases, the posterior distribution can be calculated analytically (an example in the field of dating ice cores is the Datice model [Lemieux-Dudon et al., 2010]). When the posterior distribution is less well-defined, but we can be sure that it contains a well-defined probability maximum, and that probability generally decreases as we move away from the maximum, gradient-based methods can be applied. In these methods, the gradient of the posterior probability density is usually estimated for a small number of sets of parameters different by a small amount. Then, a step is made in the direction of the largest gradient, and the gradient is recalculated. This continues until the gradient is estimated to be close enough to zero that we can be confident that we have found a probability maximum. An example of a gradient-based algorithm applied to dating paleoclimate archives is the non-linear least squares algorithm used in the original IceChrono model [Parrenin et al., 2015].

It is not always possible to generally characterize posterior probability distributions to the degree necessary to apply an analytical or gradient-based solution method. Any of the three components of Bayes’ theorem: the model prior, the probability distributions of the observations, or even the model itself, can introduce complexity into the posterior, in the form of multiple modes of high probability, for example.

Gradient-based methods often have difficulty exploring multiple modes, because the gradient of the probability density is zero at a local maximum. Similarly, analytical solutions often require assumptions about the shape of the posterior probability distribution that are not generally applicable. Stochastic methods, on the other hand, are capable in theory of better exploring probability densities without necessarily relying on assumptions about the shape of the density.

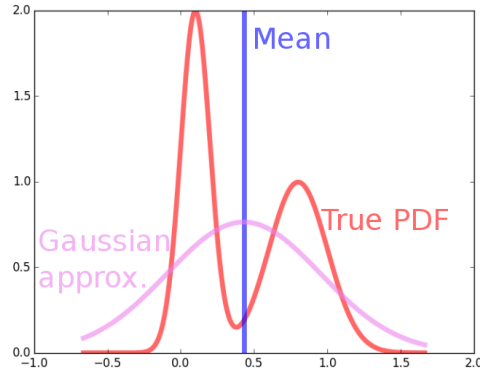


Figure 1: A schematic illustrating the oversimplification of a multimodal probability density using a Gaussian (normal) approximation. The true probability density function has two modes of high probability, each relatively narrow. The Gaussian approximation places the mean in between the two modes of the true probability density, in an anti-mode of the true probability distribution, and the scale parameter, or standard deviation, overestimates the width of the probability distribution because the two modes of the true distribution are separate.

Monte Carlo and Metropolis-Hastings Algorithms

The class of stochastic algorithms known as Monte Carlo methods [Metropolis and Ulam, 1949] is useful for solving Bayesian inverse problems, particularly when the posterior distribution associated with the inverse problem is complex and difficult to approximate using analytical or gradient-based methods.

At the most basic level, Monte Carlo methods use stochastic generation to approximate some numerical quantity. A commonly used example (see Tarantola [2005], among many others) is that of estimating the area of a complex shape. Suppose that we have drawn a square, and can calculate its surface area by measuring one of the sides. Inside this square lies another shape, with irregular boundaries, whose surface area is difficult to calculate analytically. We would like to estimate the surface area of the irregular shape. To do so using a stochastic, Monte Carlo-like method, we could throw a large number of darts, at random, onto the two shapes. At the end, the ratio of the number of darts in the irregular shape to the total number of darts in the square should approximate the ratio of the unknown area of the shape to the known area of the square. This allows us to estimate the area of the shape.

Another key principle of stochastic algorithms can be inferred from this simple example:

the law of large numbers. When we have thrown just one dart at our shape, the ratio of the areas of the two shapes can be either zero or one, clearly a very poor approximation. When ten darts have been thrown, some will land inside the shape, and some will land outside. The ratio of darts will become a better approximation of the ratio of surface areas. By the time we have thrown a reasonably large amount of darts, the two ratios should be practically indistinguishable. We might then say that the stochastic simulation has converged.

Probability distributions, including Bayesian posteriors, can be sampled from using essentially the same principle. However, this strategy would run into two problems. The space of acceptable model parameters is often much, much smaller than the space of possible values for each parameter. This problem of scale increases with each added parameter dimension—it might be possible to estimate the probability distribution of one unknown parameter using a purely random Monte Carlo simulation, but it quickly becomes more difficult to sample the distribution for larger-scale inverse problems, scaling with the dimensionality, or number of parameters, of the problem. In dating an ice core, for example, we typically deal with parameter spaces with hundreds of dimensions. Thus, using a purely randomized search would likely be exceedingly difficult and computationally expensive. In addition, we do not know the shape of the posterior, and it is difficult to evaluate whether samples fall inside or do not. We could do so by iteratively estimating the constant by numerically integrating the likelihood of the samples up to each step, but this would be computationally very costly as well.

A set of algorithms exists that are designed to solve precisely this problem, while retaining a stochastic character that allows them to explore complex probability distributions and not rely on gradients or analytical mathematics. These are known as Markov Chain Monte Carlo algorithms. Markov chains are stochastic sequences in which the state attained by each element depends only on the state attained by the previous element. In a Monte Carlo simulation, allowing the exploration of parameters to operate as a Markov Chain translates to allowing each model proposal to depend on the last model in the simulation. In practice, this implies that MCMC algorithms can be oriented to explore a desired region, rather than the “grope-in-the-dark” strategy of simple Monte Carlo. Markov Chain Monte Carlo algorithms are (and should be) designed to ensure convergence to the probability distribution we wish to explore. In general, this means that as the simulation converges, the number of samples we select in any region is proportional to the probability mass of this region.

Developed by Nicholas Metropolis and W.K. Hastings, the Metropolis-Hastings algorithm [Metropolis et al., 1953, Hastings, 1970] and its variants are perhaps the most widely used Markov Chain Monte Carlo algorithms; we indeed make use of a parallelized variant of the Metropolis-Hastings algorithm in this thesis. For some background, we provide a cursory explanation of the algorithm; for a more rigorous treatment the reader is referred to Tarantola [2005], among many others.

The Metropolis-Hastings algorithm [Metropolis et al., 1953, Hastings, 1970] stochastically generates models, here \mathbf{X}^n , from a proposal distribution $q(\mathbf{X}^n|\mathbf{X}^{n-1})$. Importantly, this proposal distribution need not be related to the prior or posterior distribution in the Bayesian inversion.

The proposed models are accepted or rejected with probability:

$$P_{X^n \rightarrow X_{n+1}} = \min \left(\frac{L(\mathbf{X}^n)q(\mathbf{X}^n|\mathbf{X}^{n-1})}{L(\mathbf{X}^{n-1})q(\mathbf{X}^{n-1}|\mathbf{X}^n)}, 1 \right) \quad (26)$$

$$L(\mathbf{X}^n) = \rho_{data}(\mathbf{X}^n)\rho_{model}(\mathbf{X}^n) \quad (27)$$

where ρ_{data} is the prior probability distribution of the data. If a gaussian prior probability distribution is taken, for example, we can rewrite:

$$L(\mathbf{X}^n) = k \exp(-J_n) \quad (28)$$

where k is a normalizing constant which cancels out in the computation of $P_{X^n \rightarrow X_{n+1}}$. Furthermore, when proposal distribution q is symmetric, equation 26 simplifies to:

$$P_{X^n \rightarrow X_{n+1}} = \min \left(\frac{L(\mathbf{X}^n)}{L(\mathbf{X}^{n-1})}, 1 \right) \quad (29)$$

This acceptance probability is the classic Metropolis algorithm from [Metropolis et al., 1953] – the Metropolis-Hastings algorithm is a more general version. The samples produced via the Metropolis-Hastings selection criteria can be shown to converge to the posterior probability distribution

$$\pi_{model}(\mathbf{X}^n) = kL(\mathbf{X}^n) \quad (30)$$

where k is a normalizing constant.

Note that we do not need to know the distribution $\pi_{model}(\mathbf{X}^n)$ to sample from it – in the kinds of inverse problems that Metropolis-Hastings algorithms are used for, we never do! We only need to be able to construct $L(\mathbf{X}^n)$, a much easier task.

Why does the Metropolis-Hastings algorithm work?

Introducing a few mathematical concepts can help us understand the workings of the Metropolis-Hastings algorithm. These concepts generally have parallels for more advanced stochastic algorithms as well.

In the darts example, recall that we required a large number of dart throws to truly approximate the area of the shape. We could do so because we threw the darts uniformly across the square, allowing us to access all areas of the square. Our throws were effectively generated using a uniform distribution.

We could have calculated the area of the shape using another known distribution, as long as we accounted for the probability of each given throw when integrating the area of the shape. But if, for example, we always aimed at the position of the last dart, it would be less intuitive to guarantee that we had calculated the correct area of the shape, because it would be difficult to assess whether we could eventually explore the whole scale.

In statistical terms, our throws were independent and identically distributed (IID). The throws were independent of each other, because no throw depended on the location of any other dart, and identically distributed, because each dart had an identical probability of landing at a given location on the square. Monte Carlo algorithms with IID proposals can, eventually, access the entire space of the probability distribution they are attempting to sample.

Markov Chains, on the other hand, break this rule. Each step of a Markov Chain is allowed to depend on the last step, meaning that these steps are not independent of each other (nor are they necessarily identically distributed). So we need to be sure that a Markov Chain will eventually reach the entire probability space, in spite of not being IID.

Formally, we can do so by showing that a Markov Chain is ergodic, meaning that for a given space, the chain can, given an infinite time to evolve, reach any point in the space.

The ergodic theorem for Markov Chains

Here, we present a simple ergodic theorem, for Markov Chains operating in discrete time and discrete space. For a Markov chain looking to approximate a function f with target distribution π , this theorem can be written:

Theorem 1. *Ergodic Theorem for discrete time Markov Chains*

1. *If $[X_0, X_1, \dots, X_i, \dots, X_n]$ is an irreducible (time-homogeneous), discrete space Markov Chain with stationary distribution π , then*

$$\frac{1}{n} \sum_{i=1}^n f(x_i) \xrightarrow{n \rightarrow \infty} Ef(x) \text{ where } x \sim \pi$$

for any bounded $f: \mathfrak{X} \rightarrow \mathbb{R}$.

2. If, further, $[X_0, X_1, \dots, X_i, \dots, X_n]$ is aperiodic, then

$$P(X_n = x | X_0 = x_0) \rightarrow \pi(x)$$

The ergodic theorem outlines the conditions that can guarantee that a Markov Chain Monte Carlo algorithm can converge to a given distribution. Part 1 of the theorem describes the conditions that need to be fulfilled so that the mean value of a function f (any bounded function with state space \mathfrak{X} in the set \mathbb{R} of all real numbers) of the samples x_i converges to the expected value of the function, with x following a distribution π (this distribution can indeed be the posterior, in a Bayesian application). In the second part, conditions are referenced so that the probability distribution of the samples of the chain P should converge to the probability distribution π .

Three conditions are referenced:

1. The distribution π which we wish to explore is the **stationary** distribution of the chain we implement.
2. The chain is **irreducible** (time-homogeneous).
3. The chain is **aperiodic**.

A Markov Chain \mathbf{X}_i is called **time-homogeneous** if the probability of the chain transitioning from position a to position b , $P(X_{i+1} = b | X_i = a) = T_{ab}$ for all i and all $a, b \in \mathfrak{X}$ for some matrix T which does not depend on time.

The matrix T is often referred to as the transition matrix of a Markov Chain. This is a stochastic matrix, whose entries are non-negative and whose rows, representing all the probabilities of transition from a given state, sum to 1. In the case of continuous-time or continuous-space Markov Chains, we refer to a stochastic transition kernel (or function) rather than a matrix.

Simply, we require that all transitions from one state to another are allowed to occur with the same probability, regardless of the timestep at which they occur. We know that this is the case for the classic Metropolis-Hastings algorithm: the Metropolis-Hastings




Figure 2: <https://xkcd.com/1724/>. MCMC algorithms may seem to operate like dark magic to the uninitiated, but they can often be shown to converge using relatively intuitive principles.

ratio 26, which determines the probability of a transition from one state to another, has no time dependency.

A probability mass function π on \mathfrak{X} is a **stationary** or invariant distribution with respect to transition matrix T if $\pi T = \pi$. This can be written as:

$$\sum_{a \in \mathfrak{X}} \pi_a T_{ab} = \pi_b \forall b \in \mathfrak{X} \quad (31)$$

We will soon see that a probability mass function π is a stationary distribution of chain \mathbf{X}_i if π satisfies a condition called detailed balance with respect to the transition matrix T_{ab} of the chain.

A Markov Chain is **irreducible** if for all $a, b \in \mathfrak{X}$ there exists $t \geq 0$ such that $P(X_t = b | X_0 = a) > 0$. That is, for any states a and b in the state space \mathfrak{X} , if we start at state a , it is possible to reach state b .

Finally, a Markov Chain is **aperiodic** if for all $a \in \mathfrak{X}$ the greatest common denominator of the times t for which $P(X_t = a | X_0 = a) > 0$ is 1. That is, the chain should not, by design, revisit the same state with any periodicity.

The proof of this theorem is beyond the scope of this chapter, but is accessible in mathematics texts (i.e. Norris [1998], statslab.cam.ac.uk/~james/Markov/s110.pdf). However, we can discuss the intuition behind these three conditions for convergence. The condition of stationarity essentially states that we wish to sample from probability distribution π . The condition of irreducibility expresses that we should, eventually, be able to visit any state in the state space of this distribution. Finally, the condition of aperiodicity expresses that we should not revisit the same states repeatedly. Together, we should then be able to sample the distribution π correctly, exhaustively and without periodic bias—the law

of large numbers indicating that we should eventually be able to create an acceptable approximation of π .

Ergodicity and the Metropolis-Hastings Algorithm

A Markov Chain produced by the Metropolis-Hastings Algorithm is designed in order to approximate a target probability distribution. The chain, in other words, should be ergodic with the target distribution as the stationary distribution of the Markov Chain.

In order to show that this is indeed the case, it is useful to introduce the concept of *detailed balance*, or *reversibility*.

A probability mass function π on state space \mathfrak{X} is said to satisfy **detailed balance** with respect to transition matrix T if $\pi_a T_{ab} = \pi_b T_{ba}$ for all $a, b \in \mathfrak{X}$.

Remember that for a probability distribution π to be stationary with respect to transition matrix T , it should hold that $\sum_{a \in \mathfrak{X}} \pi_a T_{ab} = \pi_b$. Showing that π satisfies detailed balance with respect to T is equivalent to showing that π is stationary with respect to T . This is because with detailed balance, we have

$$\sum_a \pi_a T_{ab} = \sum_a \pi_b T_{ba}. \quad (32)$$

Since π_b does not depend on a , we can write

$$\sum_a \pi_a T_{ab} = \pi_b \sum_a T_{ba}; \quad (33)$$

since we know that the entries of T sum to 1, this is equivalent to

$$\sum_a \pi_a T_{ab} = \pi_b. \quad (34)$$

as above. In practice, Markov Chains with detailed balance are reversible: that is, if a Markov Chain satisfies detailed balance with respect to T , so does the chain with the same elements in reversed order.

We might think of detailed balance as an equation of conservation of probability mass: we should tend to conserve the correct proportions of probability at states a and b with transition matrix T .

We can easily show that the Metropolis algorithm (using a symmetric proposal distribution, as for equation 29), satisfies the conditions outlined in the ergodic theorem for its target distribution (the same is true for the general Metropolis-Hastings algorithm with some adaptations to the equations below).

Recall that in the Metropolis algorithm, we do not know the target distribution but rather some function proportional to the target distribution, which we denote as $\tilde{\pi}$ (in this thesis, this is generally the Likelihood function). We have also picked a symmetric proposal distribution q (which has an equivalent transition matrix Q). Equation 29 gives us the probability with which we accept proposals, that is $P_{X^n=a \rightarrow X^{n+1}=b} = \min\left(\frac{L(a)}{L(b)}, 1\right)$.

We check first if this algorithm satisfies the first condition, that π is the stationary distribution of the chain \mathbf{X} . We can do so by checking if π satisfies detailed balance with respect to the theoretical transition matrix T ; i.e. $\pi_a T_{ab} = \pi_b T_{ba}$.

In the Metropolis algorithm, two cases are possible. We might reject a transition, in which case $a = b$. Then, $\pi_a T_{ab} = \pi_a T_{aa} = \pi_b T_{ab}$, so the algorithm satisfies detailed balance.

In the second case, we accept the transition, that is $a \neq b$. In this case, the probability of the transition from b to a is equivalent to the conjunction of the probability of the proposal $q(a, b)$ and the probability with which the proposal was accepted, $\min\left(\frac{\tilde{\pi}(a)}{\tilde{\pi}(b)}, 1\right)$.

$$\pi_a T_{ab} = \pi(a)q(a, b) \cdot \min\left(\frac{\tilde{\pi}(a)}{\tilde{\pi}(b)}, 1\right). \quad (35)$$

Since the constant of proportionality between $\tilde{\pi}$ and π cancels, this is equivalent to

$$\pi_a T_{ab} = \pi(a)q(a, b) \cdot \min\left(\frac{\pi(a)}{\pi(b)}, 1\right). \quad (36)$$

We can multiply $\pi(a)$ inside the \min function to give

$$\pi_a T_{ab} = q(a, b)\min(\pi(a), \pi(b)) \quad (37)$$

Note that, replacing $\pi_a T_{ab}$ with $\pi_b T_{ba}$, the right hand side of 37 does not change. Thus, detailed balance holds, and the chain has stationary distribution π . Note that for the same can hold with a non-symmetric proposal if we accept or reject proposals with the probability as outlined in equation 26.

Irreducibility and aperiodicity are not difficult to show. We can usually assume that $\pi(x) > 0$ for all x in the state space \mathfrak{X} . Then, because transition probabilities depend only on π and q , for any irreducible Q , T is irreducible, and for any aperiodic Q , T is aperiodic. Thus, we know that the Markov Chain resulting from the Metropolis-Hastings algorithm should converge to its stationary distribution, and that this stationary distribution is indeed π .

Numerous stochastic algorithms have been designed to more efficiently explore probability distributions that are challenging because of shape or size. Indeed, we use some

of these algorithms in this thesis. These algorithms are often built using the same principles as those of the Metropolis-Hastings algorithm. It can often be shown that they should converge, in theory, to their target distribution using the principle of detailed balance, with equations not too different from those introduced in this section (although adaptive algorithms in particular violate the properties of Markov Chains to a greater degree, making the proof of convergence more involved). In any case, understanding why the Metropolis-Hastings algorithm can be shown to converge is a useful base for understanding the convergence properties of this large class of more developed stochastic algorithms.

Sampling from challenging target distributions: adaptive and parallelized MCMC algorithms

In theory, the Metropolis-Hastings algorithm and successive stochastic algorithms are capable of converging to any probability distribution, given a sufficient (infinite) amounts of time and samples. Of course, when running these algorithms, we work with finite amounts of time and samples. Necessarily, the estimates we obtain of the target distributions are approximate. The accuracy of the approximation depends on the number of samples, and on how efficient the algorithm is at exploring the probability distribution at hand.

Probability distributions can become much more difficult to explore if they are multimodal, and the proposal distribution (denoted $q(\mathbf{X}^n|\mathbf{X}^{n-1})$ above) is not well-adapted. To reach independently prominent *peaks* of high probability in a multimodal distribution, a Markov Chain must cross *valleys* of low probability. Since the probability of entering a valley from a peak is low with equation 26, it can take a very long time to cross to a secondary peak, considerably increasing the amount of simulations required to have a good estimate of the probability distribution.

This problem occurs generally when the proposal distribution is not well-adapted to the target distribution. Even in the case of a relatively simple, unimodal distribution, if the proposal distribution is biased toward sampling in a given direction, and the target distribution is oriented along another direction, convergence will be considerably slowed.

Dealing with large parameter spaces can also complicate using MCMC algorithms. Generally, the number of iterations required for a Metropolis-Hastings simulation to converge to the posterior distribution scales with N^2 for N parameters.

In this thesis, we deal with both large parameter spaces and multimodal, jagged distributions; in addition, we have little to no information beyond the prior about the scale

or orientation of these distributions.

In the exploratory phases of both the inverse problems we address, we tested both gradient-based and stochastic algorithms; the gradient-based algorithms converged to localized minima and the standard MH simulations tended to be slow to converge. In both cases, we ultimately attempted to pick MCMC algorithms that balance a good adaptation to the complexity of the posterior distribution and an efficient use of computational resources via parallelization. The algorithms that we applied to the two studies are introduced in the next section; for a broader overview Sharma [2017] reviews parallelized MCMC algorithms developed for astronomy. The selection of MCMC algorithms for challenging probability distributions is problem-dependent, and none of the algorithms we tested showed black-box applicability to our synchronization problem, for example. However, the development of advanced MCMC algorithms is a dynamic field, and algorithms will surely continue to be developed that will allow us to solve the inverse problems outlined here with greater efficiency and less implementation time.

Two selected parallelized MCMC algorithms

The Affine Invariant Ensemble Sampler [Goodman and Weare, 2010]

In chapter 3, we pose the detection of changes in trend in CO₂ and Antarctic Temperature during the last deglaciation as an inverse problem. The parameter space of this problem is not unusually large, requiring between 6 and 12 parameters. However, we solve the problem multiple times, with respect to several time series. To tackle the large computation time required for traditional MH sampling, we apply the ensemble sampler developed by Goodman and Weare [Goodman and Weare, 2010] (GW) as implemented in the python emcee library [Foreman-Mackey et al., 2013].

The GW sampler adapts the MH algorithm so that multiple model walkers can explore the probability distribution at once, making the algorithm parallelizable. Each of these walkers is a Markov Chain, but the proposal for each walker at a given step depends on the complementary ensemble of all of the other walkers. This gives the algorithm advantage of being affine invariant: that is, steps are adapted to the scale and directionality of the posterior distribution, as roughly approximated by the ensemble of walkers at any given step.

To propose updates to the walkers, we apply what GW refer to as a “stretch move”. Consider an ensemble of walkers \mathbf{X} and an individual walker \mathbf{X}_k^j in the ensemble at




Figure 3: The GW stretch move. Light dots are walkers in the complimentary ensemble. The update to X_j is proposed by stretching along the straight line between X_j and X_k . Figure and caption (adapted) from Goodman and Weare [2010].

proposal step j . Another walker \mathbf{X}_h^j is selected from the complementary ensemble $\mathbf{X}_{[k]}^j$, composed of all of the other walkers. Then, a proposal is made to update \mathbf{X}_k to W :

$$\mathbf{X}_k^j \rightarrow W = \mathbf{X}_h^j + Z \left(\mathbf{X}_k^j - \mathbf{X}_h^j \right) \quad (38)$$

GW suggest the following probability distribution to generate stochastic variable Z , which we apply:

$$g(Z) \propto \begin{cases} \frac{1}{Z} & \text{if } Z \in \left[\frac{1}{a}, a \right] \\ 0 & \text{otherwise} \end{cases} \quad (39)$$

where a is a user-defined constant. Given this probability distribution, proposals are accepted or rejected with acceptance probability:

$$P_{\mathbf{X}_k^j \rightarrow \mathbf{X}_k^{j+1}=W} = \min \left\{ 1, Z^{j-1} \frac{L(W)}{L(\mathbf{X}_k^j)} \right\}. \quad (40)$$

Note that the term Z^{j-1} is included so that detailed balance is still respected.

The Differential Evolution Markov Chain with snooker updater [ter Braak and Vrugt, 2008]

Though developed separately, Differential Evolution Markov Chain (DEMC, ter Braak and Vrugt [2008]) and [Goodman and Weare, 2010] share a structural commonality—in both algorithms, an ensemble of Markov Chains evolves together, making both algorithms easily parallelizable. In both algorithms, the proposal at each step is determined using the positions of the remaining $N - 1$ chains.




Figure 4: The parallel direction update in DE-MC, with each chain represented by a point. a) To update chain i in state \mathbf{x}_i , the proposed state \mathbf{x}^* is generated by adding the weighted difference of the two other chains $\gamma(\mathbf{x}_{R1} - \mathbf{x}_{R2})$ to \mathbf{x}_i , plus a random vector \mathbf{e} . b) The reverse jump, to show detailed balance. Figure and caption (adapted) from ter Braak and Vrugt [2008].




Figure 5: a) The parallel direction update in DE-MC, with each chain represented by a point, in this case with $\gamma = 1$. b) The case of an outlier chain. In these figures \mathbf{e} is not shown as it is small compared to the jumps. Figure and caption (adapted) from ter Braak and Vrugt [2008].




Figure 6: The DE-MC snooker move, which generates a proposal along the line through \mathbf{x}_i and the state \mathbf{z} of another chain. The proposal point \mathbf{x}^* is generated by randomly selecting two other chains (\mathbf{z}_{R1} and \mathbf{z}_{R2}), projecting them orthogonally onto the line between \mathbf{x}_i and \mathbf{z} , and adding a multiple of the difference between the projection points \mathbf{z}_{P1} and \mathbf{z}_{P2} to \mathbf{x}_i . Figure and caption text from ter Braak and Vrugt [2008].

ter Braak and Vrugt [2008] outline two proposal steps, described here. It is worth noting that several algorithms have been developed based on the steps proposed by ter Braak and Vrugt [2008]—see Laloy and Vrugt [2012] and Nelson et al. [2013], for example.

Differential Evolution Move

For an ensemble \mathbf{X}^j of N parallel Markov chains with current positions \mathbf{x}^j at timestep j , we seek to update walker \mathbf{x}_i . In order to do so, we randomly select two additional walkers, with current positions $\mathbf{x}_{\mathbf{R1}}$ and $\mathbf{x}_{\mathbf{R2}}$ from the complimentary set $\mathbf{X}_{\neq i}^j$. The proposal, denoted \mathbf{x}^* is calculated as:

$$\mathbf{x}^* = \mathbf{x}_i^j + \gamma(\mathbf{x}_{\mathbf{R1}} - \mathbf{x}_{\mathbf{R2}}) + \mathbf{e} \quad (41)$$

where γ is a scalar (user-defined), and \mathbf{e} is a random vector. In order to respect detailed balance, this vector should be drawn from a symmetric distribution; to make sure that our proposals do not veer too far from the core DE move the width of this symmetric distribution is ideally small compared to the width of the posterior.

We accept or reject each jump according to the Metropolis-Hastings ratio:

$$P_{\mathbf{x}_i^j \rightarrow \mathbf{x}_i^{j+1} = \mathbf{x}^*} = \min \left\{ 1, \frac{L(\mathbf{x}^*)}{L(\mathbf{x}_i^j)} \right\} \quad (42)$$

This move occurs on a line approximately parallel to $\mathbf{x}_{\mathbf{R1}}$ and $\mathbf{x}_{\mathbf{R2}}$ —making the assumption that the current positions of the ensemble of walkers at any iteration approximately reflects the orientation and scale of the posterior distribution. This assumption, of course, becomes more and more true as the simulation converges. Figure 4 illustrates this move.

ter Braak and Vrugt [2008] suggest two possible values of the parameter γ . Choosing $\gamma = 2.38 / \sqrt{2d}$, where d is the number of model parameters, provides an approximate acceptance probability of around 0.23 for Gaussian and Student distributions. An acceptance probability of 0.23 is commonly taken as an optimal value for MCMC, though it is worth noting that the posterior distributions we approximate in this thesis are quite far from being appropriately represented by Gaussian or Student functions.

ter Braak and Vrugt [2008] also suggest using $\gamma \approx 1$ for a certain proportion of the proposals. This value is taken with multimodality in mind – making the assumption that the distance between any two modes is approximately characteristic of the distribution as a whole (we postulate that this assumption should be more or less true for our CH_4 synchronization problem, given the recurrent nature of the climate cycles reflected in methane time series). Figure 5 illustrates this move.

Snooker Move

The last move developed by ter Braak and Vrugt [2008] is the DEMC snooker update, illustrated in figure 6. To update walker \mathbf{x}_i using this algorithm, three other chains, with current states \mathbf{z} , \mathbf{z}_{R1} and \mathbf{z}_{R2} are selected. The vectors \mathbf{z}_{R1} and \mathbf{z}_{R2} are projected orthogonally onto the difference vector $\mathbf{x}_i - \mathbf{z}$. Using the two orthogonally projected vectors \mathbf{z}_{P1} and \mathbf{z}_{P2} , the snooker proposal is given as:

$$\mathbf{x}^* = \mathbf{x}_i^j + \gamma(\mathbf{z}_{P1} - \mathbf{z}_{P2}) \quad (43)$$

Maintaining detailed balance with this move requires a modification to the basic Metropolis-Hastings acceptance ratio. In the case of the snooker move, the q terms do not cancel, since the lines $\mathbf{x}^* - \mathbf{z}$ and $\mathbf{x}_i - \mathbf{z}$ are different in magnitude. Since this is the only difference between the proposal distributions, the acceptance ratio is written as follows:

$$P_{\mathbf{x}_i^j \rightarrow \mathbf{x}_i^{j+1} = \mathbf{x}^*} = \min \left\{ 1, \frac{L(\mathbf{x}^*) \|\mathbf{x}^* - \mathbf{z}\|^{d-1}}{L(\mathbf{x}_i^j) \|\mathbf{x}_i^j - \mathbf{z}\|^{d-1}} \right\} \quad (44)$$

where d is the dimensionality of the model vector.

The orthogonal projection effectively projects the covariance between model parameters onto the proposal. In a way, we can think of this step as reducing the effective model dimensionality in the MCMC simulation: rather than waiting for each parameter to reach acceptable values proposed independently, we use the positions of other walkers to approximate their covariance and update them together.

ter Braak and Vrugt [2008] propose that the number of walkers can be significantly reduced by sampling not only from the current position of the chains, but from their past positions as well. This algorithm, referred to as DE-MC_Z is adaptive—i.e. the proposal distribution is time-dependent, violating time-homogeneity. This class of algorithms requires a different approach to the proof of convergence. In the case of many adaptive algorithms, this can be done by showing that the adaptation itself diminishes with time [Roberts and Rosenthal, 2007, 2009].

We do not, however, apply DE-MC_Z, as storing all past iterations leads to huge memory use with the large parameter dimensionality in IceChrono, where for the non-adaptive method we can periodically move the previous iterations to a file, and clear them from active memory. We find that the tradeoff with computational time is minimal, since keeping large amounts of data in memory can significantly worsen the performance of Python code.

On the movie strategy and probabilistic chronologies

Stochastic methods also differ from gradient-based and analytical methods for solving Bayesian inverse problems in the form of the result. Gradient-based or analytical solutions are often Gaussian in shape, and can be expressed as a mean and standard deviation, in the case of a unidimensional distribution, or a mean and covariance matrix in the case of a multidimensional distribution. Gaussian statistics have a considerable advantage in usability—it is easy to make calculations using a Gaussian distribution.

There is no analog that is generally applicable to all probability distributions. The stochastic solution to an inverse problem is thus the ensemble of accepted fits, which itself forms the posterior probability distribution.

For problems with only a few parameter dimensions, this ensemble can easily be expressed as a histogram. This is the case for the chapter in this thesis treating CO₂ and Antarctic Temperature. In the article, we provide maxima and 68% and 98% probability intervals, which roughly correspond to the mean, one standard deviation, and two standard deviations for a Gaussian distributions. However, this is only a rough approximation of the true solution to the inverse problem that we treat in the article—the timings of major, approximately linear changes in the records of CO₂ and Antarctic Temperature from the last deglaciation. The “true” solution to the inverse problem is in fact the full output of the MCMC procedure applied, represented as the histogram of the timings.

Representing the stochastic result of an inverse problem becomes more complicated in higher dimensions. A histogram made for a 1000-dimensional parameter space, on the order of what we might use when dating an ensemble of ice cores, would be unwieldy and mostly useless.

Tarantola [2005] proposes what he refers to as the “movie strategy” to display and understand the results of high-dimensional inverse problems. For a parameter vector like the age of ice in an ice core, rather than showing hundreds of histograms for each element in the vector, we can show an animation of the entire vector over the set of proposed chronologies accepted in a stochastic simulation. In effect, we display a movie of Ice Age to represent the best estimate of the probability density of the Ice Age parameter vector.

To provide a chronology to the paleoclimate community in “Movie Strategy” form, we can provide an ensemble of chronologies that is representative of the posterior distribution. In practice, this would be a file containing a large set of chronologies accepted in the stochastic simulation. This set could be reduced from the full simulation by removing some initial runs to eliminate any initialization bias, and selecting one out of every n




Figure 7: A corner plot, with univariate histograms on the diagonal and bivariate histograms in the other squares of the matrix, is an efficient way to represent low-dimensional multivariate probability densities.

accepted proposals, where n is on the order of the autocorrelation time of the proposals accepted by the simulation.

However, this is not to say that the Gaussian form of presenting chronologies should be discarded, even when the solution method to the chronological inverse problem is necessarily stochastic. It is certainly worth addressing, when a stochastic simulation is finished, if the density of the accepted scenarios is even approximately Gaussian. Then, an approximate mean scenario and covariance matrix could be provided to users, making calculations considerably easier.

There is use as well in providing the most basic form of a chronology, in terms of a mean vector and a standard deviation (σ) vector. Multidimensional distributions are not accurately represented when expressed by a mean and sigma for each age-depth match, since the covariances between uncertainties in the age model are not given to the user, who is forced to treat parameters that covary as effectively independent. However, simple calculations can be performed using the mean chronology, with the standard deviations used to assess the chronological uncertainty of conclusions made for an individual, discrete point. When the probability densities that result from a stochastic simulation are far from Gaussian, we can give the maximum and central 68% and 95% confidence intervals, which are roughly analogous in terms of probability to the mean, one standard deviation, and two standard deviations of a Gaussian distribution. These intervals can then be used to

assess punctual chronological uncertainty.

For the studies presented in this thesis, we follow an overall strategy that aims to make the presented result both accurate and usable. We first analyze, and provide to the user, the entire result of the stochastic simulation. In the case of the study on CO₂ treated in Chapter 3, the result is presented in the form of composite histograms, and in the case of the study on synchronization presented in chapter 4, we present animations. Then, we also treat simplifications of the probability densities (i.e. the maximum and central 68% and 95% confidence intervals in chapter 3), and provide these to the user as well, specifying their limitations and how well they represent the true probabilities. Thus, a paleoclimatologist wishing to use the statistically most-correct version of a chronology can access the ensemble version, whereas a paleoclimatologist looking to make a quicker calculation can use the simplification, keeping its limitations in mind.

References

- D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. Emcee: The mcmc hammer. *Publications of the Astronomical Society of the Pacific*, 125:306, March 2013. doi: 10.1086/670067.
- Jonathan Goodman and Jonathan Weare. Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science*, 5(1):65–80, 2010.
- W Keith Hastings. Monte Carlo sampling methods using Markov chains and their applications. 1970.
- Eric Laloy and Jasper A Vrugt. High-dimensional posterior exploration of hydrologic models using multiple-try dream (zs) and high-performance computing. *Water Resources Research*, 48(1), 2012.
- Bénédicte Lemieux-Dudon, Eric Blayo, Jean-Robert Petit, Claire Waelbroeck, Anders Svensson, Catherine Ritz, Jean-Marc Barnola, Bianca Maria Narcisi, and Frédéric Parrenin. Consistent dating for Antarctic and Greenland ice cores. *Quaternary Science Reviews*, 29(1-2):8–20, 2010.
- Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, 2(2):164–168, 1944.
- Luan Lin, Deborah Khider, Lorraine E Lisiecki, and Charles E Lawrence. Probabilistic sequence alignment of stratigraphic records. *Paleoceanography*, 29(10):976–989, 2014.

- Donald W Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
- Nicholas Metropolis and Stanislaw Ulam. The Monte Carlo method. *Journal of the American statistical association*, 44(247):335–341, 1949.
- Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- Benjamin Nelson, Eric B Ford, and Matthew J Payne. Run dmc: an efficient, parallel code for analyzing radial velocity observations using n-body integrations and differential evolution markov chain monte carlo. *The Astrophysical Journal Supplement Series*, 210(1):11, 2013.
- James R Norris. *Markov chains*. Number 2. Cambridge university press, 1998.
- Frédéric Parrenin, Lucie Bazin, Emilie Capron, Amaelle Landais, Bénédicte Lemieux-Dudon, and Valérie Masson-Delmotte. Icechronol: a probabilistic model to compute a common and optimal chronology for several ice cores. *Geoscientific Model Development*, 8(5):1473–1492, 2015.
- Gareth O Roberts and Jeffrey S Rosenthal. Coupling and ergodicity of adaptive markov chain monte carlo algorithms. *Journal of applied probability*, 44(2):458–475, 2007.
- Gareth O Roberts and Jeffrey S Rosenthal. Examples of adaptive mcmc. *Journal of Computational and Graphical Statistics*, 18(2):349–367, 2009.
- Sanjib Sharma. Markov chain monte carlo methods for bayesian data analysis in astronomy. *Annual Review of Astronomy and Astrophysics*, 55:213–259, 2017.
- Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*, volume 89. siam, 2005.
- Cajo JF ter Braak and Jasper A Vrugt. Differential evolution markov chain with snooker updater and fewer chains. *Statistics and Computing*, 18(4):435–446, 2008.
- Mai Winstrup, AM Svensson, Sune Olander Rasmussen, Ole Winther, EJ Steig, and AE Axelrod. An automated approach for annual layer counting in ice cores. *Climate of the Past*, 8(6):1881–1895, 2012.

Chapter 3

The Last Deglaciation : Antarctic Temperature and CO₂

This article has been published in *Climate of the Past*.



Antarctic temperature and CO₂: near-synchrony yet variable phasing during the last deglaciation

Jai Chowdhry Beeman¹, Léa Gest¹, Frédéric Parrenin¹, Dominique Raynaud¹, Tyler J. Fudge², Christo Buizert³, and Edward J. Brook³

¹Université Grenoble Alpes, CNRS, IRD, IGE, 38000 Grenoble, France

²Department of Earth and Space Sciences, University of Washington, Seattle, WA 98195, USA

³College of Earth, Ocean, and Atmospheric Sciences, Oregon State University, Corvallis, OR 97331, USA

Correspondence: Frédéric Parrenin (frederic.parrenin@univ-grenoble-alpes.fr)

Received: 19 March 2018 – Discussion started: 23 April 2018

Revised: 12 April 2019 – Accepted: 30 April 2019 – Published: 22 May 2019

Abstract. The last deglaciation, which occurred from 18 000 to 11 000 years ago, is the most recent large natural climatic variation of global extent. With accurately dated paleoclimate records, we can investigate the timings of related variables in the climate system during this major transition. Here, we use an accurate relative chronology to compare temperature proxy data and global atmospheric CO₂ as recorded in Antarctic ice cores. In addition to five regional records, we compare a $\delta^{18}\text{O}$ stack, representing Antarctic climate variations with the high-resolution robustly dated WAIS Divide CO₂ record (West Antarctic Ice Sheet). We assess the CO₂ and Antarctic temperature phase relationship using a stochastic method to accurately identify the probable timings of changes in their trends. Four coherent changes are identified for the two series, and synchrony between CO₂ and temperature is within the 95 % uncertainty range for all of the changes except the end of glacial termination 1 (T1). During the onset of the last deglaciation at 18 ka and the deglaciation end at 11.5 ka, Antarctic temperature most likely led CO₂ by several centuries (by 570 years, within a range of 127 to 751 years, 68 % probability, at the T1 onset; and by 532 years, within a range of 337 to 629 years, 68 % probability, at the deglaciation end). At 14.4 ka, the onset of the Antarctic Cold Reversal (ACR) period, our results do not show a clear lead or lag (Antarctic temperature leads by 50 years, within a range of –137 to 376 years, 68 % probability). The same is true at the end of the ACR (CO₂ leads by 65 years, within a range of 211 to 117 years, 68 % probability). However, the timings of changes in trends for the individual proxy records show variations from the stack, indicat-

ing regional differences in the pattern of temperature change, particularly in the WAIS Divide record at the onset of the deglaciation; the Dome Fuji record at the deglaciation end; and the EDML record after 16 ka (EPICA Dronning Maud Land, where EPICA is the European Project for Ice Coring in Antarctica). In addition, two changes – one at 16 ka in the CO₂ record and one after the ACR onset in three of the isotopic temperature records – do not have high-probability counterparts in the other record. The likely-variable phasing we identify testify to the complex nature of the mechanisms driving the carbon cycle and Antarctic temperature during the deglaciation.

1 Introduction

Glacial–interglacial transitions, or deglaciations, mark the paleorecord approximately every 100 000 years over the past million years or so (Williams et al., 1997; Lisiecki and Raymo, 2005; Jouzel et al., 2007). The last deglaciation, often referred to as glacial termination 1 (T1), offers a case study for a large global climatic change, very likely in the 3–8 °C range on the regional scale (Masson-Delmotte et al., 2013), and thought to be initiated by an orbitally driven insolation forcing (Hays et al., 1976; Berger, 1978; Kawamura et al., 2007). The canonical interpretation of this apparent puzzle is that insolation acts as a pacemaker of climatic cycles and the amplitude of glacial–interglacial transitions is mainly driven by two strong climatic feedbacks: atmospheric CO₂ and continental-ice-surface–albedo changes. However,

the mechanisms that control the CO₂ rise are still a matter of debate. Accordingly, reconstructing the phase relationship (leads and lags) between climate variables and CO₂ during the last termination has become important, and has a substantial history in ice core research (Barnola et al., 1991; Raynaud and Siegenthaler, 1993; Caillon et al., 2003; Pedro et al., 2012; Parrenin et al., 2013).

Global temperature has been shown to lag CO₂ on average during T1 (Shakun et al., 2012), supporting the importance of CO₂ as an amplifier of orbitally driven global-scale warming. But Antarctic temperature and CO₂ concentrations changed much more coherently as T1 progressed. Indeed, midway through the glacial–interglacial transition, Antarctic warming and CO₂ increase slowed and even reversed during a period of about 2000 years, coinciding with a warm period in the north called the Bølling–Allerød (B/A). The respective period of cooling in Antarctica is called the Antarctic Cold Reversal (ACR). A period of cooling in the Northern Hemisphere known as the Younger Dryas (YD), followed the B/A, coinciding with a period of warming in the Southern Hemisphere (SH).

High-latitude Southern Hemisphere paleotemperature series – including Southern Ocean temperature – varied similarly to Antarctic temperature during T1 (Shakun et al., 2012; Pedro et al., 2016). Upwelling from the Southern Ocean is thought to have played an important role in the deglacial CO₂ increases (Anderson et al., 2009; Burke and Robinson, 2012; Schmitt et al., 2012; Rae et al., 2018). The Atlantic Meridional Overturning Circulation, or AMOC, a major conduit of heat between the Northern Hemisphere and Southern Hemisphere and component of the bipolar seesaw – the umbrella term encompassing the mechanisms thought to control the seemingly alternating variations in Northern Hemisphere and Southern Hemisphere temperature (Stocker and Johnsen, 2003; Pedro et al., 2018) – is thought to have influenced Southern Ocean upwelling during the deglaciation (Anderson et al., 2009; Skinner et al., 2010). A weakening of the oceanic biological carbon pump appears to have dominated the deglacial CO₂ increase until 15.5 ka (kiloannum before 1950 CE), when rising ocean temperature likely began to play a role as well (Bauska et al., 2016).

Ice sheets are exceptional archives of past climates and atmospheric composition. Local temperature is recorded in the isotopic composition of snow/ice (NorthGRIP Project Members, 2004; Jouzel et al., 2007) due to the temperature-dependent fractionation of water isotopes (Lorius and Merlivat, 1975; Johnsen et al., 1989). The concentration of continental dust in ice sheets is a proxy of continental aridity, atmospheric transport intensity and precipitation. Finally, air bubbles enclosed in ice sheets are near-direct samples of the past atmosphere. However, the age of the air bubbles is younger than the age of the surrounding ice, since air is locked-in at the base of the firn (on the order of 70 m below the surface on the West Antarctic Ice Sheet (WAIS) Divide) at the lock-in depth (LID) (Buizert and Severinghaus, 2016).

The firn, from top to bottom, is composed of a convective zone (CZ), where the air is mixed vigorously, and a diffusive zone (DZ), where molecular diffusion dominates transport. Firn densification models can be used to estimate the LID and the corresponding age difference (Sowers et al., 1992).

Atmospheric CO₂ concentrations, recorded in the air bubbles enclosed in ice sheets, are better preserved in Antarctic ice than in Greenland ice because the latter has much higher concentrations of organic material and carbonate dust (Raynaud et al., 1993; Anklin et al., 1995). Measured on the Vostok and EPICA Dome C ice cores (EPICA is the European Project for Ice Coring in Antarctica), the long-term history of CO₂ (Lüthi et al., 2008) covers the last 800 kyr.

Early studies suggested that at the initiation of the termination around 18 ka, just after the Last Glacial Maximum (LGM), Antarctic temperature started to warm 800 ± 600 years before CO₂ began to increase (Monnin et al., 2001), a result that was sometimes misinterpreted to mean that CO₂ was not an important amplification factor of the deglacial temperature increase. This study used measurements from the EPICA Dome C (EDC) ice core (Jouzel et al., 2007) and a firn densification model to determine the air chronology. However, this firn densification model was later shown to be in error by several centuries for low-accumulation sites such as EDC during glacial periods (Loulergue et al., 2007; Parrenin et al., 2012).

Two more recent works (Pedro et al., 2012; Parrenin et al., 2013) used stacked temperature records and improved estimates of the age difference between ice and air records to more accurately estimate the relative timing of changes in Antarctic temperature and atmospheric CO₂ concentration. In the first of these studies, measurements from the higher accumulation ice cores at Siple Dome and Byrd Station were used to decrease the uncertainty in the ice–air age shift, and indicated that CO₂ lagged Antarctic temperature by 0–400 years on average during the last deglaciation (Pedro et al., 2012). The second study (Parrenin et al., 2013) used measurements from the low accumulation EDC ice core but circumvented the use of firn densification models by using the nitrogen isotope ratio $\delta^{15}\text{N}$ of N₂ as a proxy of the DZ height, assuming that the height of the CZ was negligible during the study period. CO₂ and Antarctic temperature were found to be in phase at the beginning of T1 (-10 ± 160 years) and at the end of the ACR period (-60 ± 120 years), but CO₂ was found to lag Antarctic temperature by several centuries at the beginning of the Antarctic Cold Reversal (260 ± 130 years) and at the end of the deglacial warming in Antarctica (500 ± 90 years). The end of the deglacial warming in Antarctica occurred roughly 2 centuries after the onset of the Holocene period dated at 11.7 ka according to the International Commission on Stratigraphy. However, the assumption that $\delta^{15}\text{N}$ reflects DZ height is imperfect as it may underestimate the DZ height for sites with strong barometric pumping and layering (Buizert and Severinghaus, 2016).

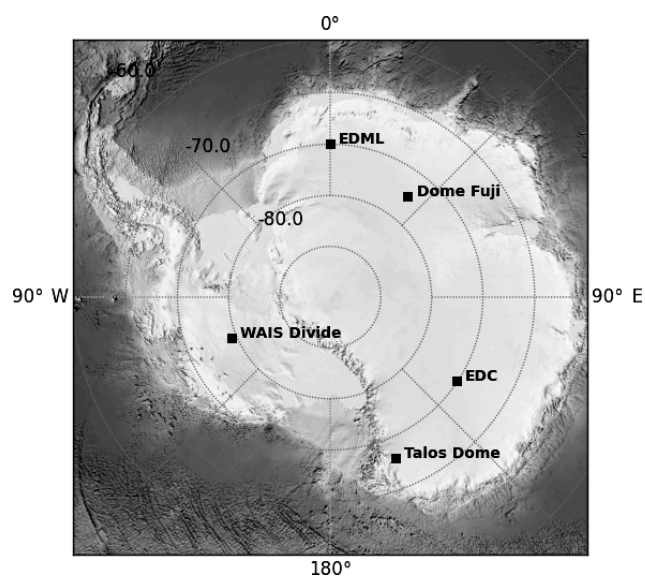


Figure 1. Drilling locations of the ice cores where the CO₂ and isotopic paleotemperature records included in this study were measured.

A new CO₂ record of unprecedented high resolution (Marcott et al., 2014) from the WAIS Divide (WD) ice core merits the reopening of this investigation. The air chronology of WAIS Divide is well constrained thanks to a relatively high accumulation rate and to accurate nitrogen-15 measurements (Buizert et al., 2015). The WAIS record evidences centennial-scale changes in the global carbon cycle during the last deglaciation superimposed on more gradual, millennial-scale trends that bear a resemblance to Antarctic temperature (Marcott et al., 2014).

The deglacial temperature rise seen at WD is structurally similar to that at other Antarctic sites. However, the rise in West Antarctic temperature shows an early warming starting around 21 ka, following local insolation (Cuffey et al., 2016). This early warming trend is much more gradual in records from East Antarctic ice cores. The difference between the two records may be related to sea ice conditions around East Antarctica and West Antarctica, and perhaps to elevation changes (WAIS Divide Project Members, 2013; Cuffey et al., 2016). The temperature record at WAIS Divide shows an acceleration in warming around 18 ka which is also present in East Antarctic records (WAIS Divide Project Members, 2013).

On the much shorter timescales of the satellite era, Jones et al. (2016) note differing temperature trends at the drilling sites of the five cores used in this study. On the other hand, the interpretation of individual isotopic records can prove to be complicated, as local effects – including those of ice sheet elevation change and sea ice extent – are difficult to correct.

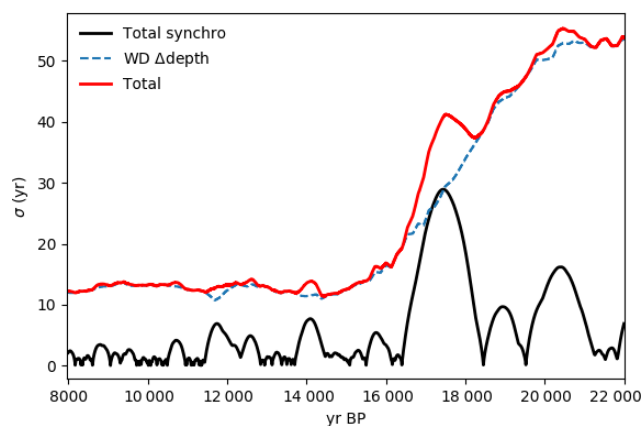


Figure 2. Relative chronological uncertainty between ATS3 and the WD CO₂ record (red line) calculated as the quadratic sum of the synchronization error from Buizert et al. (2018) (black line) and the Δ age uncertainty from WD2014 (dashed blue line).

In the present work we refine our knowledge of leads and lags between Antarctic temperature and CO₂. We use a stack of accurately synchronized Antarctic temperature records (Buizert et al., 2018) to reduce local signals, placed using volcanic matching on the WAIS Divide chronology (WD2014). We then compare the temperature stack to the high-resolution WAIS Divide CO₂ record by determining the probable timings of changes in trends, and calculate probable change point timings for the five individual isotope-derived records used in the stack as well.

2 Methods and data

2.1 Temperature stack and ice chronology

We use the $\delta^{18}\text{O}$ stack developed by Buizert et al. (2018) (referred to hereafter as Antarctic temperature Stack 3, or ATS3) to represent Antarctic temperature. The use of the stack allows us to remove local influences and noise in the individual records to the greatest extent possible. The stack contains five records: EDC, Dome Fuji (DF), Talos Dome (TD), EPICA Dronning Maud Land (EDML) and WAIS Divide (WD). The drilling site locations are shown in Fig. 1. Volcanic ties between WD and EDC, WD and TD, and WD and EDML are developed in Buizert et al. (2018); previously published volcanic ties were used between EDC and DF (Fujita et al., 2015), placing all of the records on the WD2014 chronology (Buizert et al., 2015). Notably, the Vostok record, included in the stack used by Parrenin et al. (2013), is excluded from the Buizert et al. (2018) stack: it contains additional chronological uncertainty as it is derived using records from two drilling sites. We take the quadratic sum of the synchronization error from Buizert et al. (2018) and the Δ age uncertainty from WD2014 to calculate the relative chronological error between ATS3 and the WD CO₂ record (Fig. 2).

2.2 CO₂ and air chronology

We use atmospheric CO₂ data from the WD ice core (Marcott et al., 2014) which consist of 1030 measurements at 320 depths that correspond to ages between 23 000 and 9000 years BP with a median resolution of 25 years. At WD, the age offset between the ice and air (trapped much later) at a given depth, Δage , is calculated using a firn densification model, which is constrained using nitrogen-15 data, a proxy for firn column thickness (Buizert et al., 2015). Δage ranges from 500 ± 100 years at the last glacial maximum, to 200 ± 30 years during the Holocene. Δage uncertainty is added to cumulative layer counting uncertainty to determine the total uncertainty in the air chronology.

2.3 Identifying changes in trends

We identify likely change points using piecewise linear functions. Residuals are calculated between the raw data and linear functions with a fixed number of stochastically proposed change points, which are free to explore the entire temporal range of the time series (similarly to Parrenin et al., 2013). These residuals are summed to form a cost function, which allows us to perform a Bayesian analysis of the probable timing of change points. At the base of our method is a parallelized Metropolis–Hastings (MH) procedure (Goodman and Weare, 2010; Foreman-Mackey et al., 2013). Therefore, we do not present a single “best fit” but rather analyze the ensemble of fits accepted by the routine. We plot two histograms: an upward-oriented histogram for concave-up change points and a downward-oriented histogram for concave-down change points. We use these histograms as probabilistic locators of changes in slope (Fig. 3).

The change point representations of the ATS3 and CO₂ time series are composed of a set of n specified change points $\mathbf{X}_i = (x_i, y_i) | i = 1, \dots, n$. We denote the vector of m time series observations o at time t $\mathbf{O}_l = (t_l, o_l) | l = 1, \dots, m$, and the scalar residual term J between observations and the linear interpolation between change points f_y :

$$J(\mathbf{X}_i) = \mathbf{R}^T \mathbf{C}^{-1} \mathbf{R}; \mathbf{r}_l = (f_y(t_l) - o_l)_l, \quad (1)$$

where \mathbf{R} is the vector of residuals at each data point with components \mathbf{r}_l , and \mathbf{C} is the covariance matrix of the residuals. The ATS3 series contains 1412 data points and the WD CO₂ series contains 320, each of which is considered in the residuals.

We fix $x_0 = t_0$ and $x_n = t_l$, i.e., the x values of the first and last change points are fixed to the first and last x values of the observation vector, with the y values allowed to vary. The remaining points are allowed to vary freely in both dimensions.

2.3.1 Estimating the covariance matrix \mathbf{C} : treating uncertainty and noise

Our method fits time series with piecewise linear functions, and the residual vector thus accounts for any variability that cannot be represented by these fits. Paleoclimate time series, like the CO₂ and ATS3 series used here, typically contain autocorrelated noise (see Mudelsee (2002), for example) which cannot be accurately represented by a piecewise linear function. Weighting the residuals of a cost-function-based formulation by a properly estimated inverse covariance matrix ensures that this autocorrelated noise is not overfitted, and can improve the balance of precision and accuracy of the fits.

Our time series contain two potential sources of uncertainty: measurement or observational uncertainty, related with the creation of the data series, and modeling uncertainty, related to the formulation of the fitting function. We formulate a separate covariance matrix to account for each source of uncertainty. These matrices are then summed to form \mathbf{C} . We assume the measurement uncertainty to be uncorrelated in time (i.e., a white noise process). Thus, the associated covariance matrix \mathbf{C}_{meas} is diagonal and the diagonal elements \mathbf{C}_{jj} are each equal to the variance of observation o_j and σ_j^2 , as estimated during the measurement process.

The covariance matrix of the modeling uncertainty, which we denote as \mathbf{C}_{mod} , is more complicated since the residual vector contains any autocorrelated noise in the time series that is not accounted for by the piecewise linear fits. Additionally, the time series contain outliers with respect to these linear fits and these can impact any nonrobust estimate of covariance. Finally, an initial idea of the model must be used to calculate residuals, and thus estimate their covariance. These challenges can be circumvented when data resolution is low enough to assume that residuals are uncorrelated, as in Parrenin et al. (2013); however, including the covariance matrix allows us to make use of noisy high-resolution data.

We arrive at an initial model by running a MH simulation in which \mathbf{C} is assumed equal to the identity matrix, and select the best fit of this run. Note that \mathbf{C}_{meas} is not taken into account at this point since we require an independent estimate of \mathbf{C}_{mod} . At this point, covariance could be estimated directly but tests indicated that this method was not robust, making the covariance matrix estimate sensitive to outliers and to the initial model fit. Our CO₂ data are unevenly spaced in time and developing a covariance matrix using the traditional covariance estimator would require some form of interpolation, which can introduce substantial error.

The residuals with respect to the initial model are instead used to fit an AR(1) model (Robinson, 1977; Mudelsee, 2002) which treats the autocorrelation between a pair of residuals r_i and r_{i-1} as a function of the separation between the two data points in time, $t_i - t_{i-1}$. The Robinson (1977) and Mudelsee (2002) model is expressed as follows:

$$r_i = r_{i-1} \cdot a^{t_i - t_{i-1}}, \quad (2)$$

where the constant a determines the correlation between two residuals separated by $t_i - t_{i-1}$ units of time, and minimizing the loss function:

$$S(a) = \sum_{i=1}^n \{r_i - r_{i-1} \cdot a^{t_i - t_{i-1}}\} \quad (3)$$

allows us to estimate a . We do so using a nonlinear least-squares estimate with L1-norm regularization to provide a robust estimate (Chang and Politis, 2016). We test the validity of the AR(1) hypothesis by comparing r_i with $r_{i-1} \cdot a^{t_i - t_{i-1}}$ (the Supplement). Given that the AR(1) hypothesis cannot be rejected, we can use a to calculate the theoretical correlation between two residuals, and construct a correlation matrix \mathbf{K} and the model covariance matrix \mathbf{C}_{mod} as follows:

$$\mathbf{C}_{\text{mod}} = \sigma_{\text{mod}}^2 \mathbf{K}; \mathbf{K}_{ij} = a^{t_j - t_i}, \quad (4)$$

where σ_{mod}^2 is the variance in the modeling error, assumed constant and estimated using a robust estimator based on the interquartile range (IQR), calculated as $(\text{IQR}(\mathbf{R})/1.349)^2$ (Ghosh, 2018; Silverman, 1986). Finally, the covariance matrix of the residuals C is calculated as

$$\mathbf{C} = \mathbf{C}_{\text{mod}} + \mathbf{C}_{\text{meas}}. \quad (5)$$

Rather than inverting the covariance matrix, we use Cholesky and lower–upper (LU) decompositions to solve for the cost function value J , as in Parrenin et al. (2015).

2.4 Estimating the posterior probability density

In general, the probability density of the change points cannot be assumed to follow any particular distribution, as short-timescale variations of the time series may lead to multiple modes or heavy tails, for example. Thus, stochastic methods, which are best adapted to exploring general probability distributions (for example, Tarantola, 2005), are suited to our problem.

To tackle the large computation time required for traditional MH sampling, we apply the ensemble sampler developed by Goodman and Weare (2010) (GW) as implemented in the Python emcee library (Foreman-Mackey et al., 2013). This sampler adapts the MH algorithm so that multiple model walkers can explore the probability distribution at once, making the algorithm parallelizable. It has the advantage of being affine invariant: that is, steps are adapted to the scale of the posterior distribution in a given direction.

The final task in our piecewise linear analysis is to identify the number of change points to best represent the two series we wish to analyze. The choice should reflect our goal of accurately investigating millennial-scale variability. Further, we aim for parsimony in the representation. To best balance these two goals, we apply the Bayesian information criterion (BIC; Schwarz, 1978) to the number of points we allow to fit the two series.

We apply a joint BIC – normalizing the cost function for each series by its lowest value – and arrive at the conclusion that the two series are best compared by fitting eight points. The histograms created for fits of seven, eight and nine points of the two series are remarkably similar, and we assess that our choice of eight points does not add significant uncertainty to the timing of change points, with the exception of the change point at the ACR onset. We include histograms of fits between five and nine points in the Supplement. We also include change point timings and lead-lag estimates calculated using seven-point fits in the Supplement.

The most probable timings are identified by probability peaks, or modes, for a fit of n points; we analyze the n time periods with greatest contiguous cumulative probability. Thus, we analyze a coherent number of change points, and avoid setting artificial probability thresholds. We avoid comparing incoherent modes by separating changes by the sign of the change in slope of the fits. If the slope decreases at a change point, the change in slope is negative or concave-down. These changes are indicated by the downward-facing part of the histogram graphs. Note that while this part of the histogram appears “negative”, probabilities cannot be negative; and this simply indicates that the probability is for a concave-down change point. If the slope increases at a change point, the change in slope is positive or concave-up. The probabilities of these changes are indicated by the upward-facing or “positive” part of the histogram. When we calculate leads and lags, we only do so for either a region in which there is a probability peak for a concave-down change point in both series, or for a concave-up change point in both series, but we do not treat concave-up probability and concave-down probability together.

2.5 Phasing

We estimate $\rho_{\text{lead}}^{\text{ATS3}}$, the probability that ATS3 leads CO₂ over a given interval as

$$\rho_{\text{lead}}^{\text{ATS3}} = (\rho_x^{\text{ATS3}} \circ \rho_x^{\text{CO}_2}) \star \rho^{\text{chron}}, \quad (6)$$

where ρ_x^{ATS3} is the probability of a change point at time x for ATS3; $\rho_x^{\text{CO}_2}$ is the probability of a change point at time x for CO₂; \circ is the cross-correlation operator, which is used to calculate the probability of the difference between two variables; and \star is the convolution operator, which is used to calculate the probability of the sum of two variables. ρ^{chron} is the probability distribution of the chronological uncertainty between the two records, which we take to be Gaussian centered on 0, with standard deviation $\sigma = \sigma_{\text{chron}}$ (shown in Fig. 3). The intervals associated with each change point are given in Fig. 6.

3 Results and discussion

3.1 Change point timings

The change point histograms for the ATS3 and CO₂ time series in Fig. 3 confirm that the millennial-scale changes in the two series were largely coherent. We focus on four major changes in trends which are common to both series: the onset of the deglaciation from 18.2 to 17.2 ka, the onset of the Antarctic Cold Reversal (ACR) at around 14.5 ka, the ACR end between 12.9 and 12.65 and the end of the deglaciation at approximately 11.5 ka. For each of these four changes, we calculate the probability of a lead or lag over the time interval that encompasses the continuous peaks in the two histograms. Two CO₂ change points, one centered at approximately 16 ka, and one just before the ACR onset at 14.4 ka, do not have high-probability counterparts in the ATS3 series. A low-probability change point for the temperature series after the ACR onset, centered at 14 ka, does not have a counterpart in the CO₂ series.

The deglaciation onset begins with a large, positive change point mode for Antarctic temperature, centered around 18.1 ka. The corresponding change point for the CO₂ series is centered around 17.6 ka.

The CO₂ rise peaks at around 16 ka, identified by a downward-oriented probability peak, which has no counterpart in the temperature histogram. This peak is followed by a brief plateau in CO₂ concentrations, before a gradual, accelerating resumption of the increase.

At the onset of the Antarctic Cold Reversal, an upward-facing CO₂ change point at 14.4 ka is followed by a broad, downward-facing CO₂ change point which peaks at around 14.25 ka. The first peak appears to reflect a centennial-scale change (identified by Marcott et al., 2014) and the broadness of the second peak reflects further methodological uncertainty with respect to the timing of the millennial-scale change in CO₂. An unambiguous negative temperature change also occurs at around 14.3 ka, roughly concurrent with the downward CO₂ change point. Antarctic temperature began to descend after the ACR onset, and it is worth mentioning the low-probability concave-up change point mode centered on 13.9 ka, particularly because this point is much more probable for some of the individual isotopic records. No corresponding change point is detected for CO₂.

The ACR termination is represented by large probability peaks in both series. An increase in CO₂ began at the peak occurring around 12.85 ka, while the ATS3 increase is centered at approximately 12.78 ka reaching its maximum around 12.7 ka.

The end of the deglacial warming in Antarctica is well-defined in the ATS3 series, with a large mode reaching its maximum at 11.7 ka. The corresponding CO₂ mode reaches its maximum at around 11.15 ka. Visually, we might question why the CO₂ change point is deemed to most probably occur at 11.15 ka, when a kink in the series appears to occur closer

to 11.5 ka. It is important to note that two minor spikes in probability appear to fit the rapid rise that occurs close to 11.5 ka, and that the downward spike at the end of this rise indeed indicates that a small number of accepted iterations do indeed fit a change point here.

Since the resolution of the WAIS Divide dataset decreases considerably after the Holocene onset, and only three points account for the majority of the rapid rise that occurs before the Holocene onset, most of the weight is given to the obvious line beginning at the ACR end. Adding an additional point should allow us to fit this slightly better; the CO₂ series is slightly better fit with nine points according to the individual BIC values, and there is more probability around the rapid rise in the nine-point fit, though the peak at 11.15 ka is still dominant. In any case, some methodological uncertainty exists regarding the location of this point, and the probability estimate is possibly biased by the quick change in resolution. Better resolution around this point will help identify the true location of the change.

As a second test of the timings of millennial-scale events, we use our method to fit filtered versions of the ATS3 and WAIS Divide CO₂ data. A Savitzky–Golay filter, designed to have an approximate cutoff periodicity of 500 years, is applied to the two records. Fitting change points to these two series allows us to verify that our leads and lags are not overly influenced by submillennial scale noise in the original records.

Figure 4 shows the Savitzky–Golay filtered CO₂ and ATS2 time series, and the corresponding change point histograms. The four major changes identified in both series, at the T1 onset, the ACR onset, the ACR end and T1 end, are similar in shape and center to the change points identified for the raw data. However, there are two notable differences between the two fits. First, the histograms are smoother and have broader peaks. This is not surprising, given that the Savitzky–Golay filters are designed to remove all variability with periodicities less than 500 years, whereas the covariance matrix applied to the fits of the raw data only treats an approximation of AR(1)-correlated noise. Second, the pre-ACR change in CO₂ is removed from the filtered series, which is again reasonable as it appears to mark a centennial-scale event. Savitzky–Golay filtering has its own drawbacks – data reinterpolation is required, for example, and propagating measurement uncertainty becomes difficult. However, the similarity of the two results supports our fits of the raw data.

3.2 Change point timings for individual temperature records

Histograms calculated for each of the regional $\delta^{18}\text{O}$ records are shown in Fig. 5. These histograms should still be interpreted cautiously, as additional information included in the isotopic records here assumed to represent temperature – the signal of ice sheet elevation change, for example – are not corrected for. The comparison of these histograms pro-

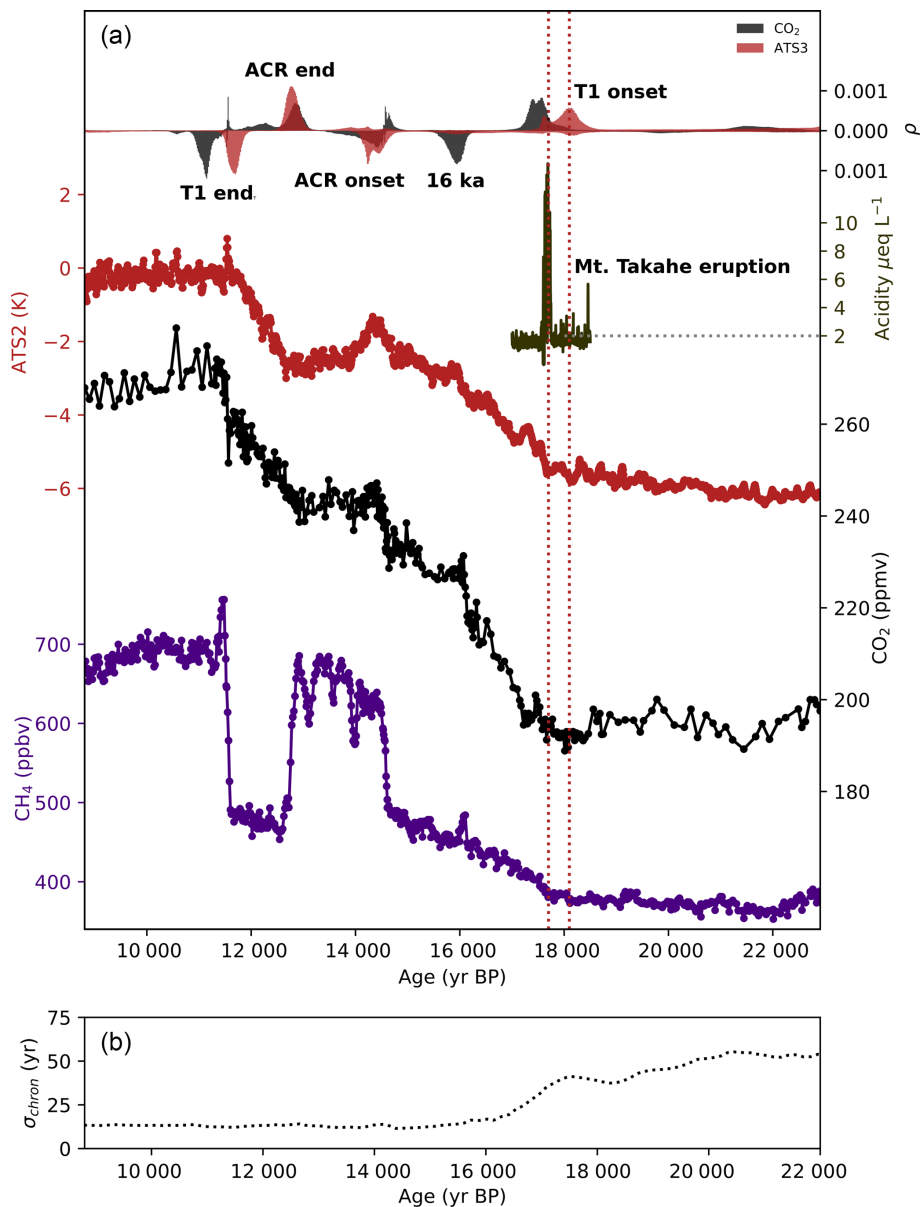


Figure 3. (a) Atmospheric CO₂ (black) and ATS3 (red) placed on a common timescale with the normalized histograms of probable change points (eight-point simulations, allowed to fit six points and two endpoints per series). Histograms are plotted downward-oriented when the rate of change decreases and upward-oriented when it increases (same colors, y axis not shown). Probabilities are normalized so that the integrated probability for a given histogram sums to 1. In four distinct time intervals, both series show concurrent probable change points. We also plot WD Acidity (army green) and WD CH₄ (violet) series. CH₄ tracks changes in Northern Hemisphere climate. CO₂ modes correspond with rapid changes in CH₄ at the ACR end, ACR onset, 16 ka rise and the rapid rise preceding the T1 end. (b) Chronological uncertainty, taken as the sum of the Δ age uncertainties and the uncertainty estimate for our volcanic synchronization.

vides an initial, exploratory picture of potential regional differences in climate change during the last termination.

Of the four changes identified as coherent between the temperature stack and CO₂, those at the deglaciation onset, the ACR end and the T1 end are expressed as probability peaks in all five records. Some ambiguity appears to exist about the timing of the ACR onset in the EDML record. It is expressed by a rather broad, low-probability mode extending

between 16 and 14 ka, though a large spike at 14 ka marks the downturn seen in the other records. The ACR onset is well-defined in all of the five other records. Three of the records – TD, EDC and WD – show a marked stabilization in temperature after the ACR onset, near 13.8 ka, which appears as a region of much lower probability in the ATS3 stack.

The WAIS Divide record is, notably, the only isotopic record in the stack from the West Antarctic Ice Sheet. We




Figure 4. (a) Savitzky–Golay filtered atmospheric CO₂ (black) and ATS3 (red) placed on a common timescale with the normalized histograms of probable change points (eight-point simulations, allowed to fit six points per series). Histograms are plotted downward-oriented when the rate of change decreases and upward-oriented when it increases (same colors, y axis not shown, probabilities range from 0 (center) to 0.0024 (top and bottom)). (b) Chronological uncertainty, taken as the sum of the Δ age uncertainties and the uncertainty estimate for our volcanic synchronization.

could thus reasonably expect it to show considerably different trends from the other records. Indeed, changes in the WD temperature record occur at 22 and 20 ka. This earlier change in the isotopic record was identified and confirmed to indeed be a temperature signal by Cuffey et al. (2016) using a borehole temperature record, though their study places the change at 21 ka. We confirm that the onset of the deglacial temperature rise in West Antarctica likely began as much as 4 kyr before the onset of temperature rise in East Antarctica. Interestingly, the WD record also shows a temperature change point around 17.8 ka, expressed slightly later than in the other records and more synchronous with CO₂. This apparent acceleration of the temperature rise is followed by a downward-facing change point not seen in any of the other records. A difference appears to exist in timing at the T1 end as well, with temperature change at WD appearing to precede the East Antarctic records and the DF temperature change, centered at 11.2 ka, occurring more synchronously with CO₂.

3.3 Leads and lags

The probability densities of leads and lags at the coherent change points between ATS3 and CO₂ are shown in Fig. 6. We report the central 68 % and 95 % probability intervals for each histogram. These values are grouped in Table 1.

ATS3 led CO₂ by 570 years, (within a 68 % interval of 128 to 751 years) at the T1 onset. Given the large range of uncertainty, though, we cannot exclude the possibility of synchrony at the 95 % level, which, interestingly, appears to be the case for the Dome Fuji record. At the ACR onset, we are not able to identify a clear lead or lag. At this point, phasing is sensitive to the number of points used to make the calculation: with 7 points, we calculate a 240 year lead of ATS3, and with 8 points, we calculate a 50 year lead. In neither of these cases can we exclude synchrony within 95 % probability, and with 8 points, it is well within 68 %. At the ACR end, CO₂ led ATS3, by $65 \pm$ years within a 68 % range of 211 years to -117 years (a temperature lead) and so again, the possibility of synchrony cannot be excluded within 68 % probability.

At the T1 end, a CO₂ lag is certain. Calculating the phasing between 12.0 and 11.0 ka, we obtain an ATS3 lead of 532 years, with a 68 % probability range of 337 to 629 years. This estimate is complicated, though, if we consider the small possibility that the true CO₂ change point occurs closer to 11.5 ka, at the end of the rapid rise. In this case, the phasing is reduced to 174 years (68 % central probability range of 65 to 280 years) and synchrony is within the 95 % central probability interval (-71 to 411 years).




Figure 5. Atmospheric CO₂ (black) and individual δ¹⁸O records placed on a common timescale, with the normalized histograms of probable change points (eight points) for each ice core used in the ATS3 stack; the locations of the drill sites are shown in the top right corner. Details of the histogram plots are as in Fig. 3. Maximum probability estimates and 68 % and 95 % probability intervals for timings of the individual records are provided in the Supplement. The δ¹⁸O records are given in per mill anomalies with respect to the last 200 years, as is the ATS3 stack.

Table 1. Maximum probabilities and central probability intervals for leads and lags at each of the selected change point intervals. Negative x-axis values indicate a CO₂ lead.

	95 % (lower)	68 % (lower)	Maximum	68 % (upper)	95 % (upper)	Parrenin et al. (2013)	σ
T1 onset	-338	127	570	751	1045	-10	160
ACR onset ^b	-357	-137	50	376	708	260	130
ACR end	-410	-211	-65	117	375	-60	120
T1 end ^a	45	337	532	629	773	500	90

^a Note the small, but distinct, probability that the CO₂ change point occurs closer to 11.5 ka would indicate a lag of 174 years, with a 68 % central probability range of 65 to 280 years. ^b The phasing at the ACR onset is sensitive to whether seven or eight points are used. Timings with seven- and eight-point fits, and for the individual isotopic records, are made available in the Supplement.

3.4 Discussion

Our results refine and complicate the timings, leads and lags identified by the most recent comparable studies (Pedro et al., 2012; Parrenin et al., 2013). We identify a CO₂ change point not treated in these studies at 16 ka and one before the ACR onset, associated with the centennial-scale rapid rises iden-

tified by Marcott et al. (2014). We also treat regional isotopic records and identify a change point occurring at 13.9 ka in three of the records. None of these change points have a marked counterpart in the other series.

During the major multi-millennial-scale changes which occur at T1 onset and T1 end, Antarctic temperature likely

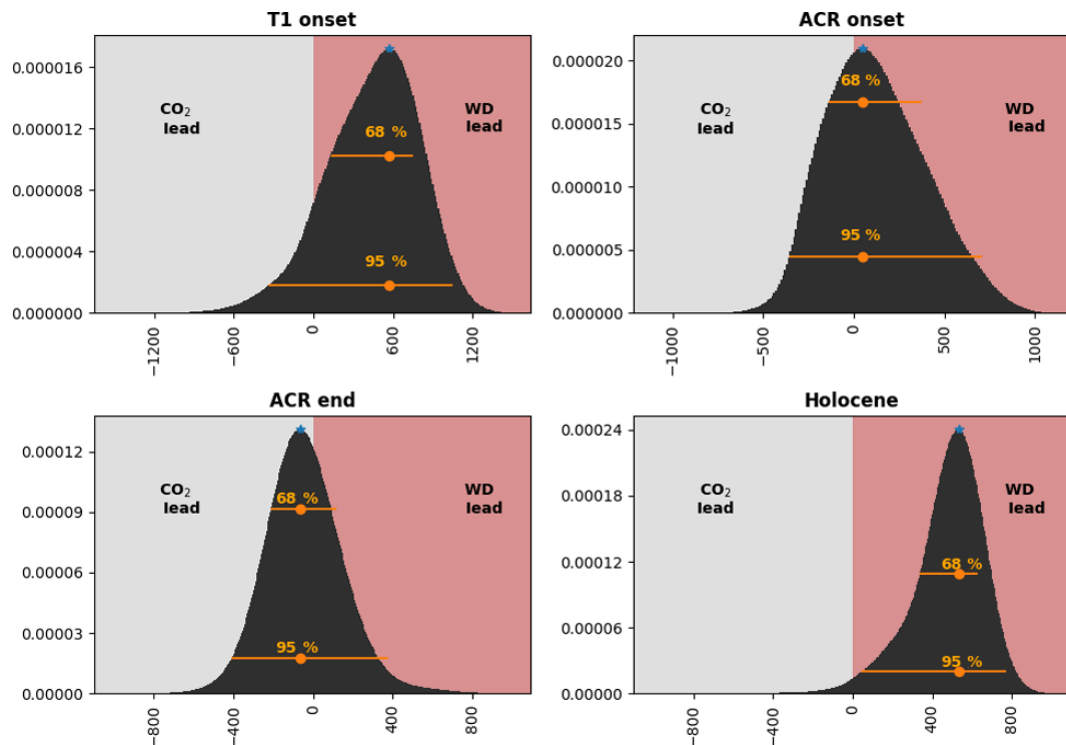


Figure 6. Probability density ρ (y axis, normalized) of an ATS lead (x-axes, in years) at each of the selected change point intervals (noted on subfigures). Negative x-axis values indicate a CO₂ lead. The maximum probability lead/lag and 68%/95% central probability intervals are indicated by the orange dot and lines on each histogram.

led CO₂ by several centuries. However, during the complex centennial-scale change at the ACR onset, we cannot calculate a clear lead of either ATS or CO₂; and at the end of the ACR, CO₂ leads temperature. Further, we neither identify a temperature analog for the CO₂ change at 16 ka nor an analog in CO₂ of the temperature stabilization in ATS3 after the ACR onset, itself not present in all of the regional $\delta^{18}\text{O}$ records, indicating at least some degree of decoupling during these changes. Additionally, the CO₂ changes at the ACR onset and T1 end are overlaid with centennial-scale substructures. Finally, synchrony is within the 2σ uncertainty range for each phasing, with the exception of the T1 end.

The changes in CO₂ occurring at the ACR onset, ACR end and the T1 end have been identified to correspond with changes in CH₄ (Marcott et al., 2014), which are thought to originate in tropical wetland sources (Chappellaz et al., 1997; Fischer et al., 2008; Petrenko et al., 2009) and are indicative of Northern Hemisphere and low-latitude temperature changes during the deglaciation (Shakun et al., 2012). Indeed, the CO₂ modes appear to demarcate the rapid changes in the WD CH₄ record, shown in Fig. 3.

The beginning of a gradual rise in CH₄ at around 18 ka appears to be near-synchronous with the T1 onset rise in Antarctic temperature. This rise is not seen in Greenland paleotemperature records, where it may have been masked by AMOC-driven wintertime cooling (Buizert et al., 2017) but

it also appears in proxy temperature stacks spanning both the Northern Hemisphere and Southern Hemisphere 0 to 30° latitude bands (Shakun et al., 2012).

Tephra from Mt. Takahe, a stratovolcano located in West Antarctica, have been detected in Antarctic ice cores during a 192-year interval around 17.7 ka. It has been postulated that this eruption may have provoked changes to large-scale SH circulation via ozone depletion, possibly triggering the transition between the gradual SH temperature rise beginning well before 18 ka and the more rapid rise marking the deglaciation (McConnell et al., 2017). The CO₂ mode we find at the deglaciation is coeval with this event within the range of dating uncertainty (Fig. 3), and CH₄ visually appears to accelerate concurrently. However, the cumulative probability of the ATS3 change point is much greater before 17.7 ka than after (approximately 80% of the probability density occurs before, see Fig. 3); hence, our results do not support the proposed volcanic forcing by McConnell et al. (2017) of the temperature change.

Though the T1 onset and the ACR end are both thought to originate in AMOC reductions (Marcott et al., 2014), our results allow for the CO₂–ATS3 phasing to be different during the two events, with the maximum probabilities reversed in directionality (i.e. with temperature leading at T1 and CO₂ leading at the ACR end, though zero phasing is within 95% error). Though CH₄ appears to change alongside CO₂ dur-

ing both intervals, the phasing between CO₂ and AT3 are opposite in direction and different in slope. This hints at a complex coupling, depending on conditions defined by multiple variables and mechanisms between CO₂ and Antarctic temperature. Bauska et al. (2016), for example, hypothesize that an earlier rise of CO₂ at 12.9 ka, driven by land carbon loss or SH westerly winds, might have been superimposed on the millennial-scale trend.

The apparent decoupling between CO₂ and AT3 at 16 ka also merits further discussion. We do not detect a change point for any of the isotopic records at 16 ka but the EDML record contains extremely broad uncertainty associated with the ACR onset peak, stretching to 16 ka, which indicates that this portion of the EDML time series is indeed notably different in shape from the other records, even if a clear signal is not identified at 16 ka by our method. EDML indeed appears to record changes in AMOC differently than the other isotopic records (Buizert et al., 2018; Landais et al., 2018). CO₂ and AT3 are similarly apparently decoupled at the temperature change point centered at 14 ka and this point could be indicative of variability specific to the Pacific and eastern Indian Ocean sectors, as it is present only in the TALOS Dome and EDC records, and slightly later, around 13.7 ka, in the WAIS Divide record, indicating a cooling trend after the ACR onset which is not clear in the DF or EDML series.

Within the range of uncertainty, the mean of our lead–lag estimates is consistent with the boundaries proposed by Pedro et al. (2012). Our results are consistent with those of Parrenin et al. (2013) for three out of the four change points addressed, but differ considerably at the T1 onset.

The considerable difference at T1 between our result and that of Parrenin et al. (2013) is most likely due to the much higher resolution of the WD CO₂ time series. It is also possible that the result of Parrenin et al. (2013) was limited to a local probability maximum of this change point in the CO₂ series. The addition of the WD paleotemperature record and removal of the Vostok record from AT3, the updated atmospheric CO₂ dataset, and our more generalized methodology are all, in part, responsible for the differences in computed time delays (the Supplement). This testifies to the importance of data resolution, methodological development, and chronological accuracy in the determination of leads and lags.

4 Conclusions

Our study is a follow-up of the studies by Pedro et al. (2012) and Parrenin et al. (2013) on the leads and lags between atmospheric CO₂ and Antarctic temperature during the last deglacial warming. We refine the results of these studies by using the high-resolution CO₂ record from WD; using Δ age computed on WD; using a new Antarctic temperature stack composed of 5 volcanically synchronized ice core isotope records, developed by Buizert et al. (2018); and using a more precise and complete probabilistic estimate to de-

termine change points. Our methodology detects four major common break points in both time series. The phasing between CO₂ and Antarctic climate is small but variable, with phasing ranging from a centennial-scale CO₂ lead, to synchrony, to a multicentennial-scale lead of Antarctic climate. This variability in phasing indicates that the mechanisms of coupling are complex. We propose three possibilities: (i) the mechanisms by which CO₂ and Antarctic temperature were coupled were consistent through the deglaciation, but can be modulated by external forcings or background conditions that impact heat transfer and oceanic circulation (and hence CO₂ release); (ii) these mechanisms can be modulated by internal feedbacks that change the response timings of the two series; and/or (iii) multiple, distinct mechanisms might have provoked similar responses in both series, but with accordingly different lags.

We also explore the hypothesis of regional differences in temperature change in Antarctica. Though the use of individual isotopic temperature records is complicated by influences other than regional temperature, including localized variations in source temperature and ice sheet elevation change we confirm that the deglacial temperature rise did not occur homogeneously across the Antarctic continent, with significant differences existing between the WAIS Divide and East Antarctic records at the onset of the termination and smaller potential differences occurring between the East Antarctic records, including a considerably later end of the deglacial warming in the Dome Fuji record.

Hypotheses of relationships between these events should now be reinvestigated with modeling studies. The relationship between CO₂ and Antarctic temperature on longer timescales and during other periods of rapid climate change is also of interest. Additional high-resolution West Antarctic paleotemperature records would allow for a robust investigation of regional differences between West Antarctica and East Antarctica and our analysis at the T1 end could be improved with continued high-resolution CO₂ measurements through the beginning of the Holocene. Finally, the continued measurement of high-resolution ice core CO₂ records is essential to understand the relationship between CO₂ and global and regional temperature during the last 800 000 years.

Code and data availability. The code and data series used in this article are available at <https://doi.org/10.5281/zenodo.1221165> (Chowdhry Beeman, 2019). The CO₂, AT3 and individual isotopic series, as well as the chronological uncertainty, are included as .txt files. The original CO₂ data from Marcott et al. (2014) are available at <https://www.ncdc.noaa.gov/paleo/study/18636>. The synchronization tie points and original AT3 stack are available in the Supplementary Data spreadsheet of Buizert et al. (2018) at <https://doi-org.insu.bib.cnrs.fr/10.1038/s41586-018-0727-5>.

Supplement. The supplement related to this article is available online at: <https://doi.org/10.5194/cp-15-913-2019-supplement>.

Author contributions. JCB and LG contributed equally to this manuscript, and should both be considered lead authors. LG, FP, DR and JCB designed the research project. JCB and FP designed the change point identification method. JCB and LG led the writing of the manuscript; and all authors contributed to analysis and discussion of the phasing and timings, with TJF, CB and EB in particular providing expertise on the ATSS3 stack, the WD CO₂ record and their climatic context; and DR and FP on the lead–lag problem and its history.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. We thank Michael Sigl, Jinhwa Shin, Emmanuel Witrant and Amaelle Landais for their support and great help discussing this work and Mirko Severi for his EDC data and support with the volcanic synchronization. This work is supported by the Fondation Ars et Cuttoli and by the LEFE IceChrono and CO2Role projects.

Financial support. This research has been supported by the Centre National de la Recherche Scientifique (grant LEFE IceChrono and LEFE CO2Role) and the Fondation Ars et Cuttoli (grant CO2Role).

Review statement. This paper was edited by Hubertus Fischer and reviewed by Joel Pedro and one anonymous referee.

References

- Anderson, R., Ali, S., Bradtmiller, L., Nielsen, S., Fleisher, M., Anderson, B., and Burckle, L.: Wind-driven upwelling in the Southern Ocean and the deglacial rise in atmospheric CO₂, *Science*, 323, 1443–1448, 2009.
- Anklin, M., Barnola, J.-M., Schwander, J., Stauffer, B., and Raynaud, D.: Processes affecting the CO₂ concentrations measured in Greenland ice, *Tellus B*, 47, 461–470, 1995.
- Barnola, J.-M., Pimienta, P., Raynaud, D., and Korotkevich, Y. S.: CO₂-climate relationship as deduced from the Vostok ice core: a re-examination based on new measurements and on a re-evaluation of the air dating, *Tellus B*, 43, 83–90, 1991.
- Bauska, T. K., Baggenstos, D., Brook, E. J., Mix, A. C., Marcott, S. A., Petrenko, V. V., Schaefer, H., Severinghaus, J. P., and Lee, J. E.: Carbon isotopes characterize rapid changes in atmospheric carbon dioxide during the last deglaciation, *P. Natl. Acad. Sci. USA*, 113, 3465–3470, 2016.
- Berger, A.: Long-term variations of daily insolation and Quaternary climatic changes, *J. Atmos. Sci.*, 35, 2362–2367, 1978.
- Buizert, C. and Severinghaus, J. P.: Dispersion in deep polar firn driven by synoptic-scale surface pressure variability, *The Cryosphere*, 10, 2099–2111, <https://doi.org/10.5194/tc-10-2099-2016>, 2016.
- Buizert, C., Cuffey, K. M., Severinghaus, J. P., Baggenstos, D., Fudge, T. J., Steig, E. J., Markle, B. R., Winstrup, M., Rhodes, R. H., Brook, E. J., Sowers, T. A., Clow, G. D., Cheng, H., Edwards, R. L., Sigl, M., McConnell, J. R., and Taylor, K. C.: The WAIS Divide deep ice core WD2014 chronology – Part 1: Methane synchronization (68–31 ka BP) and the gas age–ice age difference, *Clim. Past*, 11, 153–173, <https://doi.org/10.5194/cp-11-153-2015>, 2015.
- Buizert, C., Keisling, B. A., Box, J. E., He, F., Carlson, A. E., Sinclair, G., and DeConto, R. M.: Greenland-Wide Seasonal Temperatures During the Last Deglaciation, *Geophys. Res. Lett.*, 45, 1–10, <https://doi.org/10.1002/2017GL075601>, 2017.
- Buizert, C., Sigl, M., Severi, M., Markle, B. R., Wettstein, J. J., McConnell, J. R., Pedro, J. B., Sodemann, H., Goto-Azuma, K., Kawamura, K., et al.: Abrupt ice-age shifts in southern westerly winds and Antarctic climate forced from the north, *Nature*, 563, 681–685, 2018.
- Burke, A. and Robinson, L. F.: The Southern Ocean’s role in carbon exchange during the last deglaciation, *Science*, 335, 557–561, 2012.
- Caillon, N., Severinghaus, J. P., Jouzel, J., Barnola, J.-M., Kang, J., and Lipenkov, V. Y.: Timing of atmospheric CO₂ and Antarctic temperature changes across Termination III, *Science*, 299, 1728–1731, 2003.
- Chang, C. C. and Politis, D. N.: Robust autocorrelation estimation, *J. Comput. Graph. Stat.*, 25, 144–166, 2016.
- Chappellaz, J., Blunier, T., Kints, S., Dällenbach, A., Barnola, J.-M., Schwander, J., Raynaud, D., and Stauffer, B.: Changes in the atmospheric CH₄ gradient between Greenland and Antarctica during the Holocene, *J. Geophys. Res.-Atmos.*, 102, 15 987–15 997, 1997.
- Chowdhry Beeman, J.: Jai-Chowdhry/LinearFit-2.0-beta 2.0-CP (Version 2.0-CP), Zenodo, available at: <http://doi.org/10.5281/zenodo.2672547>, last access: 7 May 2019.
- Cuffey, K. M., Clow, G. D., Steig, E. J., Buizert, C., Fudge, T., Koutnik, M., Waddington, E. D., Alley, R. B., and Severinghaus, J. P.: Deglacial temperature history of West Antarctica, *P. Natl. Acad. Sci. USA*, 113, 14249–14254, 2016.
- Fischer, H., Behrens, M., Bock, M., Richter, U., Schmitt, J., Loulergue, L., Chappellaz, J., Spahni, R., Blunier, T., Leuenberger, M., and Stocker, T. F.: Changing boreal methane sources and constant biomass burning during the last termination, *Nature*, 452, 864–867, 2008.
- Foreman-Mackey, D., Hogg, D. W., Lang, D., and Goodman, J.: Emcee: The MCMC Hammer, *Publ. Astron. Soc. Pac.*, 125, 306–312, <https://doi.org/10.1086/670067>, 2013.
- Fujita, S., Parrenin, F., Severi, M., Motoyama, H., and Wolff, E. W.: Volcanic synchronization of Dome Fuji and Dome C Antarctic deep ice cores over the past 216 kyr, *Clim. Past*, 11, 1395–1416, <https://doi.org/10.5194/cp-11-1395-2015>, 2015.
- Ghosh, S.: Kernel smoothing: Principles, methods and applications, John Wiley & Sons, Hoboken, NJ, USA, 2018.
- Goodman, J. and Weare, J.: Ensemble samplers with affine invariance, *Comm. App. Math. Com. Sc.*, 5, 65–80, 2010.
- Hays, J. D., Imbrie, J., and Shackleton, N. J.: Variations in the Earth’s orbit: pacemaker of the ice ages, *Science*, 194, 1121–1132, 1976.

- Johnsen, S., Dansgaard, W., and White, J.: The origin of Arctic precipitation under present and glacial conditions, *Tellus B*, 41, 452–468, 1989.
- Jones, J. M., Gille, S. T., Goosse, H., Abram, N. J., Canziani, P. O., Charman, D. J., Clem, K. R., Crosta, X., De Lavergne, C., Eisenman, I., England, M. H., Fogt, R. L., Frankcombe, L. M., Marshall, G. J., Masson-Delmotte, V., Morrison, A. K., Orsi, A. J., Raphael, M. N., Renwick, J. A., Schneider, D. P., Simpkins, G. R., Steig, E. J., Stenni, B., Swingedouw, D., and Vance, T. R.: Assessing recent trends in high-latitude Southern Hemisphere surface climate, *Nat. Clim. Change*, 6, 91–926, 2016.
- Jouzel, J., Masson-Delmotte, V., Cattani, O., Dreyfus, G., Falourd, S., Hoffmann, G., Minster, B., Nouet, J., Barnola, J.-M., Chappellaz, J., Fischer, H., Gallet, J. C., Johnsen, S., Leuenberger, M., Loulergue, L., Luethi, D., Oerter, H., Parrenin, F., Raisbeck, G., Raynaud, D., Schilt, A., Schwander, J., Selmo, E., Souchez, R., Spahni, R., Stauffer, B., Steffensen, J. P., Stenni, B., Stocker, T. F., Tison, J. L., Werner, M., and Wolff, E. W.: Orbital and millennial Antarctic climate variability over the past 800,000 years, *Science*, 317, 793–796, 2007.
- Kawamura, K., Parrenin, F., Lisiecki, L., Uemura, R., Vimeux, F., Severinghaus, J. P., Hutterli, M. A., Nakazawa, T., Aoki, S., Jouzel, J., Raymo, M. E., Matsumoto, K., Nakata, H., Motoyama, H., Fujita, S., Goto-Azuma, K., Fujii, Y., and Watanabe, O.: Northern Hemisphere forcing of climatic cycles in Antarctica over the past 360,000 years, *Nature*, 448, 912–916, 2007.
- Landais, A., Capron, E., Masson-Delmotte, V., Toucanne, S., Rhodes, R., Popp, T., Vinther, B., Minster, B., and Prié, F.: Ice core evidence for decoupling between midlatitude atmospheric water cycle and Greenland temperature during the last deglaciation, *Clim. Past*, 14, 1405–1415, <https://doi.org/10.5194/cp-14-1405-2018>, 2018.
- Lisiecki, L. and Raymo, M.: A Plio-Pleistocene stack of 57 globally distributed benthic 0126-026180 records, *Paleoceanography*, 20, 1–17, 2005.
- Lorius, C. and Merlivat, L.: Distribution of mean surface stable isotopes values in East Antarctica; observed changes with depth in coastal area, *Tech. rep.*, CEA Centre d'Etudes Nucleaires de Saclay, Gif-Sur-Yvette, France, 1975.
- Loulergue, L., Parrenin, F., Blunier, T., Barnola, J.-M., Spahni, R., Schilt, A., Raisbeck, G., and Chappellaz, J.: New constraints on the gas age-ice age difference along the EPICA ice cores, 0–50 kyr, *Clim. Past*, 3, 527–540, <https://doi.org/10.5194/cp-3-527-2007>, 2007.
- Lüthi, D., Le Floch, M., Bereiter, B., Blunier, T., Barnola, J.-M., Siegenthaler, U., Raynaud, D., Jouzel, J., Fischer, H., Kawamura, K., and Stocker, T. F.: High-resolution carbon dioxide concentration record 650,000–800,000 years before present, *Nature*, 453, 379–382, 2008.
- Marcott, S. A., Bauska, T. K., Buizert, C., Steig, E. J., Rosen, J. L., Cuffey, K. M., Fudge, T., Severinghaus, J. P., Ahn, J., Kalk, M. L., McConnell, J. R., Sowers, T., Taylor, K. C., White, J. W. C., and Brook, E. J.: Centennial-scale changes in the global carbon cycle during the last deglaciation, *Nature*, 514, 616–619, 2014.
- Masson-Delmotte, V., Schulz, M., Abe-Ouchi, A., Beer, J., Ganopolski, A., González Rouco, J., Jansen, E., Lambeck, K., Luterbacher, J., Naish, T., Osborn, T., Otto-Bliesner, B., Quinn, T., Ramesh, R., Rojas, M., Shao, X., and Timmermann, A.: Information from Paleoclimate Archives, in: *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by: Stocker, T., Qin, D., Plattner, G.-K., Tignor, M., Allen, S., Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P., book section 5, 383–464, Cambridge University Press, Cambridge, UK and New York, NY, USA, <https://doi.org/10.1017/CBO9781107415324.013>, 2013.
- McConnell, J. R., Burke, A., Dunbar, N. W., Köhler, P., Thomas, J. L., Arienzo, M. M., Chellman, N. J., Maselli, O. J., Sigl, M., Adkins, J. F., Baggenstos, D., Burkhart, J. F., Brook, E. J., Buizert, C., Cole-Dai, J., Fudge, T. J., Knorr, G., Graf, H.-F., Griesman, M. M., Iverson, N., McGwire, K. C., Mulvaney, R., Paris, G., Rhodes, R. H., Saltzman, E. S., Severinghaus, J. P., Steffensen, J. P., Taylor, K. C., and Winckler, G.: Synchronous volcanic eruptions and abrupt climate change ~ 17.7 ka plausibly linked by stratospheric ozone depletion, *P. Natl. Acad. Sci. USA*, 114, 10035–10040, 2017.
- Monnin, E., Indermühle, A., Dällenbach, A., Flückiger, J., Stauffer, B., Stocker, T. F., Raynaud, D., and Barnola, J.-M.: Atmospheric CO₂ concentrations over the last glacial termination, *Science*, 291, 112–114, 2001.
- Mudelsee, M.: TAUEST: A computer program for estimating persistence in unevenly spaced weather/climate time series, *Comput. Geosci.*, 28, 69–72, 2002.
- NorthGRIP Project Members: High-resolution record of Northern Hemisphere climate extending into the last interglacial period, *Nature*, 431, 147–151, 2004.
- Parrenin, F., Barker, S., Blunier, T., Chappellaz, J., Jouzel, J., Landais, A., Masson-Delmotte, V., Schwander, J., and Veres, D.: On the gas-ice depth difference (Δ depth) along the EPICA Dome C ice core, *Clim. Past*, 8, 1239–1255, <https://doi.org/10.5194/cp-8-1239-2012>, 2012.
- Parrenin, F., Masson-Delmotte, V., Köhler, P., Raynaud, D., Pailard, D., Schwander, J., Barbante, C., Landais, A., Wegner, A., and Jouzel, J.: Synchronous change of atmospheric CO₂ and Antarctic temperature during the last deglacial warming, *Science*, 339, 1060–1063, 2013.
- Parrenin, F., Bazin, L., Capron, E., Landais, A., Lemieux-Dudon, B., and Masson-Delmotte, V.: IceChrono1: a probabilistic model to compute a common and optimal chronology for several ice cores, *Geosci. Model Dev.*, 8, 1473–1492, <https://doi.org/10.5194/gmd-8-1473-2015>, 2015.
- Pedro, J. B., Rasmussen, S. O., and van Ommen, T. D.: Tightened constraints on the time-lag between Antarctic temperature and CO₂ during the last deglaciation, *Clim. Past*, 8, 1213–1221, <https://doi.org/10.5194/cp-8-1213-2012>, 2012.
- Pedro, J. B., Bostock, H. C., Bitz, C. M., He, F., Vandergoes, M. J., Steig, E. J., Chase, B. M., Krause, C. E., Rasmussen, S. O., Markle, B. R., and Cortese, G.: The spatial extent and dynamics of the Antarctic Cold Reversal, *Nat. Geosci.*, 9, 51, 2016.
- Pedro, J. B., Jochum, M., Buizert, C., He, F., Barker, S., and Rasmussen, S. O.: Beyond the bipolar seesaw: Toward a process understanding of interhemispheric coupling, *Quaternary Sci. Rev.*, 192, 27–46, 2018.
- Petrenko, V. V., Smith, A. M., Brook, E. J., Lowe, D., Riedel, K., Brailsford, G., Hua, Q., Schaefer, H., Reeh, N., Weiss, R. F., Etheridge, D., and Severinghaus, J. P.: ¹⁴CH₄ measurements

- in Greenland ice: investigating last glacial termination CH₄ sources, *Science*, 324, 506–508, 2009.
- Rae, J., Burke, A., Robinson, L., Adkins, J., Chen, T., Cole, C., Greenop, R., Li, T., Littley, E., Nita, D., Stewart J. A., and Taylor, B. J.: CO₂ storage and release in the deep Southern Ocean on millennial to centennial timescales, *Nature*, 562, 569–573, 2018.
- Raynaud, D. and Siegenthaler, U.: Role of trace gases: the problem of lead and lag, *Global changes in the perspective of the past*, edited by: Eddy, J. A., and Oeschger, H., Wiley, Chichester, UK, 173–188, 1993.
- Raynaud, D., Jouzel, J., Barnola, J., Chappellaz, J., Delmas, R., and Lorius, C.: The ice record of greenhouse gases, *Science*, 259, 926–926, 1993.
- Robinson, P.: Estimation of a time series model from unequally spaced data, *Stoch. Proc. Appl.*, 6, 9–24, 1977.
- Schmitt, J., Schneider, R., Elsig, J., Leuenberger, D., Lourantou, A., Chappellaz, J., Köhler, P., Joos, F., Stocker, T. F., Leuenberger, M., and Fischer, H.: Carbon isotope constraints on the deglacial CO₂ rise from ice cores, *Science*, 336, 711–714, 2012.
- Schwarz, G.: Estimating the dimension of a model, *Ann. Stat.*, 6, 461–464, 1978.
- Shakun, J. D., Clark, P. U., He, F., Marcott, S. A., Mix, A. C., Liu, Z., Otto-Bliesner, B., Schmittner, A., and Bard, E.: Global warming preceded by increasing carbon dioxide concentrations during the last deglaciation, *Nature*, 484, 49–54, 2012.
- Silverman, B. W.: *Density estimation for statistics and data analysis*, Chapman and Hall, London, UK and New York, USA, 1986.
- Skinner, L., Fallon, S., Waelbroeck, C., Michel, E., and Barker, S.: Ventilation of the deep Southern Ocean and deglacial CO₂ rise, *Science*, 328, 1147–1151, 2010.
- Sowers, T., Bender, M., Raynaud, D., and Korotkevich, Y. S.: $\delta^{15}\text{N}$ of N₂ in air trapped in polar ice: A tracer of gas transport in the firm and a possible constraint on ice age-gas age differences, *J. Geophys. Res.-Atmos.*, 97, 15683–1697, 1992.
- Stocker, T. F. and Johnsen, S. J.: A minimum thermodynamic model for the bipolar seesaw, *Paleoceanography*, 18, 1087, <https://doi.org/10.1029/2003PA000920>, 2003.
- Tarantola, A.: *Inverse problem theory*, SIAM, Philadelphia, USA, 2005.
- WAIS Divide Project Members: Onset of deglacial warming in West Antarctica driven by local orbital forcing, *Nature*, 500, 440–444, 2013.
- Williams, D., Peck, J., Karabanov, E., Prokopenko, A., Kravchinsky, V., King, J., and Kuzmin, M.: Lake Baikal record of continental climate response to orbital insolation during the past 5 million years, *Science*, 278, 1114–1117, 1997.

Chapter 4

Methane Synchronization

This article is in preparation for submission to *Geoscientific Model Development*.

Automated and probabilistic ice core synchronisation in IceChrono 2.0.

Jai Chowdhry Beeman¹, Frédéric Parrenin¹, Emmanuel Witrant², Amaëlle Landais³, Xavier Fain¹, Benoît Tournadre⁴, and Robert Mulvaney⁵

¹Université Grenoble Alpes – CNRS – IRD – Institut des Géosciences de l’Environnement, F-38000 Grenoble, France

²Université Grenoble Alpes – CNRS – Grenoble Images Parole Signal Automatique, F-38000 Grenoble, France

³Laboratoire des Sciences du Climat et de l’Environnement, CEA-CNRS-UVSQ, Gif sur Yvette, France

⁴MINES ParisTech, PSL Research University, O.I.E. - Centre Observation, Impacts, Energy, 06904 Sophia Antipolis, France

⁵British Antarctic Survey – NERC – Cambridge, UK

Abstract. Accurate chronologies allow us to access the wealth of paleoclimate information in ice cores. Synchronization using time series measurements—to other ice cores, to paleoclimate records such as speleothems or marine sediment cores, or to orbital targets—is a powerful tool for improving ice core chronologies. Here, as part of the IceChrono dating model (Parrenin et al., 2015), we propose an automated method for synchronization. The method minimizes residuals between time series, taking chronological observations and prior information into account, and a newly included parallelized Markov Chain Monte Carlo method evaluates the posterior probability of chronologies, including the synchronization term. We make two tests of our method using CH₄ time series. For the first test, we synchronize the TALOS Dome and EPICA Dome C ice cores in the air phase, including both cores in the dating experiment. For the second test, we synchronize the Fletcher Promontory and WAIS Divide ice cores in the air phase, treating WAIS Divide, which is well dated, as an external record. The method is able to accurately synchronize the gas records in the two tests and provides an objective, continuous estimate of the synchronization uncertainty.

1 Introduction

Deep polar ice cores record past atmospheric concentrations in air bubbles trapped in the ice, alongside climate proxies in both the ice and air phases. The relationships between these two archives are most accurately understood when they are dated with realistic estimates of uncertainty. Synchronization, with other ice cores, other paleoclimate records, or astronomical targets, is a valuable source of information in the construction of ice core chronologies. In the ice phase of ice cores, volcanic events recorded as acidity peaks are often used to align ice cores located within the range of atmospheric transport of volcanic material (Vinther et al., 2006; Severi et al., 2007; Rasmussen et al., 2008; Severi et al., 2012; Parrenin et al., 2012; Svensson et al., 2013; Fujita et al., 2015; Buizert et al., 2018). Ice cores have also been synchronized in the ice phase using Beryllium 10 (Raisbeck et al., 2006). To synchronize the air trapped in ice cores, well-mixed trace gases like CO₂ and CH₄ are commonly used (Monnin et al., 2004; Blunier et al., 2007; Buiron et al., 2011; Schüpbach et al., 2011), with CH₄ now preferred because of the rapid transitions present in CH₄ series, which are relatively easy to match by eye; atmospheric $\delta^{18}\text{O}$, though less well-mixed than

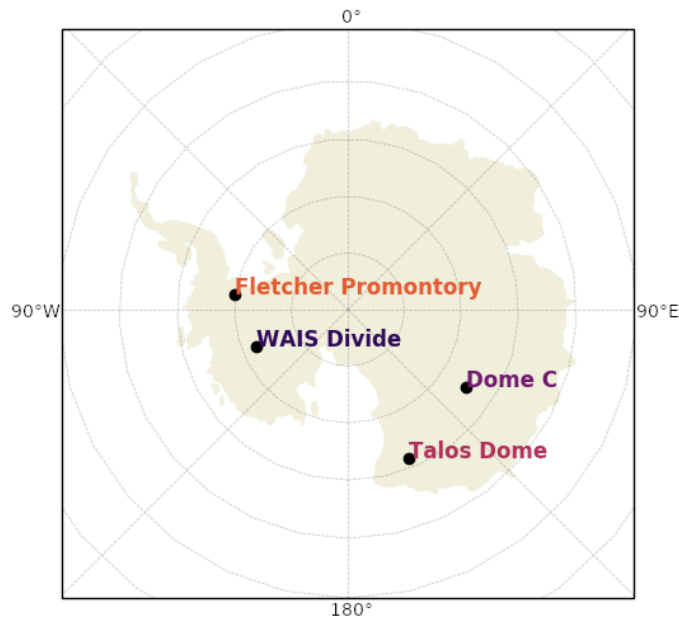


Figure 1. Drilling site locations of the Fletcher Promontory (-82.605° , -77.901667°), WAIS Divide (-112.085° , -79.467°), EPICA Dome C (123.332° , -75.010°) and Talos Dome (159.183° , -72.817°) ice cores on the Antarctic Continent.

CH₄ or CO₂, has also been used as it records rapid transitions accurately (Capron et al., 2010). Orbital tuning – matching proxy series to the earth’s astronomical parameters, which are well-known far into the past – has been used to date many deep ice cores, and is indeed a form of external synchronization. These parameters are recorded in ice cores by proxies such as the O₂/N₂ ratio (Bender, 2002; Suwa and Bender, 2008; Lipenkov et al., 2011; Landais et al., 2012), Total Air Content (Raynaud et al., 2007; Lipenkov et al., 2011) and/or $\delta^{18}\text{O}_{atm}$ (Petit et al., 1999; Dreyfus et al., 2007). Ice cores can also be synchronized to well-dated external paleoclimate archives, such as speleothems (i.e. Buizert et al. (2015); Extier et al. (2018)).

Synchronization between paleoclimate records is usually done manually. Two strategies are common: matching peaks in the time series, and matching slope midpoints during fast transitions. Each of these methods attempts, visually, to minimize the distance between the two records. Often, an effort is made to change the individual ice core chronologies as minimally as possible. However, visually synchronized chronologies can be method-subjective and difficult to reproduce (Martinson et al., 1987), as is the case for approximate visual dating methods more generally (Winstrup et al., 2012).

Several previous studies have developed and tested automatic or semi-automatic methods for the synchronization of paleoclimate records (Lisiecki and Lisiecki, 2002; Raynaud et al., 2007; Blunier et al., 2007). The Match algorithm of Lisiecki and Lisiecki (2002) is widely used in the field of paleoceanography, and the HMM-Match algorithm developed by Lin et al. (2014) is of particular relevance to this study, as it also calculates the probability of synchronization scenarios, between marine sediment cores, allowing for the definition of a sedimentation prior.

In IceChrono 2.0, we build upon the ice core dating model IceChrono 1.0 (Parrenin et al., 2015) to develop a method adapted to the synchronization of ice cores, and to the synchronization of paleoclimate archives in general. Like its predecessors IceChrono 1.0 and Datice (Lemieux-Dudon et al., 2010), IceChrono 2.0 is built on the Bayesian inversion of three glaciological variables, and allows for the inclusion of discrete chronological information. It is developed to be adaptable to a variety of optimization methods, and thus adaptable to a variety of formulations of the chronological inverse problem. Importantly, within the framework of calculating a probabilistic chronology, we can synchronize a core to an external target, or two (or more) cores with each other.

As a test, we apply our method to the gas phase synchronization of two deep Antarctic ice cores: the European Project for Ice Coring in Antarctica Dome C (EPICA Dome C, or EDC) ice core, (EPICA Community Members, 2004) and the Talos Dome Ice Core (TALDICE, <http://www.taldice.org>) (Figure 1). In the gas phase, the atmospheric lifetime of methane (CH_4), on the order of a decade, makes it an ideal candidate for synchronization. The lifetime of methane is short enough for concentrations to respond to and record rapid changes in global climate, but long enough so that atmospheric pole-to-pole transport times, on the order of a year, do not add large geographic variability to the records (Blunier et al., 2007). Loulergue et al. (2008), Buiron et al. (2011) and Schüpbach et al. (2011) performed manual synchronizations using methane measurements as part of the TALDICE 1 chronology, and methane measurements of appropriate resolution are available for both ice cores up to ~ 45 ka BP (Before Present, defined as years before 1950). These two cores have also been dated as part of the Antarctic Ice Core Chronology 2012 (AICC2012; Bazin et al. (2013); Veres et al. (2013)). AICC2012 was calculated using the DATICE model (Lemieux-Dudon et al., 2010), following a similar Bayesian framework to IceChrono (though coded differently, and using an analytical gradient). In DATICE, as in IceChrono 1.0, synchronization is treated by defining discrete stratigraphic links between the two records.

Gas records in ice cores are smoothed by the air trapping process in the firn column, which is controlled largely by the accumulation rate at the drilling site. While the accumulation at the TALDICE site was most likely about three times as high as the accumulation at EDC during the Holocene, this difference is relatively small when compared to much higher accumulation cores like WAIS Divide, where accumulation is as much as an order of magnitude higher than at EDC (Buizert et al., 2015; Sigl et al., 2016). The correct treatment of gas trapping is out of the scope of this paper, but we consider it more prudent to compare two cores with closer accumulation rates.

2 IceChrono: making an ice core age model

The IceChrono dating model (Parrenin et al., 2015) constructs ice and gas ages for ice cores by inverting three glaciological variables, which we refer to as the forward model. The formation of a polar ice sheet, ice cap or glacier begins when snow accumulates on its surface – in IceChrono, this process is represented by the accumulation vector, denoted as a , expressed in $\text{m}\cdot\text{yr}^{-1}$ ice equivalent. As the snow densifies into firn and then ice, air that circulates through the firn is trapped in bubbles in the ice at the Lock-in Depth (LID), denoted as l , expressed in m. Finally, as ice flows at the site of a drilling, it effectively thins vertically – this is represented in IceChrono by an ice thinning function, denoted as τ , which is unitless.

These three glaciological variables allow us to build an ice core chronology for each core k at depth z . First, we calculate the ice age χ_k , by integrating the number of accumulated and thinned ice layers above depth z :

$$\chi_k(z_k) = \int_0^{z_k} \frac{D_k(z'_k)}{a_k(z'_k)\tau_k(z'_k)} dz'_k \quad (1)$$

where D_k is the relative density with respect to pure ice at a given depth in the core. Then, we calculate the air age Ψ :

$$\Psi_k(z_k) = \chi_k(z_k - \Delta d(z_k)), \quad (2)$$

where $\Delta d(z_k)$, or Δ depth, is the depth difference between the air trapped at a given depth and the ice with the same age, which is located shallower in the core. $\Delta d(z_k)$ is derived from the Lock-in Depth l using the following relationships:

$$\int_{z_k - \Delta d_k(z_k)}^{z_k} \frac{D(z'_k)}{\tau(z'_k)} dz'_k \approx \int_0^{l_k(z_k) D_k^{firn}(z_k)} \frac{1}{\tau_k(z_k^{ie})} dz_k^{ie}, \quad (3)$$

$$z_k^{ie} = \int_0^{z_k} D_k(z'_k) dz'_k \quad (4)$$

10 where $D_k^{firn}(z_k)$ is the firn density when the ice at depth z_k was at lock-in depth, and z_k^{ie} is the ice equivalent depth.

We define a general relationship $age(z)$ for the chronology as a whole. $age(z^{ice})$ takes the value of $\chi(z)$ in the ice matrix, and $age(z^{air})$ the value of $\Psi(z)$ for air bubbles. Similarly, $age(z_k)$ indicates the chronology for core k , and $age(z_j)$ indicates the chronology for core j . This notation allows us to express the synchronization residuals simply, without separate equations for ice or air ages. The correct age for the synchronization is determined by the species used (i.e. air age for CH_4).

15 3 Defining a cost function

In the chronological inverse problem, we estimate the posterior probability distributions of the three glaciological variables, which we refer to as forward model variables, using the available chronological information. We therefore formulate the problem in a Bayesian inverse framework.

20 Before we perform the inversion, we compile the available chronological observations, to which we associate the probability density ρ_D . We also specify a prior probability density for the three glaciological variables (denoted together as \mathbf{X}), based roughly on our overall knowledge from previous modeling and experimental studies, which we denote ρ_M . Then, Bayes'

theorem allows us to calculate the overall state of information about a set of model parameters: the posterior probability density σ_M .

$$\sigma_M(\mathbf{X}) = \kappa \rho_M(\mathbf{X}) \rho_D(f(\mathbf{X})) \quad (5)$$

where κ is a normalizing constant and f is the forward relationship between the model parameters and the data. In our case, we may have data in units of age, or the concentration of a particular chemical species measured at a given depth in the core, for example. For data in units of age, f is equivalent to $age(z)$; for data in concentrations (for synchronization in particular), we must additionally associate a concentration with each value of $age(z)$. Assuming that the uncertainties associated with the observation values and the model prior are Gaussian (importantly, we have made no such assumptions about f itself), we can rewrite equation 5 as:

$$\sigma_M(\mathbf{X}) = \kappa \frac{1}{2} \exp(-J(\mathbf{X})), \quad (6)$$

where J is a cost function that sums residual terms with respect to the model priors and chronological information.

The cost function we develop for IceChrono 2.0 includes residual terms for each core k and core pair k, m that represent each source of available chronological information. We can group this information into three classes:

- residuals J^{prior} with respect to the forward variable priors
- residuals J^{obs} for discrete chronological observations, e.g. dated horizons
- and finally, residuals J^{sync} for time series synchronization (new to IceChrono 2.0).

3.1 Forward variable residuals

To calculate J^{prior} , we first define a correction function for each glaciological variable. In the case of accumulation, we write this function as

$$a_{corr} = \ln\left(\frac{a}{a_{prior}}\right); \quad (7)$$

the correction functions for τ and l are written similarly. These correction functions allow us to transform these three variables from Jeffreys variables, which are uniformly positive and described by log-normal probability distributions, to Cartesian variables, which can be described by Gaussian probability distributions (Tarantola, 2005). This allows us to define a Gaussian uncertainty (i.e. standard deviation) for the prior value of each of the three variables at a given depth.

The cost function and residual terms for accumulation for each core k are written as

$$J_k^a = (R_k^a)^T C_a^{-1} R_k^a \quad (8)$$

$$R_k^a = \frac{a_{corr}}{\sigma_k^a}. \quad (9)$$

σ_k^a is the uncertainty with respect to accumulation, and C_a is a matrix that accounts for the autocovariance of the accumulation residuals, if any is expected. The residuals for τ and l take the same form (all three are treated in greater detail in the description paper for IceChrono 1.0, Parrenin et al. (2015)). Finally, the sum of the residuals J_k^a , J_k^τ and J_k^l gives us the first residual class for ice core k , J_k^{prior} . Note that residuals with respect to the priors are defined only for individual cores, and not for core pairs.

3.2 Discrete chronological observation residuals

The second class of residuals, discrete observations or J^{obs} , is defined both for individual cores and for core pairs. These discrete observations can be dated ice or air horizons, dated ice or air intervals, stratigraphic links between two cores (in the ice or gas phases), or any combination of these. These residuals are defined individually in the appendix of Parrenin et al. (2015), and generally take the following form, for individual ice core k :

$$J_k^{obs} = (R_k^{obs})^T C_{obs}^{-1} R_k^{obs} \quad (10)$$

$$R_k^{obs} = \frac{age(z_k^{obs}) - age_k^{obs}}{\sigma_k^{obs}}. \quad (11)$$

As above, C_{obs} is the residual autocorrelation matrix, age_k^{obs} is the observed age, $chron(z_k^{obs})$ is the modeled age at the depth of the observation, and age_k^{obs} is the age uncertainty with respect to the observations for k . For stratigraphic links between ice cores k and m , we similarly write:

$$J_{k,m}^{obs} = (R_{k,m}^{obs})^T C_{obs}^{-1} R_{k,m}^{obs} \quad (12)$$

$$R_{k,m}^{obs} = \frac{age(z_k^{obs}) - age(z_m^{obs})}{\sigma_{k,m}^{obs}} \quad (13)$$

3.3 Synchronization residuals

The third class of residuals is used to synchronize ice cores using measurement series. This class of residuals is new to IceChrono 2.0; the residual classes described above were included in IceChrono 1 as well. Within IceChrono 2.0, two types of synchronization can be performed. The first is alignment to an external target, whose chronology does not change. The second is synchronization with another core also included in the chronological inverse problem, in which case the chronologies of both cores are allowed to vary.

3.3.1 Synchronization to an external target

For alignment to an external target, a target function is defined before the inversion. This function should take ages as an argument, and produce an output with the same units as the time series we wish to use to synchronize. This could mean using the same type of measurement (synchronizing methane series, for example), or different measurements, transformed to best match the synchronization series.

We denote the series of observations of each synchronization proxy p as obs_p . The depths corresponding to these observations are denoted z_i^{obs} . For proxy p , we denote the target function as $tar_p(age)$. At each optimization step, we evaluate the residuals between the observations at depths z_i^{obs} and the target. At each step, the proposed chronology $age(z)$ allows us to evaluate the value of the target function at the ages assigned to the depths z_i^{obs} :

$$\mathbf{R}_{p,i}^{sync} = \left(\frac{obs_{p,i} - tar_{p,i}(age(z_{p,i,k}^{obs}))}{\sigma_{p,i}^{sync}} \right) \quad (14)$$

Here, $\sigma_{p,i}^{sync}$ accounts for measurement and method uncertainty. To transform the residual vector into a residual scalar, we take its product with itself, transformed by \mathbf{C}_p^{sync} , the residual autocorrelation matrix defined for proxy p .

$$J_p^{sync} = (\mathbf{R}_p^{sync})^T (\mathbf{C}_p^{sync})^{-1} \mathbf{R}_p^{sync}; \quad (15)$$

3.3.2 Synchronization between two cores

In the case of synchronization between two cores i and j that are both included in the chronological inverse problem, we can no longer define a fixed target. Instead, the target for each series is the equivalent series on the other core, adjusted for the proposed chronology at each optimization step. This means that for core i , the target function $tar_{p,i}(age)$ changes at each step according to core j , and vice versa. This functionality is tested in the synchronization between Talos Dome and EDC.

3.4 The cost function

We sum the prior and discrete observation residual terms for each core k , the discrete observation residual terms for each core pair k, m , and the synchronization residuals for each method p to arrive at the total cost function:

$$J = \sum_k (J_k^{obs} + J_k^a + J_k^T + J_k^l) + \sum_{k,m} J_{k,m}^{obs} + \sum_p J_p^{sync}$$

With the cost function defined, we can now evaluate the likelihood

$$L = \exp(-J(\mathbf{X})), \quad (16)$$

which is proportional to the posterior probability density as given in equation 6. Constructing a likelihood function is sufficient to apply many Bayesian solution algorithms.

4 Exploring the posterior probability

Many of the time series used to synchronize paleoclimate records, like Methane or Sulfate time series, tend to be jagged in character. Additionally, the pairs of peaks that should be matched are not always unique—sometimes, a match can be ambiguous, with a peak or abrupt change in one series potentially matching two peaks in the other series. In practice, this means that the target functions f^{tar} used for synchronization is often highly nonlinear, and the posterior probability density will most likely be non-Gaussian and potentially multi-modal. The synchronization term in the cost function, J^{sync} , may have many local minima, where the values of the two series are locally close, and shifting one of the series by a small amount worsens the alignment (though the alignment could eventually be improved by a large shift). As such, convex, gradient-based optimization methods, like the Levenberg-Marquardt algorithm (Levenberg, 1944; Marquardt, 1963) used in IceChrono 1.0, are not necessarily applicable to methane synchronization, as they use information about the local gradient of the cost function to reach a minimum. In some cases, these methods may be useful to find local minima. To explore the posterior probability density, though, a stochastic (Monte-Carlo) sampling method that searches beyond local minima is necessary.

Even when using a stochastic sampling method, we cannot entirely escape the challenge of multimodality. Many Markov Chain Monte Carlo (MCMC) sampling methods like the Metropolis-Hastings algorithm (Metropolis et al., 1953; Hastings, 1970) can become stuck in local minima (Sharma, 2017), greatly increasing the number of iterations necessary to produce independent samples of the posterior.

We face a second challenge as well—that of dimensionality. Because stochastic algorithms require repeated model realizations and cost function evaluations, the required computational time generally increases with both the number of parameters and number of observation. Thus, we not only need a sampling method capable of dealing with multimodality of the posterior, but also with high-dimensional parameter spaces.

We design a solution algorithm for this application based on the Python emcee package (Foreman-Mackey et al., 2013). Emcee was originally written as a Python implementation of the affine-invariant ensemble Metropolis-Hastings algorithm developed by Goodman and Weare (2010). This algorithm is parallelizable, making it useful for the application of stochastic optimization to large-scale problems. However, the proposal steps in the Goodman and Weare (2010) algorithm are not well-suited to multimodal posterior distributions (Sharma, 2017). Thus, we implement Differential Evolution Markov Chain (DEMC) proposals, as in ter Braak (2006) and ter Braak and Vrugt (2008), leaving out the adaptive component, that are included in the Foreman-Mackey et al. (2013) package. We tune the proposals in order to adapt to the multimodality of our problem. The algorithm is parallelized using the Schwimmbad module (Price-Whelan and Foreman-Mackey, 2017).

In DEMC, N chains evolve together, allowing us to maintain a parallel structure. For a given chain, the proposal at each step is determined using the positions of the remaining $N - 1$ chains. These proposals are accepted or rejected using the Metropolis-Hastings criterion, and each chain thus becomes a Markov chain whose target distribution is the posterior. In ter Braak and

Vrugt (2008), two proposal steps are proposed: a standard DE-MC move, and a snooker move. These moves are described below; for further detail with respect to the algorithm we refer the reader to ter Braak and Vrugt (2008). Similar combinations of proposals are used in many studies using algorithms based on ter Braak and Vrugt (2008)—see Laloy and Vrugt (2012) and Nelson et al. (2013), for example.

5 4.1 Differential Evolution Move

Consider a population of N parallel markov chains with current positions \mathbf{x}^j at timestep j , forming a set \mathbf{X}^j . To update a given walker \mathbf{x}_i , randomly select the current positions of two additional walkers \mathbf{x}_{R1} and \mathbf{x}_{R2} from the complimentary set $\mathbf{X}_{\neq i}^j$ which includes all members of \mathbf{X}^j except \mathbf{x}_i^j . The proposal \mathbf{x}^* is generated using:

$$\mathbf{x}^* = \mathbf{x}_i^j + \gamma(\mathbf{x}_{R1} - \mathbf{x}_{R2}) + \mathbf{e} \quad (17)$$

10 where γ is a scalar, and \mathbf{e} is a stochastic component drawn from a symmetric distribution with small variance with respect to the posterior. Using this proposal distribution, each chain respects detailed balance with respect to the posterior distribution if we accept or reject each jump according to the Metropolis-Hastings ratio:

$$P_{\mathbf{x}_i^j \rightarrow \mathbf{x}_i^{j+1} = \mathbf{x}^*} = \min \left\{ 1, \frac{L(\mathbf{x}^*)}{L(\mathbf{x}_i^j)} \right\} \quad (18)$$

15 This move is nearly parallel to the line between \mathbf{x}_{R1} and \mathbf{x}_{R2} . As the move uses walkers from the current generation, it adapts to the scale and the orientation of the posterior distribution as the walkers converge toward the posterior.

For our synchronization test, we use two variants of the classical Differential Evolution move. For the first, we take the value of γ to be $2.38 / \sqrt{2d}$, where d is the number of model parameters, as suggested by ter Braak and Vrugt (2008) to provide an approximate acceptance probability of around 0.23, a heuristic reference for problems with large parameter dimensionality. In the emcee framework, multiple moves are randomly selected with user-specified probability, which we assign by testing for
20 optimal acceptance probability. We assign this move a selection probability of 0.5 – that is, it should be applied for about half of all proposals.

For the second variant of the classical Differential Evolution move, we take γ to be 0.98. This value reflects the probable multimodality of the synchronization posterior – ter Braak and Vrugt (2008) suggest that a value close to 1 allows a walker in one mode to easily transit to another. Kupiainen-Määttä (2016) uses a value of 0.98, which we find to slightly improve
25 acceptance probability in our tests. We assign this move a selection probability of 0.1 – that is, it should be applied for about 10 % of proposals. This probability could be adjusted according to the degree of multimodality of the synchronization at hand.

4.2 Snooker Move

For the remaining ~ 40 % of proposals, we use the DEMC snooker update proposed by ter Braak and Vrugt (2008). To update walker \mathbf{x}_i using this algorithm, three other chains, with current states \mathbf{z} , \mathbf{z}_{R1} and \mathbf{z}_{R2} are selected. Then, we orthogonally

project the vectors \mathbf{z}_{R1} and \mathbf{z}_{R2} onto the vector $\mathbf{x}_i - \mathbf{z}$, giving two projected vectors \mathbf{z}_{P1} and \mathbf{z}_{P2} . Then, the proposal is calculated as:

$$\mathbf{x}^* = \mathbf{x}_i^j + \gamma(\mathbf{z}_{P1} - \mathbf{z}_{P2}) \quad (19)$$

To maintain detailed balance, the Metropolis-Hastings ratio when this move is applied is modified to:

$$P_{\mathbf{x}_i^j \rightarrow \mathbf{x}_i^{j+1} = \mathbf{x}^*} = \min \left\{ 1, \frac{L(\mathbf{x}^*) \|\mathbf{x}^* - \mathbf{z}\|^{d-1}}{L(\mathbf{x}_i^j) \|\mathbf{x}_i^j - \mathbf{z}\|^{d-1}} \right\} \quad (20)$$

where d is the dimensionality of the model vector.

The orthogonal projection is designed to account for the covariance between model parameters, allowing them to be updated coherently, and effectively reducing the model dimensionality. Thus, ter Braak and Vrugt (2008) propose $\gamma = 2.38/\sqrt{2}$ for this move, considering the dimensionality to be effectively 1. In our tests, however, we found that $\gamma = 2.38/\sqrt{2d}$ led to acceptance rates much closer to the reference value, and lower autocorrelation times for the model variables. This may be due to the jagged, multimodal nature of our posterior probability distribution. The local covariance of model parameters, within individual modes, is possibly stronger and structurally different than the covariance across modes. Thus, this step is likely more appropriate for smaller, local moves, in our particular case. However, we retain the orthogonal projection, as it should still allow us to explore local modes efficiently.

4.3 Alternative methods for exploring the posterior

The parameters of the MCMC steps given above, the MCMC steps selected, and even the nature of the solution method used are problem-dependent. In the case of a synchronization problem thought to be close to unimodal – if we have a considerable amount of discrete dated points and intervals, for example, a simpler MCMC step might be preferable. Without the synchronization module, convex optimization is still ideal, given that it takes considerably less computational effort.

IceChrono is written in a modular, object-oriented Python framework. In practice, this means that it is not difficult to implement various solution methods. The Levenberg-Marquardt gradient-based optimization method is carried over from IceChrono 1.0. The basinhopping method of Wales and Doye (1997), which combines stochastic sampling and gradient-based optimization, is implemented. The additional moves included in the emcee package are also easily available, including the classic Metropolis-Hastings algorithm and the original Goodman and Weare (2010) move. The python Kombine package, an algorithm based on a kernel-density estimate (Farr and Farr, 2015) is included, as is a version of Parallel Tempering in the ptemcee package (Vousden et al., 2015). The latter two are likely useful as well for multimodal distributions. Thorough tests of these algorithms are outside the scope of this paper, given the time and computational resources necessary to apply a single algorithm, but we implement these solution methods to make IceChrono more broadly applicable for future studies.




Figure 2. Prior, age (yr BP) vs. CH₄ (ppmv) CH₄ series on the AICC2012 prior scenarios for the EDC (black) and Talos Dome (blue) ice core records. The cumulative tuning uncertainty is shown using the brown shading on the EDC series.

4.4 Representing the posterior probability

This probability density is best represented as the ensemble of all model vectors accepted in the MCMC simulation. In IceChrono 2.0, we save the age-depth relationships that correspond to these vectors in the HDF5 format (Folk et al., 2011) as implemented in the Python h5py package by Collette (2013) and the emcee package by Foreman-Mackey et al. (2013). The most accurate chronology produced by our method is indeed this ensemble of model vectors, rather than a mean chronology with covariances, as would be the case for IceChrono 1.0. The ensemble chronology can be represented as an animation, as indicated by Tarantola (2005). Alternatively, we can make histograms of individual variables or multiple-line plots of the series of glaciological variables or age-depth relationships.

5 Test case: synchronization of EDC and Talos Dome ice cores using CH₄

5.1 Setting up the synchronization inverse problem

We test our method using CH₄ series to synchronize the EDC and Talos Dome ice cores in the gas phase over the last ~45 ky. The current standard chronology for both the Talos Dome and EDC ice cores is AICC2012 (Bazin et al., 2013; Veres et al., 2013). In AICC2012, the manual tie points derived by Louergue et al. (2008) and Buiron et al. (2011) were not included, as both cores were independently synchronized to the NGRIP methane record (Capron et al., 2010) up to 50 ka BP. However, we can still use the manual tie points developed by Louergue et al. (2008) (the tie points derived by Buiron et al. (2011) are older than our test period) to cross-verify our code with the results of a manual synchronization.

We develop our synchronization based on the AICC2012 scenario, using the correlation matrices and priors as defined for the AICC2012 experiment in Parrenin et al. (2015), with updated uncertainties for the priors. We limit the chronologies to the first 1000 m of the EDC core, and the first 1100 m of the EDML core. Age grids for the LID and Accumulation correction functions are defined from 0 to 50,000 yr BP. We operate the dating experiment at very low resolution to minimize the necessary computation time (8 nodes per core for the thinning correction function, and an age interval of 2500 yrs for the LID and accumulation correction functions).

We then remove any chronological information that could interfere with or "help" the synchronization module. Apart from the CH₄ series for both cores, we retain only three ice dated points for EDC so that the absolute ages stay reasonable.

The CH₄ series used for EDC was measured by Loulergue et al. (2008), and the TALDICE series measured by Buiron et al. (2011). In addition to the CH₄ measurement uncertainty, an uncertainty of 10 ppbv is taken into account for the synchronization. This uncertainty is representative of the maximum offset for individual CH₄ peaks due to differences in the gas trapping processes between the two records (see, for example, the Dansgaard-Oeschger events in figures 2 and 3).

To correctly balance the terms of the cost function in the AICC2012 scenario, the priors were assigned uncertainties that prove to be too large for our scenario. Preliminary testing showed that leaving the original uncertainties from AICC2012 in place causes our problem to be ill-posed, with inappropriately low cost function values assigned to chronologies in which the methane series were well-synchronized but the glaciological scenarios physically unrealistic (thinning function values far from 1 near the ice surface, for example, or large changes in accumulation anticorrelated with trends in the isotopic series used to determine the priors).

We considerably reduce the uncertainties for all three variables: by a factor of 10 for the thinning function and EDC lock-in depth, and by a factor of 5 for the accumulation and TALDICE lock-in depth priors (note that these uncertainties are relative uncertainties of the raw variables, so these correspond to smaller changes for the log-transformed correction functions). We should note that these uncertainties are still very large, allowing for, in some cases, negative values of accumulation (which is generally impossible at the EDC and TALDICE sites). We also set the uncertainty with respect to the thinning function at the surface to 0 (impossible in the DATICE code), as the most recently accumulated snow is entirely unthinned. Tests with the original AICC2012 priors showed similar results for the CH₄ synchronization, but allowed LID to vary too far from the prior for both cores. Interestingly, in spite of the large changes in LID, the Depth-Ice Age and Depth-Air Age results are similar regardless of the prior. The correlation matrices for the three glaciological variables are left unchanged.

The Talos Dome and EDC CH₄ series on these two prior scenarios are shown in Figure 2. Visually, the two series are relatively well-matched in the shallowest parts of the core (from the onset of the Holocene period around ~ 11.5 ka BP onward). However, as we continue to older ages in the cores, the match between the two CH₄ series worsens. None of the Dansgaard-Oeschger events, between 40 ka BP and 20 ka BP, are particularly well-matched, with several false matches (i.e. around 30 ka) between different events. These false matches create local minima in the synchronization term of the cost function, and the ability to escape them is an important test of the true usability of our algorithm.

5.2 Convergence diagnostics

While MCMC simulations, in theory, are designed to converge to their target distribution (in our case, the posterior probability distribution), it is difficult in practice to guarantee that convergence has occurred, since the information we have about the posterior beyond the simulation itself is limited. For parallelized, multi-chain MCMC methods, Goodman and Weare (2010) recommend the integrated autocorrelation time of the variables as a diagnostic of the convergence of a simulation. We use this metric, as estimated in Foreman-Mackey et al. (2013) to evaluate the sampling error associated with each variables in the simulation. The sampling error, σ^2 , is expressed as:

$$\sigma^2 = \frac{act_f}{N} Var(f(\theta)), \quad (21)$$

where act_f is the autocorrelation time, N is the number of samples, and $f(\theta)$ the chain in question. Thus, we seek a small value of $\frac{act_f}{N}$ to minimize sampling error. Complicating this diagnostic, the estimate used by Foreman-Mackey et al. (2013) for act_f is biased for small values of $\frac{act_f}{N}$, asymptotically reaching correct values; once $N \sim 50act_f$ the estimate is usually robust. Of course, given unlimited computational resources, it is best to let the simulations run for much longer. The estimator should still be treated with care, as it is sensitive as well to parameters used in its own calculation.

In our test simulations, we found four potential ways to improve convergence:

1. Changing the MCMC steps. Including the variants of the Differential Evolution and snooker steps, particularly the mix of $\gamma = 1$ and $\gamma = 2.38 \sqrt{2d}$ greatly improved autocorrelation times by preventing the walkers from becoming stuck in local modes. MCMC steps well adapted to the shape of the posterior solution are a powerful tool to avoid astronomical autocorrelation times.
2. Increasing the number of iterations. This is the easiest solution, particularly when the simulation has nearly converged.
3. Adjusting the number of walkers. More walkers help the simulation start out in a larger sub-region of the posterior, but more walkers also require more iterations to converge appropriately.
4. Lowering the resolution of the dating experiment. If it is not imperative to optimize the chronologies at high resolution, it is much easier to converge to a 100-dimensional probability distribution than to a 1000-dimensional distribution.

5.3 Results: standard case

We run 128 model walkers for 1,000,000 iterations of the Markov Chain Monte Carlo algorithm as described above, and save one out of every 1,000 iterations to reduce memory usage. After 1,000,000 iterations, the two methane series are visually well-synchronized, as visible in figure 3 and the supplementary interactive figures. Indeed, our tests show that appropriate synchronizations can be reached by some walkers with as few as 1,000 iterations. Using the autocorrelation estimate as described above, the variables have all been sampled for more than 50 autocorrelation times – though we should keep in mind that the estimator of autocorrelation times is itself not necessarily reliable.




Figure 3. Methane series for TALDICE and EDC, on the maximum prior probability (left) and maximum posterior probability (right) scenarios. The cumulative tuning uncertainty is shown using the brown shading on the EDC series. Note that in each figure, only one synchronization is shown – see the supplementary interactive figures for the full ensemble, which is representative of the posterior distribution.

Line plots of the accepted scenarios are shown in figures 4 and 5. Each accepted scenario is plotted as a thin, semi-transparent line, so that in the regions where more scenarios are accepted, these plots show greater color density. It is important to note that the dependencies between variables are not shown in these plots (unlike in a multivariate histogram, for example). These plots, instead, show the MCMC estimate of the marginal distribution of each variable. The interdependencies/covariances between glaciological variables are better reflected in the ice ages and air ages as shown in figure 5, as these variables are integrative; and they are best understood using the composite animations in the digital supplement, which show the ensemble of all variables and calculated ice/air ages for each scenario. The composite animations can be downloaded at https://github.com/Jai-Chowdhry/IceChrono_new/tree/master/Interactive.

The accumulation and thinning functions do not deviate considerably from the prior for either scenario, which is in part by design—the synchronization, which is the largest component of the cost function, takes place in the air phase, and the age of the air phase can in theory be adjusted by changing only the LID. In practice, the thinning function and accumulation will be affected by the synchronization as well, in order to balance the prior terms in the cost function.

Both accumulation records remain relatively close to the prior, and do not present unrealistic values. Accumulation should be mostly constrained by the age of the ice matrix, which ultimately depends on the three dated points for EDC, and the three EDC dated points plus the synchronization for TALDICE. It is thus limited to realistic values by the model, which would produce incoherent results and high cost function values for negative accumulation values, for example, even though the range of the prior distribution contains unrealistic values.

Neither LID record has changed dramatically, with the TALDICE record appearing to have changed the most at around 35,000 yrs BP with respect to its prior. During this period, the prior scenarios were particularly badly synchronized.

Finally, the thinning function τ remains relatively similar to its prior for both cores.

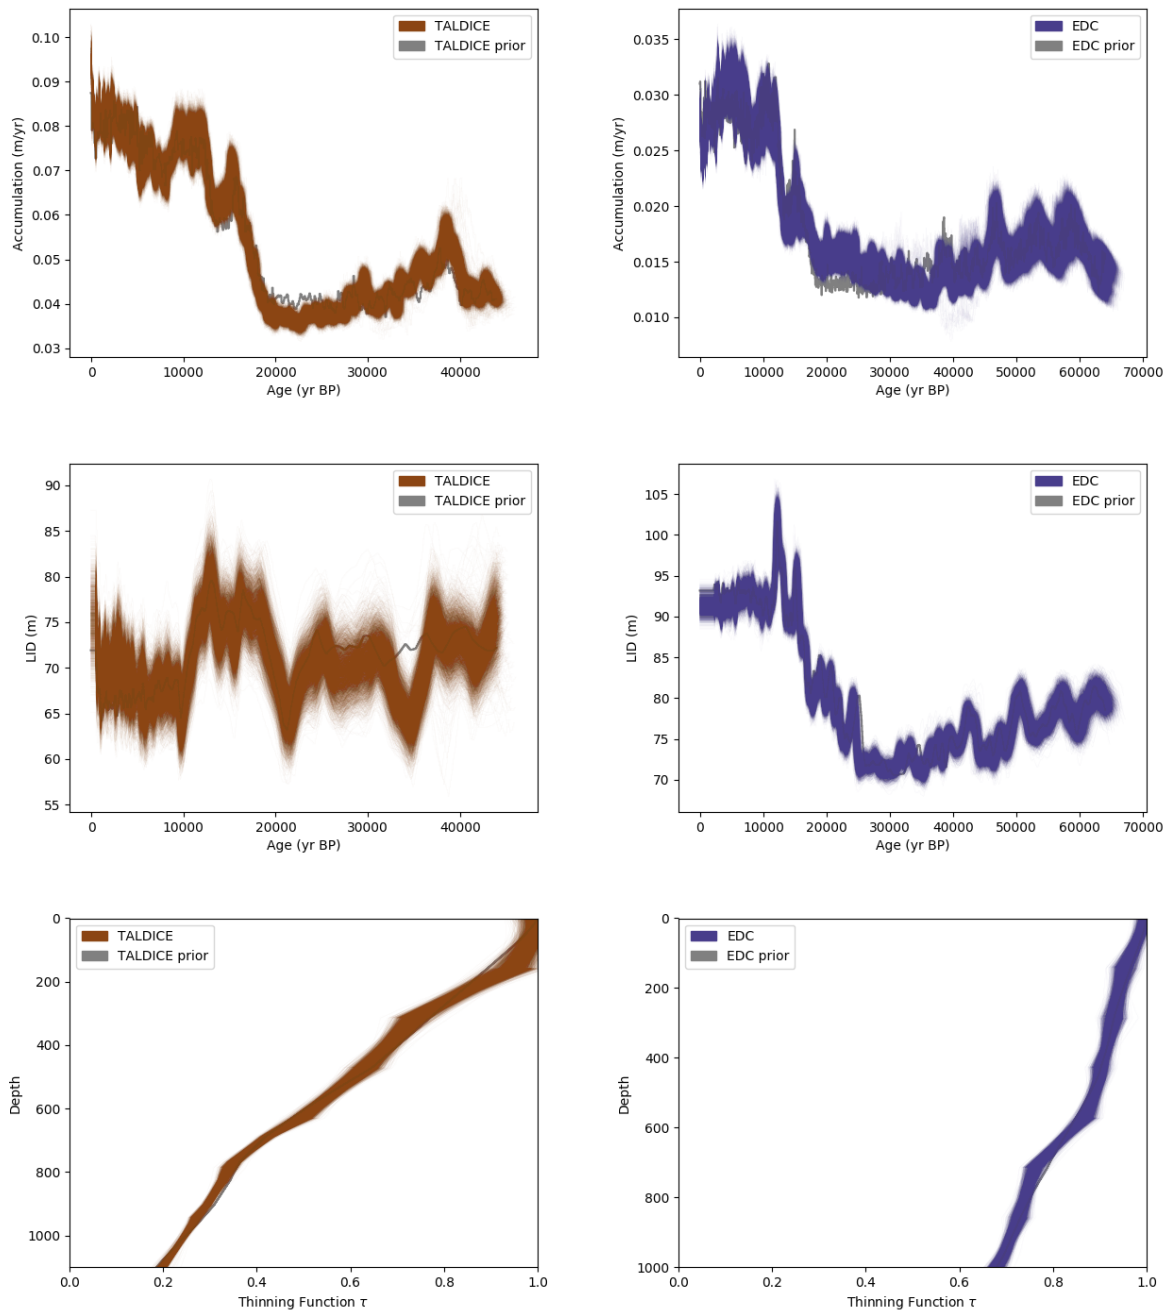


Figure 4. From top to bottom: Accumulation, Lock-In Depth and Thinning Function τ for TALDICE (left) and EDC (right). In the figures, we show the a sampling, reduced by a factor of 1000, of the 44,800,000 models accepted in the MCMC simulation, of 350,000 iterations for 128 walkers (44,800 models are shown in these figures). The models are plotted semi-transparently, so that an area with higher probability density appears denser in color.

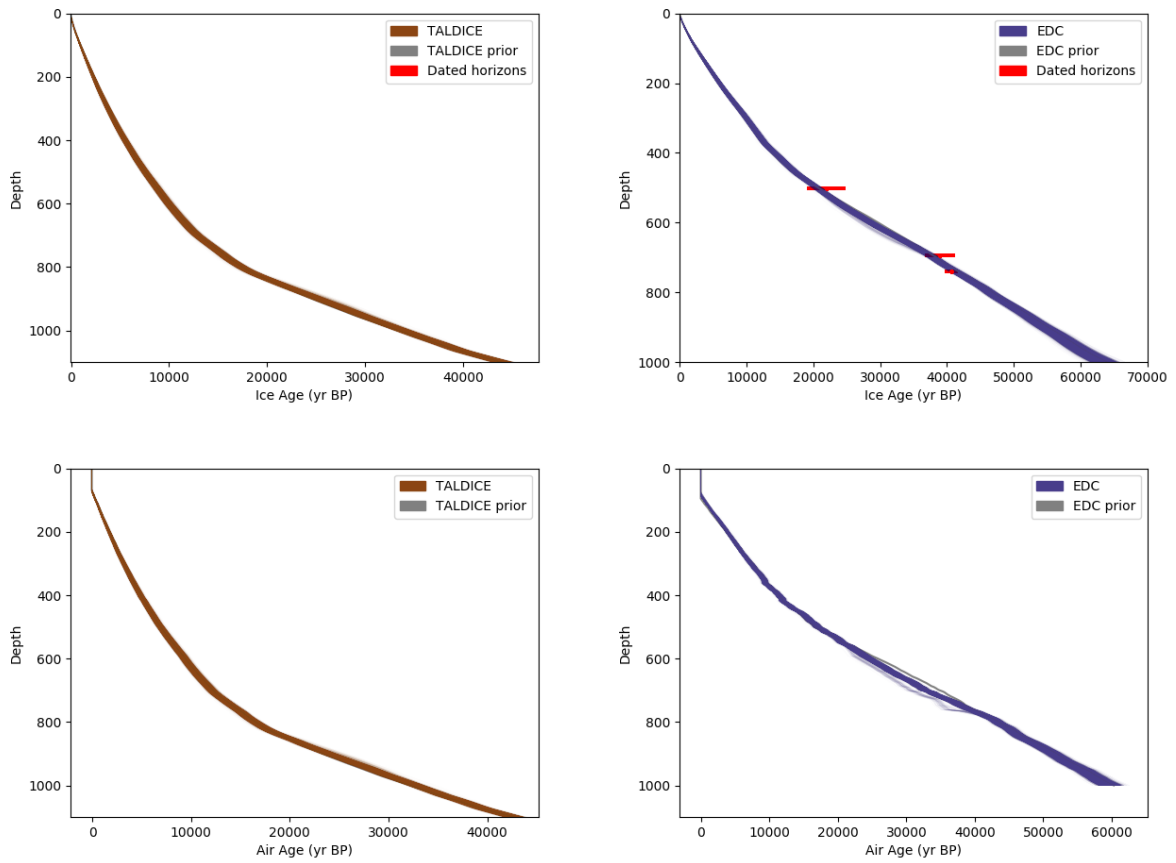


Figure 5. From top to bottom: ice age vs. depth and air age vs. depth for TALDICE (left, brown) and EDC (right, blue). In the figures, we show the a sampling, reduced by a factor of 1000, of the 44,800,000 models accepted in the MCMC simulation, of 350,000 iterations for 128 walkers (44,800 models are shown in these figures). The models are plotted transparently, so that an area with higher probability density appears denser in color.

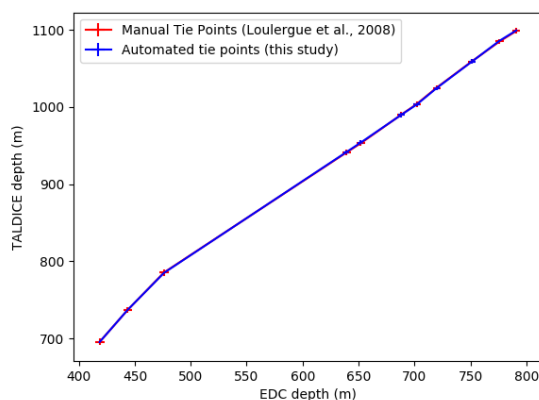


Figure 6. Comparison of the manual tie points between TALDICE and EDC derived by Loulergue et al. (2008) (red) and the same points recreated from the synchronization in this study (blue). The uncertainties estimated by eye (Loulergue et al., 2008) and by calculating the standard deviation of the ensemble of fits (this study) are shown as error bars.

In spite of the broad range of uncertainty for the priors, the ranges of air age and ice age for both TALDICE and EDC remain well constrained, as shown in figure 5. These variables integrate the three glaciological parameters, and we can expect that if the inverse problem is not too badly posed (i.e. the cost function terms are more or less appropriately balanced), the three glaciological parameters will covary so that a change in one is compensated by a change in the other, leading to an acceptable age scenario as per the discrete chronological information, as shown in the EDC Ice Age plot.

The methane series for each core for all accepted scenarios are shown in the supplementary interactive figures. This representation is useful to check whether our synchronization module works as expected. Indeed, the entire deglaciation is well-matched for the vast majority of accepted scenarios, as are each of the Dansgaard-Oeschger (D-O) events. Some events smaller than the D-O events are also well-synchronized—for example, the dip after the beginning of the Holocene period, around 8 ka BP. We confirm that our residual formulation works as expected.

We further compare our synchronization scenario to the manual tie points derived by Loulergue et al. (2008). Manual tie points are, in general, specified by identifying corresponding depths along the two cores, with uncertainty given either in terms of depth or age. In order to be able to compare our synchronization with the manual tie points, we take the EDC depths for each tie point, and identify the TALDICE depth with the same age for each scenario in the ensemble. We then simplify this information by calculating the mean and standard deviation of the corresponding TALDICE depths. The tie points derived from our synchronization and those derived by Loulergue et al. (2008) are shown in figure 6 and table 1. It is important to note that the Loulergue et al. (2008) tie points are derived at the midpoints in rapid transitions. Since this technique, unlike our cost function term, does not rely on concentration values, the results of the two synchronizations could be expected to differ.

The tie points derived from our synchronization correspond with those of Loulergue et al. (2008). The error bars derived by our synchronization tend to be smaller than those of Loulergue et al. (2008) in spite of having assigned a large additional

EDC Depth (m)	TALDICE Depth (m) Loulergue et al., 2008	1σ	TALDICE Depth (m) Automated	1σ
418.50	695.40	4.29	695.28	0.23
443.50	736.80	4.00	737.18	0.22
476.10	785.50	4.06	784.90	0.26
639.10	941.00	3.72	941.93	0.19
651.90	953.00	3.81	953.97	0.19
688.10	990.00	3.55	990.14	0.56
702.10	1004.00	3.53	1004.00	0.72
719.70	1025.00	3.29	1025.00	0.97
751.30	1059.20	3.43	1059.04	1.44
776.00	1085.00	3.57	1085.83	0.77
791.00	1098.60	3.63	1099.40	1.21

Table 1. Tie points as derived by Loulergue et al. (2008) between EDC and TALDICE, and the same points recreated from the synchronization in this study. The tie points are expressed for EDC depths (in m depth); the corresponding TALDICE depths and uncertainty estimates are given in m depth as well.

uncertainty to the synchronization (20 ppbv) to very roughly account for gas trapping differences and offsets between the two records. This is likely due to the continuous nature of our synchronization—we do not rely exclusively on midpoints, but rather match the entire series, leaving less room for ambiguity for the largest events. The last four matches, corresponding with smaller, more ambiguous events, have the largest uncertainties.

5 Three important points should be brought up regarding the estimated sigmas for this test. The first, and most important, is that the uncertainty related to smoothing and the gas trapping process is not taken into account here. This means that the uncertainties given by the automatic method still contain a major source of error, which should be corrected in future studies.

Second, the uncertainties presented here are calculated after the dating experiment, while the uncertainties for manual synchronization are estimated beforehand.

10 The third point is related to the shape of the posterior probability density. In presenting the automatic tie points using a mean and standard deviation, we assume the shape of their distribution should be approximately Gaussian, which is likely untrue; we also cannot express the covariance between tie points. As mentioned throughout the paper, the ensemble of accepted scenarios should always be preferred to any statistical summary of the synchronization.

5.4 Results: test of sensitivity to resolution

15 We run a test of the sensitivity of the synchronization solution to the resolution of the correction functions for the three glaciological variables. In this experiment, we reduce the resolution of the correction functions for LID and accumulation to 5,000 years, and use eight nodes for the thinning correction function. Since the resolution of this experiment is lower, it takes less iterations to converge, and we run the experiment for 400,000 iterations, after which all but one of the variables

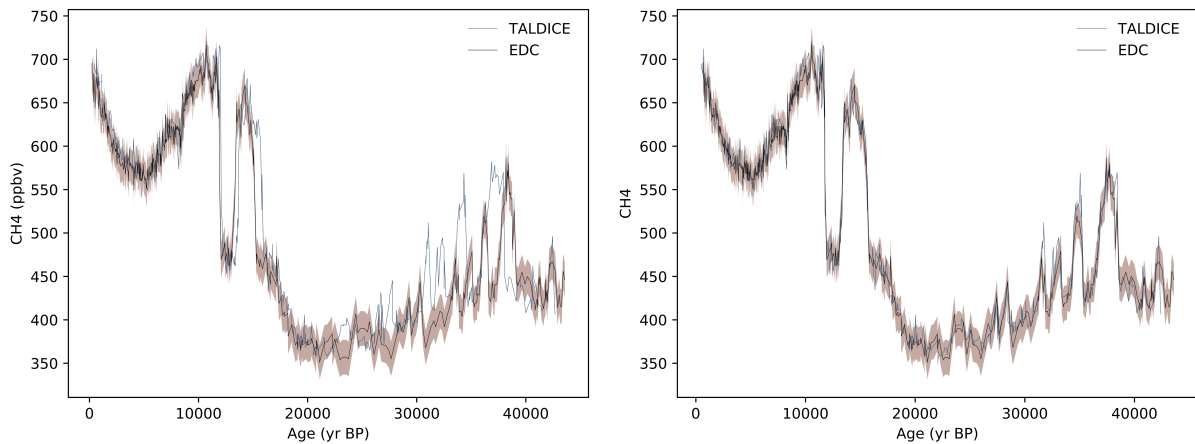


Figure 7. Methane series for TALDICE and EDC, on the maximum prior probability (left) and maximum posterior probability (right) scenarios. The cumulative tuning uncertainty is shown using the brown shading on the EDC series. Note that in each figure, only one synchronization is shown – see the supplementary interactive figures for the full ensemble, which is representative of the posterior distribution.

have crossed the threshold of 50 autocorrelation times. The synchronization of the two cores on the accepted scenario with the lowest cost function value in the run is shown in figure 8. The best synchronization for the low-resolution experiment is visually indistinguishable from the standard experiment.

Because of the considerable change in resolution, corrections to the prior chronology are made more coarsely, and the posterior densities for the three glaciological variables themselves may be different. This is confirmed by figure S1, where the posterior density of LID for EDC appears to deviate slightly from those sampled for the standard experiment; the posterior densities of the thinning function vary significantly as well. We seek to identify how these changes affect the proposed chronologies and synchronization scenarios.

Figure S2 shows a sampling of the ensemble of chronologies accepted during the low-resolution experiment. The ensembles are not markedly different from the standard experiment, though we should note that the air age at EDC between 40 ka BP and 20 ka BP has not explored as large of a region as in the standard experiment.

Tie points generated by the low-resolution experiment are compared to those calculated by Louergue et al. (2008) in table S1 and figure S3. The automated tie points correspond well to those estimated by Louergue et al. (2008) and the standard experiment. However, we note a significant difference in the estimation of the associated uncertainty. The low-resolution experiment tends to underestimate the uncertainty of the matches, particularly for the six deepest tie points, with respect to the higher-resolution experiment. This is likely because the very low resolution – 5,000 years – means that each node of the LID correction function changes the chronology on scales larger than the individual Dansgaard-Oeschger events. While we are still able to match the two series, the events themselves are not allowed to change in shape, thus leading to an error in the uncertainty estimation. We can infer that, for an appropriate estimation of error, the resolution of the dating experiment should be greater than the typical frequency of the events we wish to match. It may be possible to conduct very low-resolution

synchronization experiments by including an error term accounting for the potential bias induced by the resolution, but this would require further study.

6 External Synchronization test case: Fletcher Promontory to WAIS Divide

5 As a test of the capability of our method to synchronize an ice core dated within IceChrono to an external record, we use a high-resolution CH₄ series from the Fletcher Promontory (FP) ice core, measured at IGE, to synchronize FP to the WAIS Divide 2014 (WD2014) chronology (WAIS Divide Project Members, 2015). We use the WAIS Divide methane record as a target record: that is, we do not allow the WD2014 dating scenario to change (though were it to be allowed to change, the very low chronological uncertainty of WD2014 would not allow for much modification in any case).

6.1 Prior glaciological scenario for Fletcher Promontory

10 The glaciological priors are chosen to be roughly correct within the extent of our knowledge of the ice properties at the FP drilling site, using simple semi-physical descriptions. We later confirm that our choice of priors is very roughly appropriate by plotting the Fletcher Promontory CH₄ record on the prior age scenario, and checking whether the oldest and youngest ages for the series correspond approximately. For details on the definitions of the priors and their uncertainties, we refer the reader to the supplementary materials.

15 6.2 Discrete observations

We include five discrete ice age observations in the optimization to correctly date the ice phase. Sensitivity tests showed that excluding these points did not have a significant effect on the air chronology.

20 A small measurement offset exists between the WAIS divide and Fletcher Promontory cores. We make use of an iterative procedure to correct this offset: we allow the synchronization of the two cores to run for 10,000 iterations, and then correct the Fletcher Promontory series by the mean percentage offset from the WAIS divide core using the maximum probability scenario. After two iterations, the mean percentage offset is an order of magnitude smaller than its standard deviation, so we use this correction for our synchronization.

6.3 Result

25 We ran the FP-WAIS experiment with 256 walkers at 500-year resolution for 2.5 million iterations. Like the two-core experiments, we retained one out of every 1,000 iterations. Out of 218 variables, 10 had autocorrelation times estimated longer than 50,000 iterations, though the mean autocorrelation time registered lower, at 17,900 iterations. In any case, the result of this experiment should be treated with care, given the larger potential estimation error.

30 In this experiment, the prior scenario is highly mismatched from the WD chronology, in spite of our efforts to improve the prior glaciological variables. The MC simulation takes more iterations to leave the region of the prior (in sensitivity tests, the walkers all finally reach the WD-chronology region of probability after about 1 million iterations). Because of this, we run an




Figure 8. Methane series for FP and WAIS Divide, on the maximum prior probability (left) and maximum posterior probability (right) scenarios. The predefined tuning uncertainty is shown using the blue shading on the WD series. Note that in each figure, only one synchronization is shown – see the supplementary interactive figures for the full ensemble, which is representative of the posterior distribution.

"initialization" period of 25,000 iterations, after which the walkers are re-initialized by sampling from a small gaussian-shaped region close to the position of the lowest cost-function walker (details included in the supplement) and allowed to run for 2.5 million iterations. In an ideal scenario, the priors would be improved to be much closer to the appropriate synchronization, though this is not always practically possible.

- 5 The interactive figures in the supplement show good agreement between the CH₄ series, showing that our method is capable of external synchronization between the relatively high-resolution series that result from continuous flow analysis, and even with very approximate (incorrect) priors. We can identify a few "outlier" scenarios for the glaciological variables and ages, but their number, and the associated probability density, is minimal. Better informed and constrained priors, like those used in the EDC-TALDICE experiments, would likely eliminate these scenarios.

10 7 Assumptions and Limitations

- We emphasize that this is strictly a test of the synchronization module, and does not provide a new chronology for the Fletcher Promontory, Talos Dome or EDC records. In order to make sure that the synchronization of the series was indeed produced by the synchronization term in the cost function, and not simply matched by good supporting chronological information, we removed all additional chronological information in the EDC-TALDICE experiments, with the exception of three ice dated
 15 points for EDC, which anchor the EDC-TALDICE chronologies in a generally acceptable region, and should not greatly affect the synchronization, which is in the gas phase.

It should be noted that the EDC-TALDICE chronologies produced here are thus not necessarily correct (note, for example, that the Bolling-Allerod period occurs too early). In order to produce correct chronologies, we should include all of the addi-

tional information for the pair of cores we wish to match; this is outside the scope (and indeed contradictory to the verification purposes) of this study. Otherwise, a well-dated external record, like the WD series used in the external experiment, can be used to create a "correct" chronology.

Nor do we test the validity of the principles of methane synchronization – we simply cross-verify our result against a manual synchronization to confirm that our code can appropriately match two series. This does not mean that CH₄ synchronization is itself perfect – the treatment of gas trapping distributions in the cores, for example, would be a significant improvement.

A few assumptions inherent to our model should be highlighted:

- We assume that the two series used for synchronization can be assumed to be similar in shape and timing. The degree of truth of this assumption depends on the two series used. Methane series, for example, should come from ice cores that have similar gas trapping characteristics (or the additional uncertainty should be accounted for). It is, of course, possible to introduce two highly dissimilar series, but with the risk of introducing a high degree of complexity into the posterior, making the problem ill-posed and the solution vague. In the case that we know two series should correspond at certain points, but these series are dissimilar, we can either transform these series to be more similar (and account for the uncertainty added by the transformation) or simply use discrete tie points, which would likely be more reliable.
- We assume that the uncertainties associated with the prior and the observations can be represented by a Gaussian distribution.

We identify a few limitations, as well:

- Our method remains limited by computation time. Performing a federative dating experiment like AICC2012 at high resolution with appropriate convergence diagnostics would require significant computational resources. IceChrono is uploaded on the Dahu platform of the CIMENT computational cluster (<https://ciment.ujf-grenoble.fr>), and the computations performed here required between 7 and 24 hours on 48 cores. A much larger experiment would likely require further innovation for the MCMC procedure.
- Our method is not an independent synchronization method – that is, we do not derive tie points without creating chronologies. This has its advantages and disadvantages—the multiple sources of information in a chronology are all considered together, diminishing the potential for error and maximizing the reproducibility of chronologies, but we need considerable additional information. If the end goal is to simply derive tie points, the Dynamic Time Warping-based methods of Lisiecki and Lisiecki (2002) and Lin et al. (2014) are likely more applicable.
- Defining appropriate parameters for the priors and solution algorithm is delicate; the definition of the prior will always be prone to some subjectivity. In order to improve the priors, further glaciological study at the drilling sites is required.

8 Conclusions

Our method can synchronize two ice cores (or other paleoclimate archives) by adjusting accumulation, lock-in depth, and the thinning function. In the case of CH₄ synchronization, we were able to further constrain the chronology with respect to a manual synchronization because our method is continuous – we synchronize the entire series, rather than a few selected tie
5 points. This can add significant detail to chronologies. The method is capable of synchronizing an ice core to an external record as well.

Importantly, though, this is a dating method which *includes* a synchronization term, and not exclusively a synchronization method. Thus, we can also, and indeed must also include complimentary information to create a complete chronology. The glaciological priors must be defined carefully to avoid subjectivity, and require significant prior study. However, including the
10 priors and discrete age markers allows us to create robust chronologies that are consistent with all sources of chronological information.

Our method has potential for broad application, and we can envisage many perspectives. The model should be tested for the synchronization of volcanic peaks. We could also envisage a test of orbital tuning using isotopic ratios or Total Air Content. It might possible also to use the algorithm to tune to seasonal cycles, that is use it as an automated annual layer counting
15 algorithm. Finally, it will be important to investigate how to better account for the effects of diffusion and the gas trapping process when synchronizing in the gas phase.

Several further developments of IceChrono can be envisaged for the near future. For large dating experiments with multiple ice cores, we hope to implement a more efficient MCMC sampler for high-dimensional, multimodal densities. One way to do so may be to decrease the computational time per iteration. This could be done by modifying our code to allow for GPU
20 computing, and/or by training neural networks to replace the computationally intensive proposal step, as in Habib and Barber (2018). Additional possible developments for IceChrono include coupling with automated layer counting as in Winstrup et al. (2012); the development of on-line accumulation models; and modifications to better include paleoclimate archives beyond ice cores; this last development is in progress, and a manuscript describing the new model that can include many different paleoclimate archives, now called Paleochrono, is in preparation. Our synchronization module will eventually be included in
25 the new Paleochrono model as well.

9 Code availability

The code used in this study is available at https://github.com/Jai-Chowdhry/IceChrono_new.

10 Result availability

The full ensemble of accepted iterations from the MCMC simulation is available from the authors upon request, due to large
30 file size.

11 Interactive figures

Interactive figures are available at <http://pp.ige-grenoble.fr/pageperso/beeman/Synchronization/>. These figures are made with the Python Bokeh library, and are saved in .html format.

12 Acknowledgements

- 5 We thank Christo Buizert and Tyler J. Fudge for their input regarding our study. The computations presented in this paper were performed using the Dahu platform of the CIMENT infrastructure (<https://ciment.ujf-grenoble.fr>), which is supported by the Rhône-Alpes region (GRANT CPER07 13 CIRA), the OSUG@2020 labex (reference ANR10 LABX56) and the Equip@Meso project (reference ANR-10-EQPX-29-01) of the programme Investissements d'Avenir supervised by the Agence Nationale pour la Recherche. The development of the synchronization method was funded by the CNRS/INSU/LEFE IceChrono and
- 10 CO2Role projects, and the FP CH4 analyses were supported by the ANR-10-RPDOC-002-01 grant.

References

- Bazin, L., Landais, A., Lemieux-Dudon, B., Kele, H. T. M., Veres, D., Parrenin, F., Martinerie, P., Ritz, C., Capron, E., Lipenkov, V., et al.: An optimized multi-proxy, multi-site Antarctic ice and gas orbital chronology (AICC2012): 120-800 ka, *Climate of the Past*, 9, 1715–1731, 2013.
- 5 Bender, M. L.: Orbital tuning chronology for the Vostok climate record supported by trapped gas composition, *Earth and Planetary Science Letters*, 204, 275–289, 2002.
- Blunier, T., Spahni, R., Barnola, J.-M., Chappellaz, J., Loulergue, L., and Schwander, J.: Synchronization of ice core records via atmospheric gases, *Climate of the Past Discussions*, 3, 365–381, 2007.
- Buiron, D., Chappellaz, J., Stenni, B., Frezzotti, M., Baumgartner, M., Capron, E., Landais, A., Lemieux-Dudon, B., Masson-Delmotte, V.,
10 Montagnat, M., et al.: TALDICE-1 age scale of the Talos Dome deep ice core, East Antarctica, *Climate of the Past*, 7, 1–16, 2011.
- Buizert, C., Cuffey, K., Severinghaus, J., Baggenstos, D., Fudge, T., Steig, E., Markle, B., Winstrup, M., Rhodes, R. H., Brook, E., et al.: The WAIS Divide deep ice core WD2014 chronology—Part 1: Methane synchronization (68–31 ka BP) and the gas age–ice age difference, *Climate of the Past*, 11, 153–173, 2015.
- Buizert, C., Sigl, M., Severi, M., Markle, B. R., Wettstein, J. J., McConnell, J. R., Pedro, J. B., Sodemann, H., Goto-Azuma, K., Kawamura,
15 K., et al.: Abrupt ice-age shifts in southern westerly winds and Antarctic climate forced from the north, *Nature*, 563, 681, 2018.
- Capron, E., Landais, A., Lemieux-Dudon, B., Schilt, A., Masson-Delmotte, V., Buiron, D., Chappellaz, J., Dahl-Jensen, D., Johnsen, S., Leuenberger, M., et al.: Synchronising EDML and NorthGRIP ice cores using $\delta^{18}\text{O}$ of atmospheric oxygen ($\delta^{18}\text{O}_{\text{atm}}$) and CH_4 measurements over MIS5 (80–123 kyr), *Quaternary Science Reviews*, 29, 222–234, 2010.
- Collette, A.: *Python and HDF5*, O'Reilly Media, 2013.
- 20 Dreyfus, G., Parrenin, F., Lemieux-Dudon, B., Durand, G., Masson-Delmotte, V., Jouzel, J., Barnola, J.-M., Panno, L., Spahni, R., Tisserand, A., et al.: Anomalous flow below 2700 m in the EPICA Dome C ice core detected using $\delta^{18}\text{O}$ of atmospheric oxygen measurements, *Climate of the Past*, 3, 341–353, 2007.
- EPICA Community Members: Eight glacial cycles from an Antarctic ice core, *Nature*, 429, 623–628, 2004.
- Extier, T., Landais, A., Bréant, C., Prié, F., Bazin, L., Dreyfus, G., Roche, D. M., and Leuenberger, M.: On the use of $\delta^{18}\text{O}_{\text{atm}}$ for ice core
25 dating, *Quaternary science reviews*, 185, 244–257, 2018.
- Farr, B. and Farr, W. M.: *kombine*: a kernel-density-based, embarrassingly parallel ensemble sampler, in prep, 2015.
- Folk, M., Heber, G., Koziol, Q., Pourmal, E., and Robinson, D.: An overview of the HDF5 technology suite and its applications, in: *Proceedings of the EDBT/ICDT 2011 Workshop on Array Databases*, pp. 36–47, ACM, 2011.
- Foreman-Mackey, D., Hogg, D. W., Lang, D., and Goodman, J.: *Emcee*: The MCMC Hammer, *Publications of the Astronomical Society of
30 the Pacific*, 125, 306, doi:10.1086/670067, 2013.
- Fujita, S., Parrenin, F., Severi, M., Motoyama, H., and Wolff, E.: Volcanic synchronization of Dome Fuji and Dome C Antarctic deep ice cores over the past 216 kyr, *Climate of the Past*, 11, 1395–1416, 2015.
- Goodman, J. and Weare, J.: Ensemble samplers with affine invariance, *Communications in applied mathematics and computational science*, 5, 65–80, 2010.
- 35 Habib, R. and Barber, D.: *Auxiliary Variational MCMC*, 2018.
- Hastings, W. K.: *Monte Carlo sampling methods using Markov chains and their applications*, 1970.

- Kupiainen-Määttä, O.: A Monte Carlo approach for determining cluster evaporation rates from concentration measurements, *Atmospheric Chemistry and Physics*, 16, 14 585–14 598, 2016.
- Laloy, E. and Vrugt, J. A.: High-dimensional posterior exploration of hydrologic models using multiple-try DREAM (ZS) and high-performance computing, *Water Resources Research*, 48, 2012.
- 5 Landais, A., Dreyfus, G., Capron, E., Pol, E., Loutre, M.-F., Raynaud, D., Lipenkov, V. Y., Arnaud, L., Masson-Delmotte, V., Paillard, D., et al.: Towards orbital dating of the EPICA Dome C ice core using $\delta O_2/N_2$, *Climate of the Past*, 8, 191–203, 2012.
- Lemieux-Dudon, B., Blayo, E., Petit, J.-R., Waelbroeck, C., Svensson, A., Ritz, C., Barnola, J.-M., Narcisi, B. M., and Parrenin, F.: Consistent dating for Antarctic and Greenland ice cores, *Quaternary Science Reviews*, 29, 8–20, 2010.
- Levenberg, K.: A method for the solution of certain non-linear problems in least squares, *Quarterly of applied mathematics*, 2, 164–168,
 10 1944.
- Lin, L., Khider, D., Lisiecki, L. E., and Lawrence, C. E.: Probabilistic sequence alignment of stratigraphic records, *Paleoceanography*, 29, 976–989, 2014.
- Lipenkov, V. Y., Raynaud, D., Loutre, M., and Duval, P.: On the potential of coupling air content and O_2/N_2 from trapped air for establishing an ice core chronology tuned on local insolation, *Quaternary Science Reviews*, 30, 3280–3289, 2011.
- 15 Lisiecki, L. E. and Lisiecki, P. A.: Application of dynamic programming to the correlation of paleoclimate records, *Paleoceanography*, 17, 2002.
- Loulergue, L., Schilt, A., Spahni, R., Masson-Delmotte, V., Blunier, T., Lemieux, B., Barnola, J.-M., Raynaud, D., Stocker, T. F., and Chappellaz, J.: Orbital and millennial-scale features of atmospheric CH_4 over the past 800,000 years, *Nature*, 453, 383–386, 2008.
- Marquardt, D. W.: An algorithm for least-squares estimation of nonlinear parameters, *Journal of the society for Industrial and Applied
 20 Mathematics*, 11, 431–441, 1963.
- Martinson, D. G., Pisias, N. G., Hays, J. D., Imbrie, J., Moore, T. C., and Shackleton, N. J.: Age dating and the orbital theory of the ice ages: development of a high-resolution 0 to 300,000-year chronostratigraphy, *Quaternary research*, 27, 1–29, 1987.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E.: Equation of state calculations by fast computing machines, *The journal of chemical physics*, 21, 1087–1092, 1953.
- 25 Monnin, E., Steig, E. J., Siegenthaler, U., Kawamura, K., Schwander, J., Stauffer, B., Stocker, T. F., Morse, D. L., Barnola, J.-M., Bellier, B., et al.: Evidence for substantial accumulation rate variability in Antarctica during the Holocene, through synchronization of CO_2 in the Taylor Dome, Dome C and DML ice cores, *Earth and Planetary Science Letters*, 224, 45–54, 2004.
- Nelson, B., Ford, E. B., and Payne, M. J.: RUN DMC: an efficient, parallel code for analyzing radial velocity observations using N-body integrations and Differential Evolution Markov Chain Monte Carlo, *The Astrophysical Journal Supplement Series*, 210, 11, 2013.
- 30 Parrenin, F., Petit, J.-R., Masson-Delmotte, V., Wolff, E., Basile-Doelsch, I., Jouzel, J., Lipenkov, V., Rasmussen, S. O., Schwander, J., Severi, M., et al.: Volcanic synchronisation between the EPICA Dome C and Vostok ice cores (Antarctica) 0–145 kyr BP, *Climate of the Past*, 8, 1031–1045, 2012.
- Parrenin, F., Bazin, L., Capron, E., Landais, A., Lemieux-Dudon, B., and Masson-Delmotte, V.: IceChrono1: a probabilistic model to compute a common and optimal chronology for several ice cores, *Geoscientific Model Development*, 8, 1473–1492, 2015.
- 35 Petit, J.-R., Jouzel, J., Raynaud, D., Barkov, N. I., Barnola, J.-M., Basile, I., Bender, M., Chappellaz, J., Davis, M., Delaygue, G., et al.: Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica, *Nature*, 399, 429, 1999.
- Price-Whelan, A. M. and Foreman-Mackey, D.: schwimmbad: A uniform interface to parallel processing pools in Python, *The Journal of Open Source Software*, 2, doi:10.21105/joss.00357, <https://doi.org/10.21105/joss.00357>, 2017.

- Raisbeck, G., Yiou, F., Jouzel, J., and Stocker, T.: Direct north-south synchronization of abrupt climate change record in ice cores using Beryllium 10, *Climate of the Past*, 3, 541–547, 2006.
- Rasmussen, S. O., Seierstad, I. K., Andersen, K. K., Bigler, M., Dahl-Jensen, D., and Johnsen, S. J.: Synchronization of the NGRIP, GRIP, and GISP2 ice cores across MIS 2 and palaeoclimatic implications, *Quaternary Science Reviews*, 27, 18–28, 2008.
- 5 Raynaud, D., Lipenkov, V., Lemieux-Dudon, B., Duval, P., Loutre, M.-F., and Lhomme, N.: The local insolation signature of air content in Antarctic ice. A new step toward an absolute dating of ice records, *Earth and Planetary Science Letters*, 261, 337–349, 2007.
- Schüpbach, S., Federer, U., Bigler, M., Fischer, H., and Stocker, T.: A refined TALDICE-1a age scale from 55 to 112 ka before present for the Talos Dome ice core based on high-resolution methane measurements, *Climate of the Past*, 7, 1001–1009, 2011.
- Severi, M., Becagli, S., Castellano, E., Morganti, A., Traversi, R., Udisti, R., Ruth, U., Fischer, H., Huybrechts, P., Wolff, E., et al.: Synchronisation of the EDML and EDC ice cores for the last 52 kyr by volcanic signature matching, *Climate of the Past Discussions*, 3, 409–433, 2007.
- 10 Severi, M., Udisti, R., Becagli, S., Stenni, B., and Traversi, R.: Volcanic synchronisation of the EPICA-DC and TALDICE ice cores for the last 42 kyr BP, *Climate of the Past*, 8, 509–517, 2012.
- Sharma, S.: Markov Chain Monte Carlo methods for Bayesian data analysis in astronomy, *Annual Review of Astronomy and Astrophysics*, 15, 213–259, 2017.
- Sigl, M., Fudge, T. J., Winstrup, M., Cole-Dai, J., Ferris, D., McConnell, J. R., Taylor, K. C., Welten, K. C., Woodruff, T. E., Adolphi, F., et al.: The WAIS Divide deep ice core WD2014 chronology—Part 2: Annual-layer counting (0–31 ka BP), *Climate of the Past*, 12, 769–786, 2016.
- Suwa, M. and Bender, M. L.: Chronology of the Vostok ice core constrained by O₂/N₂ ratios of occluded air, and its implication for the 20 Vostok climate records, *Quaternary Science Reviews*, 27, 1093–1106, 2008.
- Svensson, A., Bigler, M., Blunier, T., Clausen, H. B., Dahl-Jensen, D., Fischer, H., Fujita, S., Goto-Azuma, K., Johnsen, S., Kawamura, K., et al.: Direct linking of Greenland and Antarctic ice cores at the Toba eruption (74 ka BP), *Climate of the Past*, 9, 749–766, 2013.
- Tarantola, A.: *Inverse problem theory and methods for model parameter estimation*, SIAM, 2005.
- ter Braak, C. J.: A Markov Chain Monte Carlo version of the genetic algorithm Differential Evolution: easy Bayesian computing for real 25 parameter spaces, *Statistics and Computing*, 16, 239–249, 2006.
- ter Braak, C. J. and Vrugt, J. A.: Differential evolution Markov chain with snooker updater and fewer chains, *Statistics and Computing*, 18, 435–446, 2008.
- Veres, D., Bazin, L., Landais, A., Toyé Mahamadou Kele, H., Lemieux-Dudon, B., Parrenin, F., Martinerie, P., Blayo, E., Blunier, T., Capron, E., et al.: The Antarctic ice core chronology (AICC2012): an optimized multi-parameter and multi-site dating approach for the last 120 30 thousand years, *Climate of the Past*, 9, 1733–1748, 2013.
- Vinther, B. M., Clausen, H. B., Johnsen, S. J., Rasmussen, S. O., Andersen, K. K., Buchardt, S. L., Dahl-Jensen, D., Seierstad, I. K., Siggaard-Andersen, M.-L., Steffensen, J. P., et al.: A synchronized dating of three Greenland ice cores throughout the Holocene, *Journal of Geophysical Research: Atmospheres*, 111, 2006.
- Vousden, W., Farr, W. M., and Mandel, I.: Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations, 35 *Monthly Notices of the Royal Astronomical Society*, 455, 1919–1937, 2015.
- WAIS Divide Project Members: Precise inter-polar phasing of abrupt climate change during the last ice age, *Nature*, 520, 661–665, 2015.
- Wales, D. J. and Doye, J. P.: Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms, *The Journal of Physical Chemistry A*, 101, 5111–5116, 1997.

Winstrup, M., Svensson, A., Rasmussen, S. O., Winther, O., Steig, E., and Axelrod, A.: An automated approach for annual layer counting in ice cores, *Climate of the Past*, 8, 1881–1895, 2012.

Chapter 5

Conclusions and Perspectives

At the base of this thesis lies a Bayesian framework, developed for two case studies in ice core paleoclimatology. This framework relies on the construction of a posterior probability density function, which, as stated by Tarantola [2005] represents a state of knowledge with respect to the physical system under study. Here, to conclude, we outline the implementation of the framework in both studies, the limitations of performing these studies using this framework, and finally some conclusions about the larger significance to the paleoclimate community.

In our first study, we estimate the timings of millennial-scale changes in CO₂ and Antarctic temperature during the last deglaciation and the phasings between the two series. The construction of the posterior is simple in this study – the probability of the observations is defined by the CO₂ and temperature series themselves and their uncertainties. We do not explicitly include a prior probability density, which implicitly means that we have picked a uniform prior probability distribution across the space of the potential change points and piecewise linear representations. In theory, we might in fact have some prior knowledge about where the changes occurred. However, this knowledge is largely visual – that is, we can identify the change points by eye – or depends on other models, and thus the inclusion of a non-uniform prior would have introduced a large degree of subjectivity. The data used in this case are the CO₂ series and individual and stacked Antarctic temperature time series. Probability distributions are specified with respect to these data according to measurement uncertainty, and in the case of the stack, uncertainty resulting from the stacking procedure. We model the series using piecewise linear functions, which can approximately represent the millennial-scale changes during the deglaciation.

The question addressed by this work has now been answered several times – perhaps because of its relevance to the carbon cycle and contemporary anthropogenic climate change. Notably, the studies of Monnin et al. [2001], Pedro et al. [2012] and Parrenin et al. [2013] each addressed the question of phasing between Antarctic temperature and CO₂ using the latest available data. Like our study, each of these studies developed a method to model the time series and calculate timings and phasings; and each study, ours included, relies upon the assumptions inherent to these models. In our case, for example, the changes in CO₂ and Antarctic temperature, while approximately linear in shape, were surely not exactly linear in reality. We also faced challenges of data resolution and uncertainty as well, though perhaps to a lesser extent than previous work. The treatment of climatic red noise in this work, for example, remains inexact, and could certainly be improved upon.

The studies of Monnin et al. [2001], Pedro et al. [2012] and Parrenin et al. [2013] show an incremental evolution of the estimates of timings and phasings between CO₂ and Antarctic temperature during T1, and our study again adjusts these estimates. We calculate that temperature most likely led CO₂ by ~ 0.5 kyr at the onset and end of Termination 1, and that the onset and end of the Antarctic Cold Reversal were nearly synchronous, with CO₂ showing a slight lead at the end. With these results, we can improve our understanding about the carbon cycle mechanisms at play, linking Antarctic Temperature and CO₂. The series seem to show two, if not three distinct lag timescales, indicating modulations or modifications of the mechanisms linking them.

The closest of the previous studies in terms of methodology to ours was that of Parrenin et al. [2013]. This study used piecewise linear fits to represent the time series as well, and the respective posteriors for the series at hand were sampled using simple Metropolis-Hastings simulations with 1000 iterations each. Though a Metropolis-Hastings simulation should in theory converge (eventually) to its target distribution, we know in practice that MCMC methods that are not adapted to multimodality can easily become stuck in local modes. In this study, initial change points were specified by hand before running the Metropolis-Hastings algorithm. Thus, the estimated change points of the study were likely correct, or at least within an appropriate range, since manual identification of changes in a time series with marked linear trends should generally be accurate. On the other hand, modes of probability beyond the local mode corresponding to the initial fit may not have been explored. Finally, Parrenin et al. [2013] presented the timings and phasings of the changes in Gaussian form.

In our study, we implement a more powerful MCMC method, allowing us to appropriately sample a complex posterior and avoid specifying an initial guess position (we do so stochastically). We make an effort to largely avoid approximations to the probability densities of the change point locations. Given the assumption that the changes themselves are well-represented by piecewise linear fits, we are able to treat their potential timings using histograms of the ensemble of fits accepted by the MCMC simulation; the calculation of the phasing is also done using the ensemble of fits. Though we give 68 % and 95 % central probability intervals for the phasings, if we wanted to make a further calculation using the lead-lag results, we could easily use the saved ensemble of fits to directly apply the stochastic selection.

A major limitation of both Parrenin et al. [2013] and our study is the assumption of linearity of the changes. Indeed, as we comment in the article, this assumption may bias

the calculation of the CO₂ change point at the end of the deglaciation, for example – it is difficult to confirm the location of this point visually, and a secondary probability peak seems to indicate that the true change might have occurred earlier. But our method could easily be adapted to accept a non-linear model of the changes – our only requirement is to be able to take residuals between the model and the data series. The resulting probability density functions may be complex and multimodal, but we have shown, particularly in the second study, that we can begin to explore such probability density functions using specialized Markov Chain Monte Carlo methods. Such a study is within close reach, and we would simply need to propose the non-linear model of the changes. This could perhaps be done by relying on the physical characteristics of CO₂ changes – since we know, for example, that the carbon cycle is characterized by feedbacks, we could model the changes using exponential functions, pieced together with linear functions for longer trends. The residual framework could even be applied directly to the output of simple carbon cycle models designed to calculate series of CO₂ and global or Antarctic temperature – a sort of very simple reanalysis. This could be an effective way to investigate the probabilities of the mechanisms included in the models to represent the changes that occurred.

In the CH₄ synchronization study, we construct a much more complicated posterior probability density function. We work with a glaciological model, defining prior scenarios for the three glaciological variables that describe a chronology. Since rigorous scientific study has been performed concerning each of these variables, the construction of the priors, though it still contains a degree of subjectivity, has a clear initial direction. The sources of information are also much more numerous – indeed, the idea behind IceChrono is that we should be able to include as many sources of chronological information as possible, to provide the most comprehensive chronologies possible. In the synchronization study, we develop a robust method to add a significant source of information that was included very roughly in previous chronologies. Only the work of Lisiecki and Lisiecki [2002], Lin et al. [2014] and Ahn et al. [2017] had achieved continuous synchronization in the context of paleoclimate records, and the combination of continuous synchronization with many additional sources of chronological information and a glaciological model had seemed to be a difficult problem, requiring the exploration of a complex, possibly multimodal probability density function.

Yet, with our framework, this seemingly tough problem is possible to explore. We give a first working solution – and a template by which other complex, continuous sources of chronological information might be included in formalized chronologies. The next imme-

diate advances will be the application of this method to volcanic synchronization. The volcanic synchronization problem is slightly more difficult for two reasons: the first, the jagged nature of the ice core record of volcanic events; and the second, that these events are recorded at the same times in ice cores, but the concentrations at which they are recorded depend on the distance to the event source. The first problem may require modifications to the MCMC algorithm, or even the use of a more powerful algorithm. An improved solution algorithm would be of great use, even beyond our field, but would be a significant mathematical study (on the scale of a doctoral thesis or beyond). However, applied mathematicians in many fields beyond paleoclimate work on similar problems as well, and we can surely expect the choice of powerful MCMC algorithms to expand in the coming years.

The second problem will likely be mitigated by the use of glaciological variables and external chronological information in IceChrono. Unlike a pure matching algorithm, the variable priors and supplementary observations allow us to constrain the synchronization in a way that is physically correct. Thus, we reduce the number of physically possible false matches. In some cases, the use of external information can lead to the identification of one correct match, as in Svensson et al. [2013], even for a match between Antarctic and Greenland cores. Given that a large amount of constraining information is available, it could certainly be included in an IceChrono synchronization.

We can also certainly envisage the application of our method to orbital tuning, itself a form of external synchronization. For orbital tuning, the largest challenge is that we begin to deal with truly inexact proxies – that is, the chemical species that we can measure in ice cores, commonly $\delta^{18}\text{O}$, Total Air Content and $\delta \text{O}_2/\text{N}_2$, are far from being exactly representative of the astronomical parameters that they are thought to track. Rather, they are linked to these parameters by physical processes that vary, and that may contain uncertain lags on the order of thousands of years [Bazin et al., 2016]. Here, the largest challenge is not the synchronization itself, but rather the construction of appropriate target functions / relationships between the proxies measurable in ice cores and the astronomical parameters themselves.

It may also be possible to apply our synchronization method to annual layer counting, though annual layer counting is not truly a form of synchronization. For each record used in layer counting, we would develop a “template” of the annual cycles of the record, extend this template well beyond the expected age of the ice core, and synchronize the record externally to the template series, which would remain fixed in age. In the use

of a template, this would not be dissimilar to the automated method of Winstrup et al. [2012]. However, the definition of a template poses a challenge in itself, and the increase in resolution necessary to synchronize on the annual scale would demand significantly more computational resources, or more efficient solution methods.

In general, for the broader application of automated synchronization for ice cores, we come upon two research problems that remain to be tackled:

1. **Computation:** The problem of computation can be addressed from two angles: improvements to efficiency of the computational method, or the addition of computational resources. In our case, either type of improvement would allow our method to become more generally applicable. At the time of writing, the synchronization method, used on the medium-scale computational resources of the CIMENT Froggy computational platform (making computations with up to 128 computational nodes, over times as long as 48 hours) can synchronize a small number of ice cores at low resolution. This is not yet sufficient for a dating experiment on the scale of AICC2012, where we would require several ice cores to be dated at high resolution. In order to do so, we would need either to implement a significantly more adapted method to sample from the posterior, or to call upon more substantial computational resources. For large-scale chronologies, an intermediate solution may be to determine tie points using a low-resolution implementation of our method, and then include these tie points as discrete sources of information in a high-resolution experiment. We would thus include a Gaussian approximation the synchronization, which would inherently mean losing some of the accuracy and precision associated with our method. On the other hand, it would allow for the use of much more efficient gradient-based methods, but still retain some of the objectivity of automated synchronization.
2. **Definition of targets:** To perform synchronizations using series beyond those assumed to perfectly match requires some work on the target series. This work, essentially, would be composed of studying the physical processes linking series to their targets. The work of Extier et al. [2018], for example, links atmospheric $\delta^{18}\text{O}$ in ice cores to its equivalent in speleothems, providing a new potential source for synchronization.

In chapter two, we mentioned what Tarantola [2005] refer to as the “movie strategy” to understand the stochastic results of high-dimensional inverse problems. This strategy informs how we treat the results of both of the studies presented in this thesis. We might

also refer the “ensemble” treatment of stochastic results. Briefly, this means that we treat the results of a stochastic sampling of the posterior probability density functions associated with a problem as the final result themselves, avoiding the loss of information via summary statistics (or treating summary statistics with care, and acknowledging their shortcomings, if they are required).

In the CO₂ - Antarctic Temperature study, the most accurate representation of the timings and phasings is not the set of maximum probability timings with 68 % and 95 % central probability intervals, but rather the set of histograms of the timings and phasings themselves. Even better than the histograms are the sets of accepted MCMC iterations, as the timings and phasings of the individual change points as determined by the piecewise linear fits for a given series may be correlated model (a point that we did not treat in the paper, but should surely treat in future studies).

In the synchronization study, the “true” chronology, similarly, is presented in the HDF5 file containing all of the accepted MCMC iterations. This HDF5 file can be used by studies wishing to use the full stochastic chronology to perform calculations with each of the accepted scenarios (or a thinned subset of these accepted scenarios), thus accurately propagating chronological uncertainty.

In this study, we also develop interactive figures to best reflect the “movie” strategy. Like the original animations created by Tarantola [2005], in these interactive figures, the user can click on the time series representing a given variable, and visualize the corresponding time series for the other variables as well. This visualization permits us to show complex covariances in a very high-dimensional space not conducive to multivariate histograms, for example. In this study, we take gaussian approximations of the synchronization result in order to be able to compare with the result of a manual synchronization, but this is of course only a simplification of the true result.

As the computational resources and methods necessary to tackling problems, like automated synchronization for ice cores, become available and commonly used, the treatment of ensemble results as acceptable, final results themselves will become common practice in the paleoclimate community. We hope that we have provided the community with some ideas of how to treat and present these results.

The method developed for the CO₂ - Antarctic Temperature study is likely of use to the broader paleoclimate community as well. It can fit approximately piecewise-linear time series with any number of change points, which could be of use in determining phasings between different paleoclimate proxies (speleothem records, ocean sediments etc.) across

many periods of climate change. Glacial terminations, in particular, often present changes in many paleoclimate variables that are piecewise-linear in character (see Landais et al. [2013], for example). The calculations we use to approximate the leads and lags between probability density functions could be of use, for example, in comparing the timings of dated glacial moraine stages, which often have age distributions that are not exactly Gaussian, to paleoclimate proxies to determine the response times of glaciers to climate events.

The study of CO₂ and Antarctic temperature has some significance beyond the field of paleoclimate as well. It seems likely from our results that variable timescales exist in the interaction between CO₂ and temperature in Antarctica, and that these timescales can vary from region to region in Antarctica as well. In addition, we identify lags (if small) in both directions. In the context of modern climate change, this result seems to confirm that we can expect longer-timescale feedbacks in the carbon cycle to respond to warming temperatures – it could be reasonable, for example, to expect an oceanic release of CO₂ on the timescale of several hundred years after the initial climate perturbation, as was the case at the onset of the deglaciation. Of course, the results are complex enough that it is difficult to point to one particular response or response mechanism that should be expected – the main point being that, indeed, we can probably expect several major feedback responses to take place.

The development of our synchronization method does not, itself, have large implications for the study of climate. However, it is of use precisely to studies like the first, allowing us to reduce and more accurately characterize the relative chronological uncertainties between paleoclimate records. To understand the complex workings of the climate system in the past, we will need, precisely, to be able to compare paleoclimate records from across the globe with low relative uncertainty between them. Our synchronization method is a significant step in this direction.

Bibliography

- Seonmin Ahn, Deborah Khider, Lorraine E Lisiecki, and Charles E Lawrence. A probabilistic pliocene–pleistocene stack of benthic $\delta^{18}\text{O}$ using a profile hidden markov model. *Dynamics and Statistics of the Climate System*, 2(1), 2017.
- Lucie Bazin, Amaelle Landais, Emilie Capron, Valérie Masson-Delmotte, Catherine Ritz, Ghislain Picard, Jean Jouzel, Marie Dumont, Markus Leuenberger, and Frédéric Prié. Phase relationships between orbital forcing and the composition of air trapped in Antarctic ice cores. *Climate of the Past*, 12(3):729–748, 2016.
- Thomas Extier, Amaelle Landais, Camille Bréant, Frédéric Prié, Lucie Bazin, Gabrielle Dreyfus, Didier M Roche, and Markus Leuenberger. On the use of $\delta^{18}\text{O}_{\text{atm}}$ for ice core dating. *Quaternary science reviews*, 185:244–257, 2018.
- A. Landais, G. Dreyfus, E. Capron, J. Jouzel, V. Masson-Delmotte, D. M. Roche, F. Prié, N. Caillon, J. Chappellaz, Markus Leuenberger, A. Lourantou, F. Parrenin, D. Raynaud, and G. Teste. Two-phase change in CO_2 , antarctic temperature and global climate during termination ii. *Nature Geoscience*, 6(12):1062–1065, 2013.
- Luan Lin, Deborah Khider, Lorraine E Lisiecki, and Charles E Lawrence. Probabilistic sequence alignment of stratigraphic records. *Paleoceanography*, 29(10):976–989, 2014.
- Lorraine E. Lisiecki and Philip A. Lisiecki. Application of dynamic programming to the correlation of paleoclimate records. *Paleoceanography*, 17(4), 2002.
- Eric Monnin, Andreas Indermühle, André Dällenbach, Jacqueline Flückiger, Bernhard Stauffer, Thomas F Stocker, Dominique Raynaud, and Jean-Marc Barnola. Atmospheric CO_2 concentrations over the last glacial termination. *Science*, 291(5501):112–114, 2001.
- Frédéric Parrenin, Valérie Masson-Delmotte, Peter Köhler, Dominique Raynaud, Didier Paillard, Jakob Schwander, Carlo Barbante, Amaëlle Landais, Anna Wegner, and Jean

Jouzel. Synchronous change of atmospheric CO₂ and Antarctic temperature during the last deglacial warming. *Science*, 339(6123):1060–1063, 2013.

Joel B Pedro, Sune Olander Rasmussen, and Tas D van Ommen. Tightened constraints on the time-lag between Antarctic temperature and CO₂ during the last deglaciation. *Climate of the Past*, 8(4):1213–1221, 2012.

Anders Svensson, Matthias Bigler, Thomas Blunier, Henrik Brink Clausen, Dorte Dahl-Jensen, Hubertus Fischer, S Fujita, K Goto-Azuma, SJ Johnsen, K Kawamura, et al. Direct linking of Greenland and Antarctic ice cores at the Toba eruption (74 ka bp). *Climate of the Past*, 9(2):749–766, 2013.

Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*, volume 89. siam, 2005.

Mai Winstrup, AM Svensson, Sune Olander Rasmussen, Ole Winther, EJ Steig, and AE Axelrod. An automated approach for annual layer counting in ice cores. *Climate of the Past Discussions*, 8(6):1881–1895, 2012.

Appendix A

Supplementary Material

Supplement of *Clim. Past*, 15, 913–926, 2019
<https://doi.org/10.5194/cp-15-913-2019-supplement>
© Author(s) 2019. This work is distributed under
the Creative Commons Attribution 4.0 License.



Supplement of

Antarctic temperature and CO₂: near-synchrony yet variable phasing during the last deglaciation

Jai Chowdhry Beeman et al.

Correspondence to: Frédéric Parrenin (frederic.parrenin@univ-grenoble-alpes.fr)

- [cp-15-913-2019-supplement-title-page.pdf](#)
- [supplement](#)
 - [leadlagtable.xlsx](#)
 - [supplement.pdf](#)

The copyright of individual parts of the supplement might differ from the CC BY 4.0 License.

Details of the stochastic simulation

A schematic of how residuals between the change point representation and the time series are calculated is included in Figure 1.

In order to eliminate the need to specify an initial fit to estimate the autocorrelation matrix, we perform two subsequent
 5 optimizations. During the initialization, we use very rough estimates of the autocorrelation matrix of the residuals and the modeling uncertainty: identity and the standard deviation of the observations, respectively. An automated first guess is proposed by randomly generating the $n x_i$ values and interpolating the y_i values from the observations. This quickly orients the optimization toward the correct region. The optimization is then allowed to run for an initialization period of 25,000 iterations, with 50 individual walkers.

10 The best fit encountered in the initialization is used to start the true optimization (250,000 iterations with 50 walkers): modeling uncertainty is re-estimated as the mean of the absolute residuals on the best fit; and the autocorrelation matrix is reestimated by taking the autocorrelation of the residuals of the best fit. A histogram of the proposals accepted by the second optimization is taken as the final result.

To conserve computational time, the individual temperature series are run with 10,000 initialization iterations and 100,000
 15 optimization iterations, and the Savitsky-Golay filtered series with 1000 initialization iterations and 50,000 optimization iterations. Tests show that runs with 10,000 optimization iterations achieve reasonable convergence.

Parallel MCMC methodology

To propose updates to the walkers, we apply what Goodman and Weare (2010) refer to as a "stretch move". Consider the ensemble of walkers, in this case representing potential piecewise linear fits \mathbf{X} and an individual walker \mathbf{X}_k^j in the ensemble at
 20 proposal step j . We select another walker \mathbf{X}_h^j from the complementary ensemble $\mathbf{X}_{[k]}^j$, composed of all of the other walkers. Then, a proposal is made to update \mathbf{X}_k to W :

$$\mathbf{X}_k^j \rightarrow W = \mathbf{X}_h^j + Z \left(\mathbf{X}_k^j - \mathbf{X}_h^j \right) \quad (1)$$

GW define the following probability distribution to generate stochastic variable Z :

$$g(Z) \propto \begin{cases} \frac{1}{Z} & \text{if } Z \in \left[\frac{1}{A_Z}, A_Z \right] \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

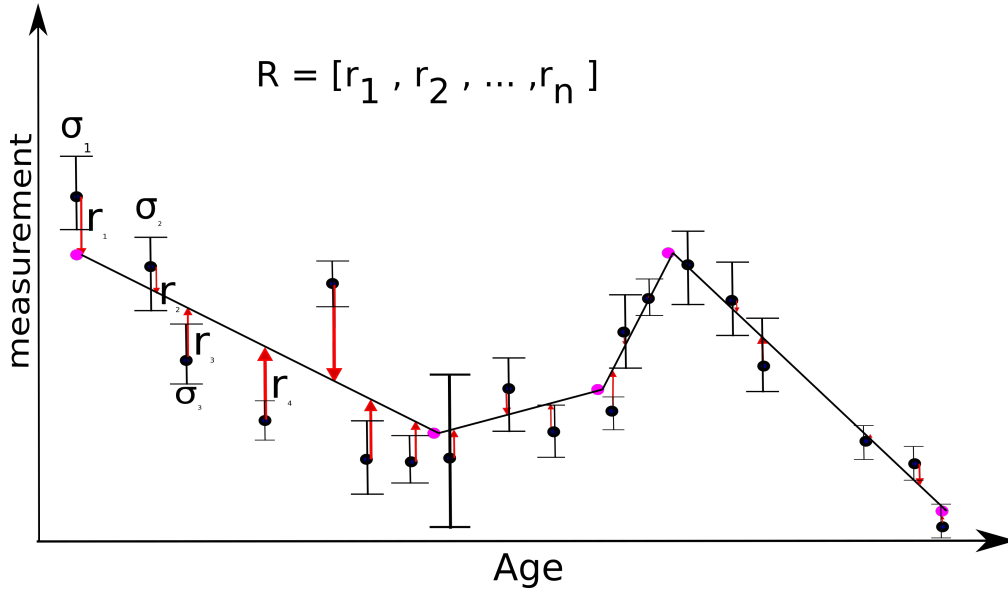


Figure 1. Schematic of the calculation of residuals between a data series and its change point representation. Here, the data series is shown as black points, with 1σ error bars. Change points are shown in pink, and the interpolations between them as black lines. The differences between the data points and the interpolations are shown as red arrows; the residual vector is composed of this distance divided by the uncertainty σ at each point.

5 where A_Z is a user-defined constant. In this study, the constant A_Z is set to 2, following Foreman-Mackey et al. (2013). Proposals are accepted or rejected with acceptance probability

$$P_{X_k^j \rightarrow X_k^{j+1}=W} = \min \left\{ 1, Z^{j-1} \frac{\exp(-J(W))}{\exp(-J(X_k^j))} \right\}. \quad (3)$$

BIC and change point sensitivity tests

We calculate the BIC for both series individually, and an additive, normalized BIC to pick the number of change points used. The additive BIC gives 7 as the best number of points to coherently fit both series. Using the individual criteria, 6 points would have been chosen for temperature, and 9 for CO₂, but comparing different numbers of points makes it less clear if the changes represented are actually coherent in timescale. These are shown in Figure 2.

We run tests to test the sensitivity of our results to the use of 5-9 change points rather than 8 (Figure 3). It is clear from these figures that 5-point fits cannot accurately represent either series; at 6 points, the temperature series begins to be well-represented, and at 9 points the CO₂ series is appropriately represented. The 7, 8 and 9 point fits do not show significant differences, except at the ACR onset.




Figure 2. Individual (left, CO₂ in black and ATS3 in red) and cumulative BIC (right) for 5,6,7,8,9 and 10 point (x-axis) runs.

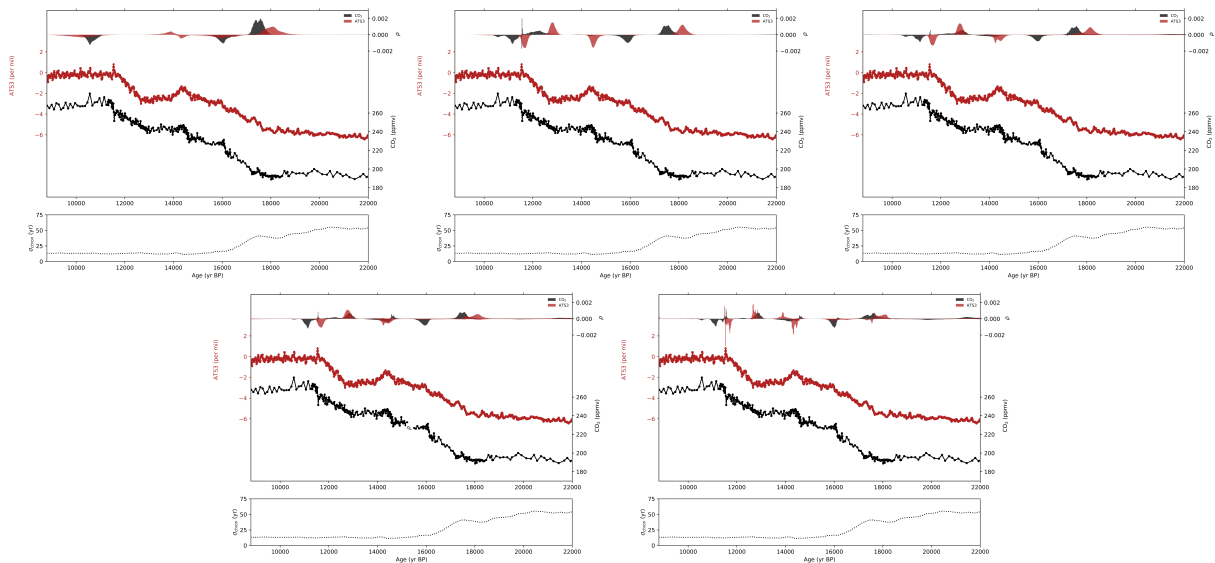


Figure 3. Atmospheric CO₂ (black) and ATS3 histograms of probable change points from the 5,6,7, 8 and 9 (from top left to bottom right) runs of LinearFit.




Figure 4. Plot of residuals r_i (x axis) against the predictions made using our AR(1) model $r_i = r_{i-1} \cdot a^{t_i - t_{i-1}}$ (y axis) shown as blue dots, for the CO₂ series after the initialization MCMC procedure. The orange line represents a perfect model fit.

Test of AR(1) residuals

We test our AR(1) model by plotting adjacent residuals of the CO₂ series against their predictions, using the robust AR1 model adapted to unevenly spaced time series described in the methods section. The results are shown in figure 4.

20 A note on Gaussian Estimates

We avoid providing estimates of change point timings in Gaussian form, i.e. as a mean \pm a standard deviation. This is because the change point histograms we calculate are often skewed, and sometimes even multimodal, with multiple, separate probability peaks in a given time period where change is likely. These peaks can be due to sub-millennial scale variations at these points in the two series.

5 A test with known lags and covariance

We use artificially generated series to test the capacity of our method to fit change points in series with four known change points (plus two end points). For each series, a covariance matrix is used to generate noise with a red component, with a correlation coefficient of 0.8 at 50 years and 0.64 at 100 years away from the central point, and a uniform white component.

10 The noise is scaled to 10% of the standard deviation of the change point values. The results of this test, along with the original change points, are shown in Figure 5.

In Figure 5, note that all of the change points are correctly identified in the histograms, with uncertainties on the order of 100 years. However, a small, but incorrect probability peak is generated around 13 ka. We conclude that gradual changes, such

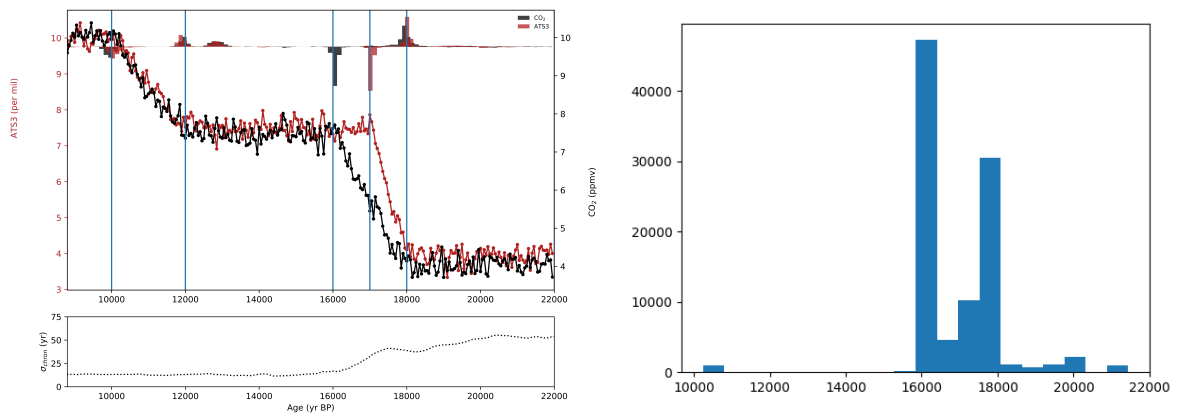


Figure 5. Artificially generated series and generated change point distributions.

as that around 12 ka, are slightly more difficult to extract from correlated noise even when this noise is modeled, adding some
 15 methodological uncertainty which is reflected in the histograms.

References

- Foreman-Mackey, D., Hogg, D. W., Lang, D., and Goodman, J.: Emcee: The MCMC Hammer, *Publications of the Astronomical Society of the Pacific*, 125, 306, doi:10.1086/670067, 2013.
- Goodman, J. and Weare, J.: Ensemble samplers with affine invariance, *Comm app math comp sci*, 5, 65–80, 2010.

Supplement to: Automated and probabilistic ice core synchronisation in IceChrono 2.0.

Jai Chowdhry Beeman¹, Frédéric Parrenin¹, Emmanuel Witrant², Amaëlle Landais³, Xavier Fain¹, Benoît Tournadre⁴, and Robert Mulvaney⁵

¹Université Grenoble Alpes – CNRS – IRD – Institut des Géosciences de l'Environnement, F-38000 Grenoble, France

²Université Grenoble Alpes – CNRS – Grenoble Images Parole Signal Automatique, F-38000 Grenoble, France

³Laboratoire des Sciences du Climat et de l'Environnement, CEA-CNRS-UVSQ, Gif sur Yvette, France

⁴MINES ParisTech, PSL Research University, O.I.E. - Centre Observation, Impacts, Energy, 06904 Sophia Antipolis, France

⁵British Antarctic Survey – NERC – Cambridge, UK

1 Fletcher Promontory continuous CH₄ data

The Fletcher Promontory, in the southwest corner of the Weddell Sea, is a small ice cap 650m thick. The Fletcher Promontory (FP) ice core was drilled in January 2012, reaching the bedrock at 654.3 m depth from the surface, as part of a UK-French collaborative project. The FP ice core sections (73 meters, from a 512-585 m depth range) were analyzed at high resolution
5 for methane concentration at IGE in 2016 over a 9-day period and using a continuous ice core melting system with online gas measurements (CFA, continuous-flow analysis). Detailed descriptions of this method have been reported before (Stowasser et al., 2012; Chappellaz et al., 2013; Rhodes et al., 2013) and the IGE CFA setup has been specifically described by Fourteau et al. (2017).

Briefly, ice core sticks of 34 by 34mm were melted at IGE at a mean rate of 3.8 cm min⁻¹, and the water and gas bubble
10 mixture was pumped toward a low-volume T-shaped glass debubbler. All the gas bubbles and approximately 15 % of the water flow were transferred from the debubbler to a gas extraction unit maintained at 30° C. The gas was extracted by applying a pressure gradient across a gas-permeable membrane (optimized IDEX in-line degasser; internal volume 1 mL). The gas pressure recorded downstream of the IDEX degasser was typically 500–600mbar. A homemade Nafion dryer with a 30 mL min⁻¹
15 purge flow of ultrapure nitrogen (Air Liquide; 99.9995 % purity) dried the humid gas sample before entry into the laser spectrometer. Online gas measurements of methane were conducted with a SARA laser spectrometer based on optical-feedback cavity-enhanced absorption spectroscopy (OF-CEAS; Morville et al. (2005); Romanini et al. (2006)). The SARA spectrometer was calibrated onto the NOAA2004 scale (Dlugokencky et al., 2005) before the CFA analyses using three synthetic air standards with known methane concentrations (Scott-Marrin, Inc.). CH₄ concentrations measured during the calibration agreed with NOAA measurements within 0.1 % over a 360–1790 ppbv range.

20 Allan variance tests (Allan, 1966; Rhodes et al., 2013) were conducted using mixtures of degassed deionized water and synthetic air standard to evaluate both the stability and the precision of the measurements. The best Allan variance was obtained

with an integration time larger than 1000 s, illustrating the very good stability of the CFA system. In this study, we used an integration time of 110 s (data resolution decreased to 7 cm) for which a precision of 0.5 ppbv (1σ) was observed.

The mixing of gases and meltwater during the sample transfer from the melt head to the laser spectrometer induces a CFA experimental smoothing of the signal. The extent of the CFA-based damping was determined by performing a step test, i.e., a switch between two synthetic mixtures of degassed DI water and synthetic air standards of different methane concentrations, following the method of Stowasser et al. (2012). It shows that the CFA system can resolve signals down to the centimeter scale (Fourteau et al., 2017).

Breaks along the core regularly let ambient air enter the system, resulting in strong positive spikes in methane concentration. In order to remove these contamination artifacts, exact times corresponding to a break running through the melt head were recorded during the measurements and later used to identify and clean the data from contamination.

In the measurement line, once the ice core is melted, the liberated gases dissolve in the melt water. However, methane has a higher solubility than oxygen or nitrogen, and it follows that CH₄ dissolves preferentially, leading to reduced methane mixing ratios in the extracted air reaching the SARA spectrometer. The solubility correction was evaluated by comparing the FP signal with WAIS Divide ice core data: a CFA record already corrected for solubility effects (Rhodes et al., 2015). We make use of an iterative procedure to correct this offset: we perform a synchronization (1000 iterations) of the two cores, and then correct the Fletcher Promontory series by the mean percentage offset from the WAIS divide core on the minimum cost function scenario. After two iterations, the mean percentage offset is an order of magnitude smaller than its standard deviation, so we use this correction for our synchronization, taking the standard deviation of the (non-percentage) offset as the modeling error, which is included alongside measurement error (by summing squares) in $\sigma_{p,i}^{\text{sync}}$. The resulting solubility correction coefficient of 1.12 was applied to the whole FP dataset.

2 Priors for the Fletcher Promontory/WD Experiment

Since the glaciological priors influence the posterior via Bayes' theorem (that is, they are included in the cost function as a regularization for the optimization problem), we choose them to be roughly correct within the extent of our knowledge of the ice properties at the FP drilling site.

The prior accumulation scenario is calculated, as in Parrenin et al. (2015) from the ice deuterium record, with the relationship:

$$a_b = a_0 \exp(8\beta\Delta D). \quad (1)$$

We take $a_0 = 0.5$ and $\beta = 0.025$. The prior LID scenario is set constant at 78 m. The prior thinning function τ is defined again as in Parrenin et al. (2015), making use of ice flow modeling concepts from Lliboutry (1979), Parrenin et al. (2006) and Parrenin and Hindmarsh (2007):

$$\tau^b = (1 - \mu)\omega + \mu. \quad (2)$$

We set μ , the ratio between accumulation and melting, to .021. The flux shape function ω is defined as

$$\omega = 1 - \frac{p+2}{p+1} (1-\zeta) + \frac{1}{p+1} (1-\zeta)^{p+2} \quad (3)$$

Where ζ is a vertical coordinate normalized by subtracting basal height and dividing by ice thickness. We set parameter p to 1.7.

5 We later confirm that our choice of priors is not totally inappropriate by regarding the Fletcher Promontory CH₄ record on the prior age scenario. As shown in the main text, this is only a first-order check that the ages of the methane record are correct to within a few thousands of years, and the ages themselves are still quite far from the correct scenario. The relative uncertainties (and covariance matrices) on the priors are set as in Parrenin et al. (2015), with diagonal variance set to 20 % for accumulation and lock-in depth, and across-diagonal linear decrease duration of 4000 yrs as in Bazin et al. (2013). The
10 covariance matrix of the thinning function, as in Parrenin et al. (2015), is assumed linearly related to the ratio of ice-equivalent depth and ice-equivalent layer thickness by constant $k = 0.1$, and across-diagonal linear decrease duration is set to 70 m.

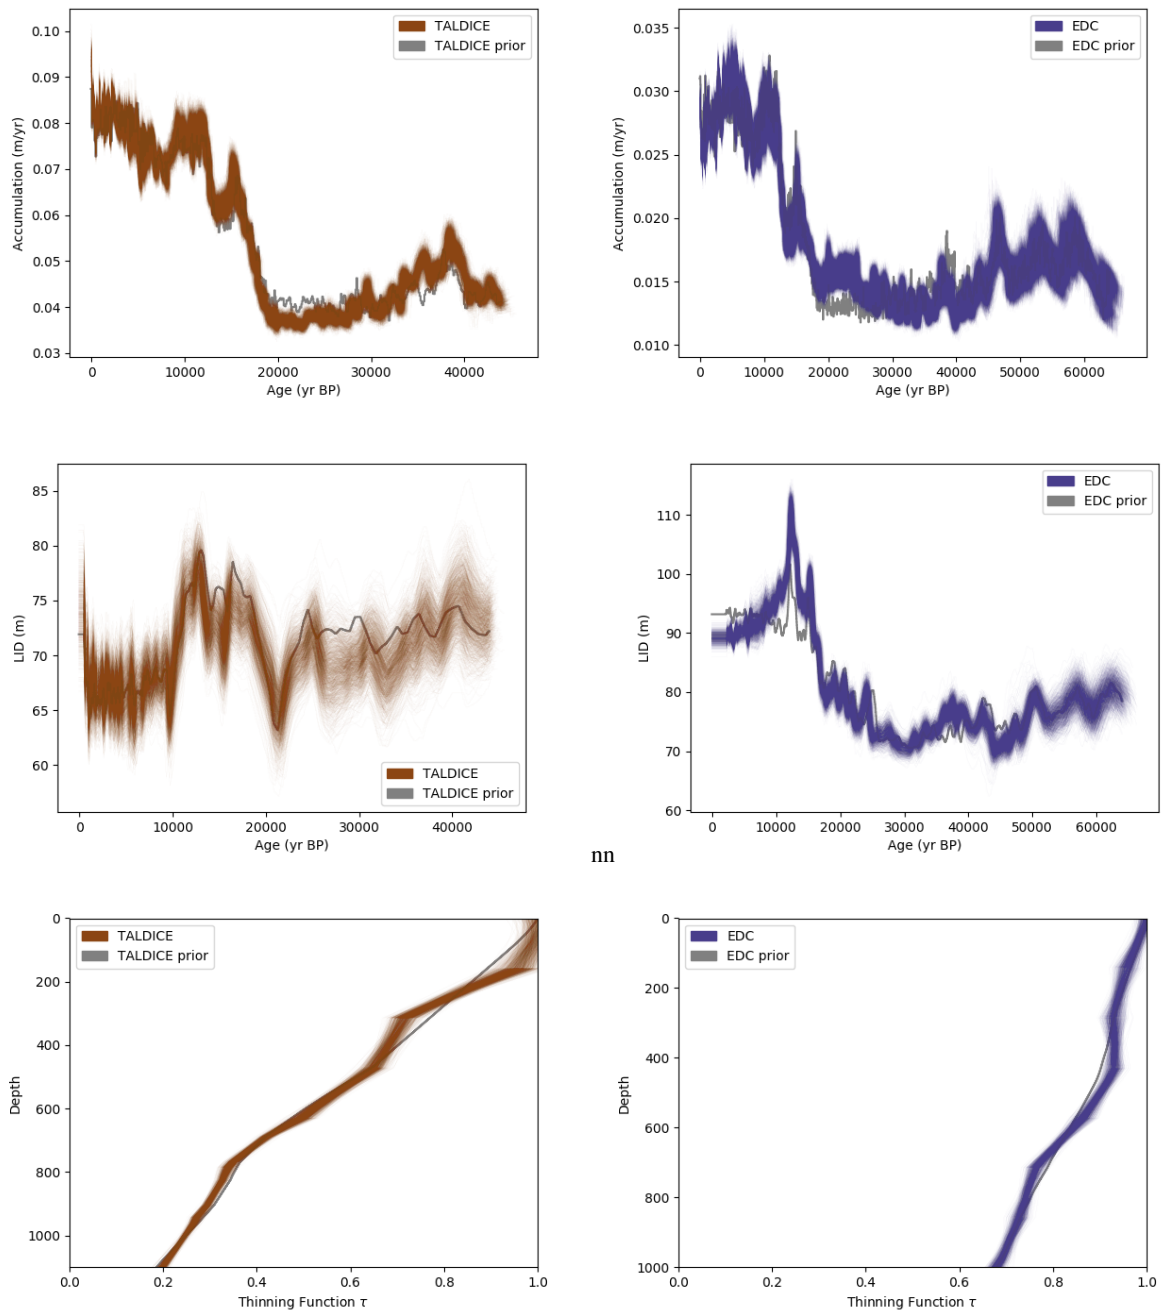
It is important to note, as mentioned in the main text, that these priors are still quite poor in terms of synchronization to WAIS Divide. We use a re-initialization procedure to quickly orient the walkers to regions of high probability and accelerate convergence. The walkers are first initialized according to the prior and allowed to evolve for 25,000 iterations.

15 3 Digital supplement

The digital supplement, including interactive figures and an animation of the synchronizations and glaciological variables for the TALDICE-EDC experiments, can be found at <http://pp.ige-grenoble.fr/pageperso/beeman/Synchronization/>.

4 Test: Sensitivity of EDC-TALDICE experiment to resolution

The supplementary figures and table for the sensitivity test are given here.



nn

Figure 1. From top to bottom: Accumulation, Lock-In Depth and Thinning Function τ for TALDICE (left) and EDC (right). In the figures, we show the a sampling, reduced by a factor of 1000, of the 44,800,000 models accepted in the MCMC simulation, of 2,500,000 iterations for 128 walkers (44,800 models are shown in these figures). The models are plotted semi-transparently, so that an area with higher probability density appears denser in color.

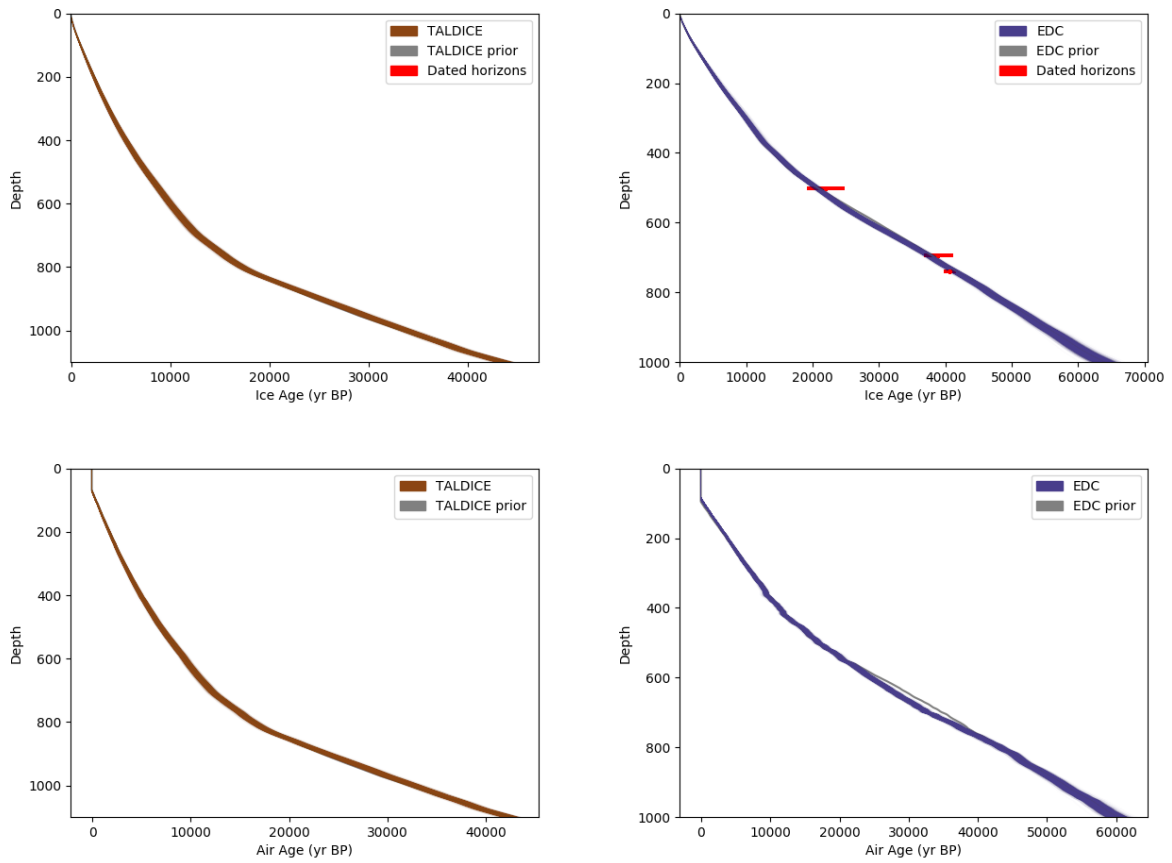


Figure 2. From top to bottom: Ice Age vs. Depth and Air Age vs. depth for TALDICE (left, brown) and EDC (right, blue). In the figures, we show the a sampling, reduced by a factor of 1000, of the 44,800,000 models accepted in the MCMC simulation, of 2,500,000 iterations for 128 walkers (44,800 models are shown in these figures). The models are plotted transparently, so that an area with higher probability density appears denser in color.

EDC Depth (m)	TALDICE Depth (m) Loulergue et al., 2008	1σ	TALDICE Depth (m) Low-Res Auto	1σ
418.50	695.40	4.29	695.39	0.13
443.50	736.80	4.00	737.08	0.15
476.10	785.50	4.06	784.53	0.25
639.10	941.00	3.72	942.90	0.33
651.90	953.00	3.81	954.25	0.30
688.10	990.00	3.55	989.17	0.12
702.10	1004.00	3.53	1004.46	0.16
719.70	1025.00	3.29	1024.14	0.07
751.30	1059.20	3.43	1059.39	0.14
776.00	1085.00	3.57	1087.05	0.64
791.00	1098.60	3.63	1099.75	0.93

Table 1. Tie points as derived by Loulergue et al. (2008) between EDC and TALDICE, and the same points recreated from the low-resolution synchronization experiment in this study. The tie points are expressed for EDC depths (in m depth); the corresponding TALDICE depths and uncertainty estimates are given in m depth as well.




Figure 3. Comparison of the manual tie points between TALDICE and EDC derived by Loulergue et al. (2008) (red) and the same points recreated from the low-resolution synchronization in this study (blue). The uncertainties estimated by eye (Loulergue et al., 2008) and by calculating the standard deviation of the ensemble of fits (this study) are shown as error bars.

References

- 5 Allan, D. W.: Statistics of atomic frequency standards, *Proc. IEEE*, p. 221–230, doi:<https://doi.org/10.1109/proc.1966.4634>, 1966.
- Bazin, L., Landais, A., Lemieux-Dudon, B., Kele, H. T. M., Veres, D., Parrenin, F., Martinerie, P., Ritz, C., Capron, E., Lipenkov, V., et al.: An optimized multi-proxy, multi-site Antarctic ice and gas orbital chronology (AICC2012): 120-800 ka, *Climate of the Past*, 9, 1715–1731, 2013.
- Chappellaz, J., Stowasser, C., Blunier, T., Baslev-Clausen, D., Brook, E. J., Dallmayr, R., Faïn, X., Lee, J. E., Mitchell, L. E., Pascual, O., Romanini, D., Rosen, J., and Schüpbach, S.: High-resolution glacial and deglacial record of atmospheric methane by continuous-flow and laser spectrometer analysis along the NEEM ice core, *Clim. Past*, 9, 2579–2593, doi:[10.5194/cp-9-2579-2013](https://doi.org/10.5194/cp-9-2579-2013), 2013.
- 10 Dlugokencky, E., Myers, R., Lang, P., Masarie, K., Crotwell, A., Thoning, K., Hall, B., Elkins, J., and Steele, L.: Conversion of NOAA atmospheric dry air CH₄ mole fractions to a gravimetrically prepared standard scale, *Journal of Geophysical Research: Atmospheres*, 110, 2005.
- 15 Fourteau, K., Faïn, X., Martinerie, P., Landais, A., Ekaykin, A. A., Lipenkov, V. Y., and Chappellaz, J.: Analytical constraints on layered gas trapping and smoothing of atmospheric variability in ice under low-accumulation conditions, *Climate of the Past*, 13, 1815–1830, 2017.
- Liboutry, L.: A critical review of analytical approximate solutions for steady state velocities and temperature in cold ice-sheets, 1979.
- Loulergue, L., Schilt, A., Spahni, R., Masson-Delmotte, V., Blunier, T., Lemieux, B., Barnola, J.-M., Raynaud, D., Stocker, T. F., and Chappellaz, J.: Orbital and millennial-scale features of atmospheric CH₄ over the past 800,000 years, *Nature*, 453, 383–386, 2008.
- 20 Morville, J., Kassi, S., Chenevier, M., and Romanini, D.: Fast, low-noise, mode-by-mode, cavity-enhanced absorption spectroscopy by diode-laser self-locking, *Applied Physics B*, 80, 1027–1038, 2005.
- Parrenin, F. and Hindmarsh, R.: Influence of a non-uniform velocity field on isochrone geometry along a steady flowline of an ice sheet, *Journal of Glaciology*, 53, 612–622, 2007.
- Parrenin, F., Hindmarsh, R., and Rémy, F.: Analytical solutions for the effect of topography, accumulation rate and lateral flow divergence on isochrone layer geometry, *Journal of Glaciology*, 52, 191–202, 2006.
- 25 Parrenin, F., Bazin, L., Capron, E., Landais, A., Lemieux-Dudon, B., and Masson-Delmotte, V.: IceChrono1: a probabilistic model to compute a common and optimal chronology for several ice cores, *Geoscientific Model Development*, 8, 1473–1492, 2015.
- Rhodes, R. H., Faïn, X., Stowasser, C., Blunier, T., Chappellaz, J., McConnell, J. R., Romanini, D., Mitchell, L. E., and Brook, E. J.: Continuous methane measurements from a late Holocene Greenland ice core: Atmospheric and in-situ signals, *Earth and Planetary Science Letters*, 368, 9–19, 2013.
- 30 Rhodes, R. H., Brook, E. J., Chiang, J. C., Blunier, T., Maselli, O. J., McConnell, J. R., Romanini, D., and Severinghaus, J. P.: Enhanced tropical methane production in response to iceberg discharge in the North Atlantic, *Science*, 348, 1016–1019, 2015.
- Romanini, D., Chenevier, M., Kassi, S., Schmidt, M., Valant, C., Ramonet, M., Lopez, J., and Jost, H.-J.: Optical-feedback cavity-enhanced absorption: A compact spectrometer for real-time measurement of atmospheric methane, *Applied Physics B*, 83, 659–667, 2006.
- 35 Stowasser, C., Buizert, C., Gkinis, V., Chappellaz, J., Schuepbach, S., Bigler, M., Fain, X., Sperlich, P., Baumgartner, M., Schilt, A., et al.: Continuous measurements of methane mixing ratios from ice cores, 2012.
- 105 Wales, D. J. and Doye, J. P.: Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms, *The Journal of Physical Chemistry A*, 101, 5111–5116, 1997.

Timing and phasing tables

The following tables were originally presented in a spreadsheet file in the supplement to the Antarctic Temperature-CO₂ study. They show the results of sensitivity tests of the timings and phasings to different numbers of change points. Phasings are shown first, followed by timings. The calculated maximum and central probability intervals are approximations of the results of MCMC simulations, which are shown in the histograms in the paper.

ATS3 8 pts	95 % (lower)	68 % (lower)	Maximum	68 % (upper)	95 % (upper)
T1 Onset	-338	127	570	751	1045
ACR Onset	-357	-137	50	376	708
ACR End	-410	-211	-65	117	375
Holocene	45	337	532	629	773

Table A.1: Phasing, calculated using fits with 8 change points, between ATS3 and CO₂. Phasing is given in years. The columns represent central intervals of probability.

Only rapid rise, 8 pts.	95 % (lower)	68 % (lower)	Maximum	68 % (upper)	95 % (upper)
Holocene	-71	65	174	280	411

Table A.2: Phasing, at the Holocene between ATS3 and CO₂, if only the end of the rapid rise at the onset of the Holocene is taken to be the true CO₂ change point. Phasing is given in years. The columns represent central intervals of probability.

ATS3 7 pts	95 % (lower)	68 % (lower)	Maximum	68 % (upper)	95 % (upper)
T1 Onset	-226	247	609	802	1073
ACR Onset	-269	3	240	502	781
ACR End	-456	-230	-56	137	363
Holocene	85	345	520	623	763

Table A.3: Phasing, calculated using fits with 7 change points, between ATS3 and CO₂. Phasing is given in years. The columns represent central intervals of probability.

WD d18O 8 pts	95 % (lower)	68 % (lower)	Maximum	68 % (upper)	95 % (upper)
T1 Onset	-439	-36	210	387	636
ACR Onset	-378	-213	-150	231	502
ACR End	-277	-89	35	206	545
Holocene	79	585	805	854	949

Table A.4: Phasing, calculated using fits with 8 change points, between WD $\delta^{18}\text{O}$ and CO₂. Phasing is given in years. The columns represent central intervals of probability.

DF d18O 8 pts	95 % (lower)	68 % (lower)	Maximum	68 % (upper)	95 % (upper)
T1 Onset	-513	-2	480	729	1054
ACR Onset	-573	-246	48	345	706
ACR End	-464	-251	-108	209	490
Holocene	-293	-25	96	408	617

Table A.5: Phasing, calculated using fits with 8 change points, between DF $\delta^{18}\text{O}$ and CO_2 . Phasing is given in years. The columns represent central intervals of probability.

EDC d18O 8 pts	95 % (lower)	68 % (lower)	Maximum	68 % (upper)	95 % (upper)
T1 Onset	-308	136	460	665	953
ACR Onset	-384	-129	-20	344	643
ACR End	-483	-299	-175	-4	360
Holocene	71	331	492	575	697

Table A.6: Phasing, calculated using fits with 8 change points, between EDC $\delta^{18}\text{O}$ and CO_2 . Phasing is given in years. The columns represent central intervals of probability.

TD d18O 8 pts	95 % (lower)	68 % (lower)	Maximum	68 % (upper)	95 % (upper)
T1 Onset	-321	196	598	800	1087
ACR Onset	-336	-159	10	338	658
ACR End	-496	-307	-174	23	403
Holocene	70	332	560	630	762

Table A.7: Phasing, calculated using fits with 8 change points, between TD $\delta^{18}\text{O}$ and CO_2 . Phasing is given in years. The columns represent central intervals of probability.

EDML d18O 8 pts	95 % (lower)	68 % (lower)	Maximum	68 % (upper)	95 % (upper)
T1 Onset	-393	53	330	590	917
ACR Onset/ 16 ka ??	-1565	-866	-518	1068	1885
ACR End	-410	-204	-48	143	390
Holocene	-27	338	552	679	844

Table A.8: Phasing, calculated using fits with 8 change points, between EDML $\delta^{18}\text{O}$ and CO_2 . Phasing is given in years. The columns represent central intervals of probability.

8 pts ATS	-95	-68	Max	68	95
T1 onset	17467	17706	18093	18218	18426
ACR onset	14068	14215	14239	14522	14699
ACR end	12583	12691	12783	12904	13118
T1 end	11384	11568	11695	11760	11877

Table A.9: Timing, calculated using fits with 8 change points, of major changes in ATS. The columns represent central intervals of probability.

8 pts WD	-95	-68	Max	68	95
T1 onset	17525	17604	17690	17809	17971
ACR onset	14113	14204	14241	14305	14387
ACR end	12739	12838	12879	12959	13436
T1 end	11691	11817	11999	11986	11998

Table A.10: Timing, calculated using fits with 8 change points, of major changes in WD $\delta^{18}O$. The columns represent central intervals of probability.

8 pts DF	-95	-68	Max	68	95
T1 onset	17141	17609	17598	18223	18448
ACR onset	13789	14046	14354	14531	14751
ACR end	12522	12627	12690	13025	13242
T1 end	11028	11169	11226	11553	11748

Table A.11: Timing, calculated using fits with 8 change points, of major changes in DF $\delta^{18}O$. The columns represent central intervals of probability.

8 pts EDC	-95	-68	Max	68	95
T1 onset	17493	17770	17945	18118	18327
ACR onset	14044	14245	14384	14450	14587
ACR end	12535	12620	12668	12764	13265
T1 end	11485	11563	11586	11698	11783

Table A.12: Timing, calculated using fits with 8 change points, of major changes in EDC $\delta^{18}O$. The columns represent central intervals of probability.

8 pts TD	-95	-68	Max	68	95
T1 onset	17388	17824	18076	18279	18451
ACR onset	14142	14204	14229	14482	14628
ACR end	12518	12594	12664	12811	13255
T1 end	11395	11561	11723	11751	11872

Table A.13: Timing, calculated using fits with 8 change points, of major changes in TD $\delta^{18}\text{O}$. The columns represent central intervals of probability.

8 pts EDML	-95	-68	Max	68	95
T1 onset	17461	17678	17871	18052	18303
ACR onset/16 ka event ???	14281	14900	16013	15791	16387
ACR end	12572	12695	12757	12943	13125
T1 end	11176	11576	11675	11821	11954

Table A.14: Timing, calculated using fits with 8 change points, of major changes in EDML $\delta^{18}\text{O}$. The columns represent central intervals of probability.

8 pts CO ₂	-95	-68	Max	68	95
T1 onset	17168	17359	17582	17689	18098
ACR onset	13756	14019	14564	14447	14561
ACR end	12559	12728	12828	12962	13127
T1 end	11011	11072	11148	11266	11534

Table A.15: Timing, calculated using fits with 8 change points, of major changes in WD CO₂. The columns represent central intervals of probability.

7 pts ATS	-95	-68	Max	68	95
T1 onset	17578	17898	18137	18259	18436
ACR onset	14112	14228	14233	14531	14704
ACR end	12592	12698	12789	12902	13058
T1 end	11445	11570	11673	11753	11861

Table A.16: Timing, calculated using fits with 7 change points, of major changes in ATS3. The columns represent central intervals of probability.

7 pts CO2	-95	-68	Max	68	95
T1 onset	17185	17359	17575	17696	18200
ACR onset	13727	13892	14216	14320	14544
ACR end	12535	12690	12901	12995	13196
T1 end	11009	11070	11150	11267	11534

Table A.17: Timing, calculated using fits with 7 change points, of major changes in CO₂. The columns represent central intervals of probability.

7 pts WD	-95	-68	Max	68	95
T1 onset	17523	17617	17676	17875	18075
ACR onset	14103	14203	14305	14333	14386
ACR end	12770	12848	12884	12965	13437
T1 end	11757	11812	11999	11980	11997

Table A.18: Timing, calculated using fits with 7 change points, of major changes in WD $\delta^{18}\text{O}$. The columns represent central intervals of probability.

7 pts DF	-95	-68	Max	68	95
T1 onset	17452	17710	17620	18231	18436
ACR onset	13947	14152	14352	14525	14743
ACR end	12524	12619	12692	12953	13183
T1 end	11043	11176	11225	11526	11651

Table A.19: Timing, calculated using fits with 7 change points, of major changes in DF $\delta^{18}\text{O}$. The columns represent central intervals of probability.

7 pts EDC	-95	-68	Max	68	95
T1 onset	17546	17798	17943	18094	18267
ACR onset	14205	14339	14438	14496	14603
ACR end	12539	12631	12688	12786	13019
T1 end	11458	11558	11606	11700	11790

Table A.20: Timing, calculated using fits with 7 change points, of major changes in EDC $\delta^{18}\text{O}$. The columns represent central intervals of probability.

7 pts TD	-95	-68	Max	68	95
T1 onset	17561	17860	18112	18247	18439
ACR onset	14140	14196	14233	14406	14553
ACR end	12516	12584	12664	12770	12995
T1 end	11523	11632	11716	11762	11857

Table A.21: Timing, calculated using fits with 7 change points, of major changes in TD $\delta^{18}\text{O}$. The columns represent central intervals of probability.

7 pts EDML	-95	-68	Max	68	95
T1 onset	17577	17714	17838	17956	18145
ACR onset/16 ka event ???	14885	14959	15394	15623	16009
ACR end	12622	12729	12855	12914	13040
T1 end	11299	11602	11670	11806	11934

Table A.22: Timing, calculated using fits with 7 change points, of major changes in EDML $\delta^{18}\text{O}$. The columns represent central intervals of probability.

Appendix B

CO₂ during MIS 6

This study, lead by Jinhwa Shin, provides the first look at millennial-scale variations of CO₂ during the penultimate glacial period. As co-author, I worked on the improved chronology for MIS 6 and the statistical analysis used in the study. The resulting article has been submitted to *Climate of the Past*.



Millennial-scale atmospheric CO₂ variations during the Marine Isotope Stage 6 period (190–135 kyr BP)

Jinhwa Shin¹, Christoph Nehrbass-Ahles^{2,4}, Roberto Grilli¹, Jai Chowdhry Beeman¹, Frédéric Parrenin¹,
Grégory Teste¹, Amaelle Landais³, Loïc Schmidely², Jochen Schmitt², Thomas F. Stocker², Hubertus
5 Fischer² and Jérôme Chappellaz¹

¹CNRS, Univ. Grenoble-Alpes, Institut des Géosciences de l'Environnement (IGE), Grenoble, France

²Climate and Environmental Physics, Physics Institute, & Oeschger Centre for Climate Change Research, University of Bern, Switzerland

³Laboratoire des Sciences du Climat et de l'Environnement, LSCE/IPSL, CEA-CNRS-UVSQ, Université Paris-Saclay, 91191,
10 Gif-sur-Yvette, France

⁴Godwin Laboratory for Palaeoclimate Research, Department of Earth Sciences, University of Cambridge, Cambridge, UK

Correspondence to: Jérôme Chappellaz (jerome.chappellaz@univ-grenoble-alpes.fr)

Abstract. Understanding natural carbon cycle / climate feedbacks on various time scales is highly relevant to reliably predict
15 future climate changes. During the last two glacial periods, climate variations on millennial time scales were observed but the
background conditions and duration of climate variations are different. Here we make use of contrasting climatic boundary
conditions during the last two glacial periods to gain insight into the co-occurring carbon cycle changes. We reconstruct a new
high-resolution record of atmospheric CO₂ from the EPICA Dome C (EDC) ice core during Marine Isotope Stage (MIS) 6
(190 to 135 kyr BP). During long stadials in the North Atlantic (NA) region, atmospheric CO₂ appears to be associated with
20 the coeval Antarctic temperature changes at millennial time scale connected to the bipolar seesaw process. However, during
one short stadial in the NA, atmospheric CO₂ variation is negligible and the relationship between temperature variation in EDC
and atmospheric CO₂ is unclear. We suggest that the amplitude of CO₂ variation may be affected by the duration of
perturbations of the Atlantic Meridional Overturning Circulation (AMOC). In addition, similar to the last glacial period, in the
earliest MIS 6 (MIS 6e and 6d, corresponding to 189 to 169 kyr BP), Carbon Dioxide Maxima (CDM) show different lags
25 with respect to the corresponding abrupt CH₄ jumps, the latter reflecting rapid warming in the Northern Hemisphere (NH).
During MIS 6e at around 181.5±0.3 kyr BP, CDM 6e.2 lags abrupt warming in the NH by only 200±360 yrs. During MIS 6d
which corresponds to CDM 6d.1 (171.1±0.2 kyr BP) and CDM 6d.2 (175.4±0.4 kyr BP), the lag is much longer, i.e., 1,400±375
yrs on average. The timing of CO₂ variations with respect to abrupt warming in the NH may be affected by a major change in
the organization of the AMOC from MIS 6e to MIS 6d.



1 Introduction

Ice core studies allow us to considerably extend our knowledge about natural climate-carbon cycle feedbacks by directly reconstructing atmospheric CO₂ from gas preserved in Antarctic ice sheets (Lüthi et al., 2008; Petit et al., 1999). Comparing atmospheric CO₂ records from Antarctic ice cores with proxies of paleoclimate helps us to understand how atmospheric CO₂ was controlled by carbon exchange with the ocean and land reservoirs on orbital to centennial time scales (Ahn and Brook, 2008, 2014; Bereiter et al., 2012; Higgins et al., 2015; Lüthi et al., 2008; Marcott et al., 2014; Petit et al., 1999).

Previous works on polar ice core records revealed that temperature variations in Greenland and Antarctica on millennial time scales appear to be a pervasive feature during the last glacial period. While Antarctic temperature varied gradually, Greenland temperature changes occurred abruptly. A phase difference can be observed between millennial-scale variations of temperature in the NH and SH (Northern and Southern hemisphere, respectively), which is referred to as the bipolar seesaw phenomenon (Broecker, 1998; Stocker and Johnsen, 2003). Potential triggers for this climatic variability on the millennial scale are fresh water inputs into the North Atlantic (NA) or alterations of sea ice extent, surface temperature and salinity in NA (Bond et al., 1992; Broecker et al., 1992; Heinrich, 1988; McManus et al., 1998), which may reduce the strength of the Atlantic Meridional Overturning Circulation (AMOC). This would cause a reduction in heat transport from the SH to the NH, which leads to an abrupt cooling in the NH and a gradual warming in the SH (Stocker and Johnsen, 2003) and the opposite behaviour when AMOC is strengthened.

In response to the millennial temperature perturbations, existing CO₂ records show the presence of millennial-scale oscillations on the order of 20 ppm over the last glacial period (Ahn and Brook, 2008, 2014; Bereiter et al., 2012), which generally covary with the major water isotope (δD) variations in Antarctic ice cores reflecting Antarctic temperature variations (Ahn and Brook, 2008; Bereiter et al., 2012) (Figure 1). During cold periods in the NA, referred to as NA stadials, atmospheric CO₂ increased continuously and in parallel to Antarctic temperature increase. Again, at the onset of warming in Greenland, atmospheric CO₂ started to decrease (Ahn and Brook, 2008; Bereiter et al., 2012), generally in line with a co-occurring, slow Antarctic cooling. However, the CO₂ decrease did not always start at exactly the same time as the onset of the DO warming, and the lag itself varied. For example, during Marine Isotope Stage (MIS) 3, atmospheric CO₂ maxima lagged behind abrupt temperature change in Greenland by 870 ± 90 yrs. During MIS 5, the lag of atmospheric CO₂ maxima with respect to abrupt temperature warming in the NH was only about 250 ± 190 yrs (Bereiter et al., 2012). Atmospheric CO₂ variations on millennial timescales are thought to be related to the role of the Southern Ocean in carbon uptake and deep ocean ventilation on millennial timescales (Fischer et al., 2010; Marcott et al., 2014; Sigman and Boyle, 2000; Toggweiler et al., 2006). In addition, atmospheric CO₂ can be affected by changes in the AMOC, which affects the ventilation of carbon from the deep ocean (Denton et al., 2010; Sigman et al., 2007). However, the mechanisms responsible for these oscillations are still under debate, and therefore require further studies.

Comparing CO₂ changes on millennial time scales during the past two glacial periods, MIS 3 (MIS 3, 60–27 kyr BP) and early MIS 6 (early MIS 6, 185–160 kyr BP) can provide us with a better understanding of the carbon mechanisms at work, due to



the similarities but also differences of the climate changes during the last two glacial periods (see Figure S1 in SI (Supplement Information)). The bipolar see-saw mechanism has been observed to be active also during the early MIS 6 period (Cheng et al., 2016; Jouzel et al., 2007; Margari et al., 2010), however, climate boundary conditions during MIS 6 were slightly different from those during MIS 3. For example, events of iceberg discharge into the NA, which impact oceanic circulation, appear to be much more frequent during MIS 3 than during MIS 6 (Margari et al., 2010; Margari et al., 2014). During the time period around 175 kyr BP, iceberg discharge was muted; the intertropical convergence zone (ITCZ) is thought to have shifted northward, intensifying monsoon systems in low latitude regions, such as in Asia, the Appenine Peninsula and the Levant (Ayalon et al., 2002; Bard et al., 2002; Cheng et al., 2016), and leading to a generally weaker AMOC (Margari et al., 2010). The duration of NA stadials during the early MIS 6 is longer than that during MIS 3. In addition, the AMOC cell during MIS 6 might have been shallower than that during MIS 3 (Margari et al., 2010), which may have been caused by intensified hydrological cycles in low latitude regions (Ayalon et al., 2002; Bard et al., 2002; Cheng et al., 2016).

In order to investigate whether the different boundary conditions between MIS 6 and MIS 3 could have impacted the relationship between atmospheric CO₂ and climate, we reconstructed atmospheric CO₂ concentrations from the EPICA Dome C (EDC) ice core (75°06'S, 123°24'E) with 150 new data points spanning MIS 6 (190 to 135 kyr BP), significantly improving existing records previously obtained from the Vostok ice core (Petit et al., 1999). With our new reconstruction, the average temporal resolution during MIS 6 reaches 360 yr as compared to 1,000 yr in the Vostok record. Over the early time period of MIS 6, when millennial-scale variability is observed in the Antarctic climate record (185–160 kyr BP), a total of 100 new CO₂ measurements provide a temporal resolution of ~280 yr. We also improved the relative age uncertainties between ice and gas in the EDC core using $\delta^{15}\text{N}$ -based estimates of firn column thickness to better constrain leads and lags between δD composition in EDC and atmospheric CO₂ concentrations during the early MIS 6, and we established our new chronology using new $\delta^{15}\text{N}$ data during the early MIS 6 in this study and published data from Landais et al. (2013). Finally, we improved the temporal resolution of existing CH₄ data from EDC (Loulergue et al., 2008) from ~600 yr to ~350 yr to calculate the shift of Carbon Dioxide Maxima (CDM) relative to the rapid climate change in the NH. To avoid the age uncertainties between proxy data and atmospheric CO₂ data, CH₄ measurements were used as a time marker of rapid warming in the NH. Over the last glacial period, CH₄ and Greenland temperature were found to be essentially synchronous with a mean lag of CH₄ of not more than about 50 years (Baumgartner et al., 2014; Rosen et al., 2014).

2 Methods

2.1 CO₂ measurements

Atmospheric CO₂ was reconstructed using the Ball Mill dry-extraction system coupled to a gas chromatograph at the Institut des Géosciences de l'Environnement (IGE), France (Schaefer et al., 2011). Each data point presented in this study corresponds to a single 40 g ice sample which was measured five times by gas chromatography (five consecutive injections of the same



5 extracted gas). Approximately 5 mm of ice were trimmed from the ice core surfaces before extraction in order to remove the external part that could be potentially contaminated with drilling fluid or might have been subject to gas loss during storage in the freezer (Bereiter et al., 2009). The CO₂ measurements were referenced to a secondary gas standard (synthetic air from Air Liquide (Alphagaz 28416000)) containing 233.7±0.4 ppm of CO₂ in dry air, which was referenced to two primary standards (238.34±0.04 from NOAA (CB09707) and 260.26±0.2 from CSIRO (CSIRO1677)).

Blank tests using 40 g of artificial bubble-free ice were conducted every 10 measurements to quantify the precision of the system and to correct for the CO₂ contamination caused by the crushing process. Blank tests are conducted in two steps: First, to validate the baseline of the system, a gas standard with 233.7±0.4 ppm is injected over the bubble-free ice in the cell. The gas is then left equilibrating in the cell for 10 minutes. Then the gas is analysed twice by successive injections into the extraction
10 line and sample loop. Afterwards, the bubble-free ice is crushed and the gas is analysed 5 more times. The difference between the results before and after crushing was considered as the contamination effect caused by the crushing process. These values were used to estimate the precision of the system. Measured CO₂ should be corrected for contamination caused by the analytical procedure by comparing measured CO₂ in blank tests with the standard gas value. However, it is not feasible to correct CO₂ concentrations directly. The CO₂ mixing ratio is calculated as the ratio between partial pressure of CO₂ and total pressure in
15 the measurement line, which is a relative value. Thus, CO₂ concentration, and the concentration of CO₂ contamination, are dependent on the total pressure. To avoid this dependence, we correct the absolute value (partial pressure) of CO₂ in the air by the expected partial pressure of CO₂ contamination, as estimated from blank tests.

For this study, we used 4 different chambers to hold the ice core samples during crushing and measurement. Each chamber showed different contamination levels. Therefore, blank tests are conducted on each chamber. The data was corrected by the
20 average of each chamber. For these measurements, a precision of the system of at ~1 ppm on average was obtained. On average, the blank correction corresponds to a reduction of the measured CO₂ concentration by 1.7±1.0 ppm (1σ). In addition, the CO₂ record was corrected for gravitational fractionation, using the δ¹⁵N isotope ratio (Craig et al., 1988). To this end, 88 new data points together with existing δ¹⁵N measurements (Landais et al., 2013) covering the late MIS 6 (156.4–139.2 kyr) were used. δ¹⁵N data were linearly interpolated in age to each corresponding CO₂ data point. On average, the correction corresponds to
25 removing 1.2±0.1 ppm from the measured CO₂ concentration. Thus, in total, an average correction of 2.9±1.0 ppm was applied to the raw CO₂ data. 150 individual ice samples were measured for CO₂ in the depth range from 2036.7 to 1787.5 m along the EDC ice core, corresponding to the time period from 189.4 to 135.4 kyr BP on the AICC2012 chronology (Bazin et al., 2013), thus leading to a mean temporal resolution of 360 yr (280 yr over the time period 185–160 kyr BP).

2.2 CH₄ measurements

30 We measured the atmospheric CH₄ content of 63 ice core samples, using the wet extraction method at IGE described in detail in Spahni et al. (2005). This allowed us to improve the temporal resolution of existing CH₄ data (Loulergue et al., 2008) from ~600 yrs to 360 yrs on the AICC2012 chronology (see Figure S2 in SI).



The previous CH₄ dataset (Loulergue et al., 2008) from EDC was produced at both IGE and Climate and Environmental Physics (CEP), Physics Institute, University of Bern, Switzerland. A systematic offset of 6 ppb between IGE and CEP was observed (Loulergue et al., 2008). A new data correction was applied to data measured at IGE in this study, and the systematic offset between the existing (corrected) data points and their nearest neighbours in the new data set is 3.0±4.6 ppb (n=11). In addition, the mean difference between the new data and the existing data points (Loulergue et al., 2008) is 1.7±2.4 ppb (n=63, the standard error of the mean) during MIS 6, which is within data error. Thus, in this study, we did not take this systematic offset between the CH₄ records into account or revise the Bern/IGE offset correction of the previous data.

2.3 Nitrogen isotopes

Isotopes of molecular nitrogen in air bubbles were measured by a melting technique at the Laboratoire des Sciences du Climat et de l'Environnement (LSCE), France. The gas was extracted from the ice by a wet extraction technique and the released air was analysed by a dual-inlet mass spectrometer (Delta V Plus; Thermo Scientific). The analytical method and data correction are described in detail in Bréant et al. (2019). In total, 151 samples from 88 depth intervals (63 duplicates) between the depths of 2124.7 and 1875.0 m below the surface were measured, corresponding to 205 to 154 kyr BP (Figure S3 in SI). The average resolution on the AICC2012 chronology is ~580 yrs.

2.4 Ice age revision by estimating Δdepth from δ¹⁵N in air of ice core

The water isotopic signature (δD), is, unlike CO₂, measured on the ice matrix. Air enclosed in an ice core moves through the porous firn layers at the top of the ice sheet by molecular diffusion, and becomes trapped in the ice at the so-called Lock-In Depth (LID, around 100 m below the surface in the case of the EDC ice core). An age difference thus exists between the ice and the air at a given depth. For the conditions of the EDC ice core, the age difference can reach 5 kyr and is associated with a large uncertainty (several hundred years (Loulergue et al., 2007)). δ¹⁵N of molecular nitrogen in air bubbles can be used to determine the LID and to calculate the depth difference between synchronous events in the ice matrix and air bubbles, called delta depth (Δdepth), thus creating a more precise relative chronology of the gas- with respect to the ice-phase (Parrenin et al., 2013). We use the Δdepth calculation to adjust the gas chronology of EDC, while the ice chronology from AICC2012 remained unchanged.

We calculate the height of the firn column, h , from δ¹⁵N of N₂ measurements (Craig et al., 1988; Dreyfus et al., 2010; Sowers et al., 1989), by using the following Eq.:

$$h = h_{\text{conv}} + (\delta^{15}\text{N} - \Omega(T) \Delta T_{\text{diff}}) \left(\frac{\Delta m \cdot g \cdot 1000}{RT} \right)^{-1}, \quad (1)$$

In this equation, ΔT_{diff} is the temperature difference between the top and the bottom of the diffusive zone as estimated by the Goujon/Arnaud model, where surface temperature and accumulation are estimated from the stable water isotope record (Loulergue et al., 2007). $\Omega(T)$ is the thermal diffusion sensitivity, which has been estimated from laboratory measurements by Grachev and Severinghaus (2003). Δm is the mass difference between ¹⁴N and ¹⁵N (kg mol⁻¹), g is the gravitational



acceleration (9.825 m s^{-2} for Antarctica) (Parrenin et al., 2013), and R is the universal gas constant ($8.314 \text{ J mol}^{-1} \text{ K}^{-1}$). Finally, h_{conv} is the height of the convective zone at the top of the firm column, which is considered negligible at EDC according to current observations (Landais et al., 2006). The variation of the convective height is related to changes in wind stress. According to Krinner et al. (2000), wind on the East Antarctic plateau varied little during the LGM (last glacial maximum), which is analogous to late MIS 6, and the convective zone was confirmed to be very thin during the last deglaciation by Parrenin et al. (2012). Thus, we assume that h_{conv} is negligible during MIS 6. Δdepth is calculated from the height of the air column using a constant firm average density and a modelled vertical thinning function suggested by Parrenin et al. (2013). Raw $\delta^{15}\text{N}$ data cannot be used directly to calculate Δdepth , because bubbles are not trapped at exactly the same time for a given depth in the ice sheet, leading to cm-scale age scale inversions (Foureaux et al., 2017). Although the layer inversions should not strongly affect our record as its resolution is larger than the scale of inversion events (Foureaux et al., 2017), the change in Δdepth between two different depths (z) in the ice core, denoted $\partial\Delta\text{depth}/\partial z$, deduced from the raw $\delta^{15}\text{N}$ data, shows values higher than 1, that could correspond to stratigraphic inversions in the gas phase (Figure 2). Therefore, a 3-point moving average weighted by the time difference between a point and its two neighbours was applied to the $\delta^{15}\text{N}$ dataset. The weights for the three points are equal if the time difference is less than or equal to 500 years, which is close to the average sampling resolution of the $\delta^{15}\text{N}$ dataset. Otherwise, the neighbouring data points are weighted by $500/\Delta T$, where ΔT is the time difference. This avoids assigning too much weight to neighbouring points where the resolution in the record is lower, which would lead to local over-smoothing.

Δdepth as estimated using the 3-point moving average weighted by the time difference is shown in Figure 2. The difference of the Lock-In Depth in Ice Equivalent (LIDIE) calculated in AICC2012 (Bazin et al., 2013) and deduced from $\delta^{15}\text{N}$ in this study is $0.5 \pm 3.0 \text{ m}$, or $0.7 \pm 5.6 \%$ on average (see Figure S4 in SI). The AICC2012 LIDIE was calculated using a background scenario derived from δD , using a linear relationship between δD and $\delta^{15}\text{N}$. Our results show this relationship to be relatively unbiased but not entirely exact (Figure 2). The Δdepth values are used to update the EDC gas chronology from the original AICC2012 one, while the ice chronology remains unchanged. For MIS 6, this method significantly reduces the relative age uncertainty between air and ice to 900 yrs on average with respect to the original AICC2012 chronology (see Figure S5 in SI).

2.5 Age scale of the MD01–2444 marine sediment record

The MD01–2444 Marine sediment core from the Iberian margin (Margari et al., 2010) is an important proxy for the interpretation of our CO_2 record, as it is well-resolved during MIS 6, and it provides high-resolution benthic and planktonic foraminiferal records. In order to make optimal use of the record, its relative chronology with respect to EDC requires additional tuning. The original MD01–2444 age scale was built by matching the benthic $\delta^{18}\text{O}$ record (Shackleton et al., 2000) to the δD record from EDC on the EDC03 ice age scale (Parrenin et al., 2007) and using 11 tie points selected by Margari et al. (2010) as shown in Table 1. The age of the sediment record was linearly interpolated between the tie points, and the age of the tie points recalculated using the AICC2012 chronology.



2.6 Definition of NA stadial duration

Due to the absence of a Greenland temperature record for MIS 6, the durations of the six NA stadials identified during the MIS 6 period were previously defined by Margari et al. (2010) using the interval between the stadial transitions. Between the maximum and the preceding minimum of δD in the EDC record is defined as the stadial transitions. However, this approach
5 relies on the assumption that the bipolar seesaw was present during MIS 6. Here we instead use another approach which was also proposed by Margari et al. (2010).

The durations of the six NA stadials were defined using $\delta^{18}O$ of planktonic foraminifera and tree pollen in MD01–2444, which reflect temperature variability in the NH (Margari et al., 2010). The interval between the maximum and the preceding minimum of both $\delta^{18}O$ of planktonic foraminifera and tree pollen in MD01–2444 were defined as the inflection points. The time interval
10 between two inflection points were defined as the NA stadial duration. In this approach, small variations of the two records may bias the calculation of the duration of short stadials in the NH. However, the average age difference between the durations identified using the two methods is only 205 years, which is less than the sampling resolution of MD01–2444 during MIS 6. The stadials identified for MIS 6 are shown in Figure 3 and Table 2. The uncertainty of the duration of each stadial was estimated as half of the temporal difference between maxima and minima of $\delta^{18}O$ of planktonic foraminifera.

15 3 Results

3.1 The new high-resolution and high precision CO_2 record during MIS 6

Figure 4 shows the new atmospheric CO_2 record from EDC during MIS 6, compared to the existing CO_2 data from Vostok (Petit et al., 1999). Both were measured using the ball mill extraction system at IGE, and both are reported on the AICC2012 air age scale (Bazin et al., 2013). The error bars indicate the standard deviation of five consecutive injections of the gas
20 extracted from each sample into the gas chromatography added to the precision of the measurement estimated by the reproducibility of the blank tests (~ 0.8 ppm) using a quadratic sum. The average standard deviation is ~ 1.1 ppm for EDC (this study) and ~ 3 ppm for the Vostok dataset (Petit et al., 1999).

The lowest atmospheric CO_2 concentration observed in our new MIS 6 dataset is 179.5 ± 0.9 ppm at approximately 156.3 ± 0.3 kyr BP, and the highest concentration is 211.7 ± 0.3 ppm at 181.3 ± 0.2 kyr BP (Figure 4). Atmospheric CO_2 variability in the
25 EDC ice core shows similar general patterns as the one retrieved from the Vostok ice core, although new features appear thanks to the improved temporal resolution. However, CO_2 concentrations from Vostok appear systematically higher than those from EDC by 4.0 ± 5.6 ppm on average and the difference grows to more than 10 ppm at the beginning of Termination 2 at around 135 kyr BP. Two main arguments may explain this offset: contamination due to the extraction system was assumed as being negligible at the time of the Vostok sample measurements at IGE (Petit et al., 1999), while from the measurements conducted
30 during this study it possibly amounts on average to an additional reduction of ~ 1.7 ppm. Some of the differences in the two records might also be related to age scale uncertainties (Figure 4), due to the limited number of stratigraphic tie points between



the two cores (Bazin et al., 2013). Lastly, the atmospheric CO₂ measurements from the Vostok ice core are affected by a larger analytical uncertainty than the EDC ones. Here we constrain our interpretation to millennial changes in our new CO₂ record, and our conclusions are largely independent of the absolute CO₂ level.

3.2 Relationship between EDC δD and atmospheric CO₂

5 Margari et al. (2010) suggested that MIS 6 can be divided into three sections depending on the degree of climatic variability observed in δD (indicative of Antarctic climate variability) and CH₄ (which may reflect NA climate variability) in EDC : early (185.2–157.7 kyr BP), transition (157.7–151 kyr BP) and late MIS 6 (151–135 kyr BP) (Figure 5 and Figure 6). Climatic oscillations on millennial time scales are pervasive during the early MIS 6 period (185–160 kyr BP) (Barker et al., 2011; Cheng et al., 2016; Jouzel et al., 2007; Margari et al., 2010; Margari et al., 2014), which is similar to MIS 3 (Figure 1, 5 and
10 Figure S11 in SI). However, during the late MIS 6 period / penultimate glacial maximum (PGM), millennial variability is subdued, and resembles climate variability on millennial time scales during MIS 2 (Figure 1 and Figure 6). During the transitional period from 157–151 kyr BP, δD in EDC slowly increased (Jouzel et al., 2007). Like δD in EDC, CO₂ variations on millennial time scales are pervasive during the early MIS 6 period (185–157.7 kyr BP). During the transitional period from 157–151 kyr BP, atmospheric CO₂ increased slowly, while, during the late MIS 6 period, CO₂ variation is subdued.
15 Between these two periods, we can identify one low-amplitude CO₂ peak at around 150 kyr BP, which can be another potential CDM. This atmospheric CO₂ variation is of triangular shape, which follows the δD pattern. The change of direction is also associated with a CH₄ peak. This variation is similar to MIS 4 and MIS 10 (Barker et al., 2011; Nehrbass-Ahles et al., in review).

For the full MIS 6 interval, the new EDC CO₂ record and δD in EDC are synchronous within age uncertainty (160±900 yrs
20 using a maximum lag correlation coefficient estimate, and taking into account the relative age uncertainty between ice and air).

3.3 Atmospheric CO₂ variability on millennial time scale

While there was indication of millennial CO₂ variability in the Vostok record, we now present clear evidence of millennial variability of CO₂ concentrations during MIS 6 that are associated with AIM events, thanks to the improved time resolution
25 and precision of the obtained CO₂ data (Figure 5). To better discuss millennial variability in MIS 6, a Savitsky Golay filter with a 500 yr cut-off period was selected to filter out centennial-scale variability and noise (see Figure 5, Figure S10 and Table S1 in SI). Five prominent and one subdued CO₂ variations were detected in atmospheric CO₂ during early MIS 6 (Figure 5). The five prominent peaks are observed at 160.6±0.3 (CDM 6c.1), 164.2±0.3 (CDM 6c.2), 169.6±0.2 (CDM 6d.1), 174.1±0.2 (CDM 6d.2) and 181.3±0.2 (CDM 6e.2) kyr BP (the errors given here reflect the uncertainty with respect to the
30 position of each maximum, and do not include the age uncertainty of the ice core, in each case around 3.0 kyr, 1 sigma). Each CDM is associated with an Antarctic Isotope Maximum (AIM) event. The short AIM 6e.1 event corresponds to CDM 6e.1 at around 178 kyr BP (with an age uncertainty of 3.0 kyr, 1 sigma), whose amplitude was estimated by a single data point of



atmospheric CO₂ maximum because CO₂ variations becomes even less pronounced after filtering. CDM 6e.1 has an amplitude of only ~4 ppm. If this point is excluded, atmospheric CO₂ and δD composition in EDC appear completely decoupled, though it is likely that the amplitude of the CDM would have been reduced by smoothing during the gas trapping process (Fourteau et al., 2017). This conclusion seems confirmed when considering the relationship between atmospheric CO₂ change and the duration of NA stadials calculated using tree pollen and the δ¹⁸O composition of planktic foraminifera in a Iberian margin core (Margari et al., 2010) for MIS 6 (Figure 3 and Table 2), and using isotopic records from Greenland ice cores (Rasmussen et al., 2014) for MIS 3 (further details can be found in the SI), as shown in Figure 7. The magnitude of atmospheric CO₂ change is generally correlated with the NA stadial duration ($r=0.7$, $n=6$) during the early MIS 6 period.

We note a similar correlation between the NA stadial duration and atmospheric CO₂ change during MIS 3 ($r=0.85$, $n=14$). When the NA stadial duration was shorter than 1,500 yrs, atmospheric CO₂ varied less than 5 ppm (Ahn and Brook, 2014; Bereiter et al., 2012) as is the case for CDM 6e.1. Margari et al. (2010) note one exception, AIM 14 during which ice discharge may have led to a stronger perturbation to AMOC.

Both Bereiter et al. (2012) and Ahn and Brook (2014) observe that during short NA stadials which last less than 1,500 yrs, the CO₂ maxima do not appear to have a consistent phase relationship with AIMs and CO₂ and δD anomalies are not correlated (Ahn and Brook, 2014). On the other hand, in both MIS 3 and 6 periods, CO₂ is highly correlated with δD_{ice} anomalies during the long stadials ($r = 0.84$ on average) with a stronger increase in CO₂ (Ahn and Brook, 2014; Bereiter et al., 2012).

We observe that during the last two glacial periods, the amplitude of CO₂ is highly determined by the NA stadial duration ($r=0.83$, $n=20$). This observation implies that despite the different climate boundary conditions during the last two glacials, the behavior of atmospheric CO₂ was similar (see Figure S11 in SI).

3.4 Leads and lags between CO₂ and the abrupt warming in NH

To better understand the link between the bipolar see-saw mechanism and atmospheric CO₂ variability on millennial time scales, we calculated the varying time lag for each CDM following abrupt warming events in the NH (see Figure S6–S9 in SI) (Bereiter et al., 2012). Due to the lack of a direct temperature proxy record in Greenland covering MIS 6, and due to the fact that the δ¹⁸O composition of planktic foraminifera in NA sediment cores, which could be used in principle as an indicator for abrupt warming in the NH (Shackleton et al., 2000), cannot be placed with sufficient precision on a common chronology with the EDC ice core, in this work CH₄ measurements performed on the EDC ice core were used as a time marker of rapid warming in the NH (Baumgartner et al., 2014; Brook et al., 1996; Huber et al., 2006). Because CH₄ and CO₂ signals are both imprinted in the air bubbles, there is no chronological uncertainty when comparing the timing of changes of those two signals. The only remaining uncertainty is related to analytical uncertainties and to the temporal resolution of the two records. We pick intervals when CH₄ increases rapidly by at least 50 ppb over a time period of less than 1 kyr that correspond to Antarctic isotope maxima (Buizert et al., 2015; Louergue et al., 2008). The timing of abrupt CH₄ increases was defined as the midpoint between the beginning of the increase of CH₄ and its maximum. The age uncertainty of the midpoint is defined by the time difference between the midpoint and either of the two endpoints.



Figure 8 shows the shifts of CDM with respect to the onset of the abrupt warming in the NH. During the MIS 6 period, three abrupt NH warmings (as inferred from the CH₄ signal) at 181.5 ± 0.3 , 175.4 ± 0.40 and 171.1 ± 0.2 kyr BP (2σ) were found. These events correspond to CDM 6e.2, CDM 6d.2 and CDM 6d.1, respectively. CDM 6c.1, CDM 6c.2 and CDM 6e.1 do not have corresponding rapid changes in the methane record; this may be due to slow gas trapping as compared to interglacial periods, which could completely smooth out smaller changes. A synthetic Greenland temperature record (Barker et al., 2011) shows abrupt temperature jumps at CDM 6c.1, CDM 6c.2 and CDM 6e.1 as well. However, this record is calculated using EDC δD_{ice} , and the large chronological uncertainty (~ 900 yrs on average) between ice and gas phases does not allow us to make any conclusions about leads and lags using this record, since these tend to be on the order of 1 kyr. We therefore exclude these events.

From MIS 6e to MIS 6d, the lag of CO₂ with respect to abrupt warmings in the NH, which were identified from this chronological comparison between EDC CH₄ and CO₂, becomes larger. During the earliest MIS 6, atmospheric CO₂ increases rapidly (by ~ 4.2 ppm in 200 ± 360 yrs) following the abrupt CH₄ increase at 181.5 ± 0.3 kyr BP. The peak of CDM 6e.2 is nearly synchronous with the onset of the NH abrupt warming (non-significant lag of 200 ± 360 yrs, Figure 8). During MIS 6d which corresponds to CDM 6d.1 and 6d.2, CO₂ concentrations show a much slower increase over a duration of ~ 3.3 kyr. Here, CO₂ lags behind the onset of the NH abrupt warming by $1,500 \pm 280$ yrs and $1,300 \pm 450$ yrs, respectively ($1,400 \pm 375$ yrs on average, with the error calculated by propagation of the uncertainties in the individual events). Interestingly, these two CDM events occurred during MIS 6d, when iceberg discharge was muted and the intertropical convergence zone (ITCZ) is thought to have shifted northward, intensifying monsoon systems in low latitude regions, such as in Asia, the Appenine Peninsula and the Levant (Ayalon et al., 2002; Bard et al., 2002; Cheng et al., 2016), and leading to a generally weaker AMOC (Margari et al., 2010). Therefore, the two apparent CO₂ lag timescales with respect to abrupt warming in NH during MIS 6 might be explained by this difference in background climate conditions.

Indeed, these features during MIS 6 appear fully compatible with those observed during the last glacial period. To compare CO₂ variations with respect to abrupt warming in NH during the last two glacial periods, in this study we also re-estimated the abrupt CH₄ jump during the last glacial period with the same tool. Figure 8 shows the CO₂ evolution during the onset of abrupt warming in the NH during MIS 3, and MIS 5. This figure shows the whole sequence of rapid events of the last glacial period (Ahn and Brook, 2014; Bereiter et al., 2012). Atmospheric CO₂ during MIS 3, as shown in Figure 8, was reconstructed from the Talos Dome ice core (TALDICE). For MIS 5 it was obtained from Byrd and the EPICA Dronning Maud Land (EDML) ice core (Ahn and Brook, 2008; Bereiter et al., 2012). In Bereiter et al. (2012), both TALDICE and EDML records were used during MIS 3 and compared to the onset of abrupt warming in the NH. However here, we only use data from TALDICE which are more accurate due to the narrower age distribution of the gas trapped in the LID (Bereiter et al., 2012). Using the same method, the average value of CDM lag with respect to the abrupt warming in NH was calculated. The average CDM lags with respect to the abrupt warming in the NH for the MIS 3 and 5 periods are 770 ± 180 and 280 ± 240 yr (Bereiter et al., 2012). Thus, over the course of the last glaciation, the lag of CO₂ maxima with respect to the abrupt NH warming events significantly increased. We observe the same trend through the millennial events depicted during MIS 6, albeit with different absolute lags.



While the earlier CO₂ maximum corresponding to CDM 6e.2 during MIS 6 shows no significant lag (thus close to the small lag of 190±110 yr observed during MIS 5), the lag between the abrupt warming in NH and CDM 6d.1 and 6d.2 is longer than the lags of CDMs during MIS 3.

4 Discussion

5 4.1 Atmospheric CO₂ variability on millennial time scale

We found that the amplitude of atmospheric CO₂ variations is well related to the NA stadial duration during MIS 6, which implies that the amplitude of CO₂ variations might be affected by the duration of AMOC disruption during the early MIS 6 period (Margari et al., 2010). This hypothesis is also supported by a recent study using oceanic sediment cores from the Southern Ocean (Gottschalk et al., 2019b). The authors report that respired carbon levels in the deep South Atlantic decrease
10 when AMOC is weakened during both glacial periods, and the amount of carbon loss in the deep South Atlantic is highly correlated with the duration of NA stadials.

The hypothesis that perturbations of the AMOC might ultimately lead to changes in atmospheric CO₂ concentrations during MIS 6 is supported by numerical simulations (Menviel et al., 2014; Schmittner and Galbraith, 2008). Menviel et al. (2014) report that when large amounts of low-density fresh water are released into the NA, the AMOC is shut down, which strongly
15 reduces heat transport from the south to the north, and moisture transport from the Atlantic to the Pacific is reduced (Leduc et al., 2007; Peterson et al., 2000; Richter and Xie, 2010). In response, sea ice retreats in the Southern Ocean. On the other hand, the NA is cooled, and the ITCZ is shifted to the south and Summer Monsoon intensity in East Asia is reduced (Wang et al., 2008). Changes in precipitation connected to ITCZ shifts (mainly in tropical Africa and northern South America) may cause variations in terrestrial carbon stocks (Bozbiyik et al., 2011; Köhler et al., 2005; Menviel et al., 2008; Obata, 2007). Due to the
20 reduction of Summer Monsoon intensity in East Asia, salinity at the surface of the Pacific Ocean is increased. Thus, AABW and North Pacific Deep Water (NPDW) transport is enhanced (Menviel et al., 2014). Enhanced NPDW transport ventilates deep Pacific carbon via the Southern Ocean which may lead to atmospheric CO₂ increases.

As mentioned above, atmospheric CO₂ on millennial timescales can be controlled by CO₂ exchange between the ocean and the atmosphere, as well as changes of terrestrial carbon stocks. Coupled climate carbon cycle models reported that the
25 variations of atmospheric CO₂ concentration on millennial timescales are mainly dominated by deep ocean inventory, requiring a few millennia to react to climate change (Schmittner and Galbraith, 2008). These variations of CO₂ concentration might be compensated by fast changes in the terrestrial biosphere (Bouttes et al., 2012; Menviel et al., 2014; Schmittner and Galbraith, 2008). The initial response of the terrestrial biosphere and deep ocean to AMOC perturbations are opposite in the CLIMBER-2 model (Bouttes et al., 2012).

30 Based on the comprehensive review of Gottschalk et al. (2019a), some models show a CO₂ increase during stadial conditions and others show a decrease, with a preference for those showing an increase. During short stadials, a change in the carbon stock of the terrestrial biosphere can compensate the slow response of the deep Ocean (Gottschalk et al., 2019a; Menviel et



al., 2014; Schmittner and Galbraith, 2008). This might explain our observation that atmospheric CO₂ variation is rather muted during short stadials, even though other SH climate proxies like dust flux and δD composition in EDC show significant variations (Figure 1). However, during long stadials, modelling studies indicate that the deep ocean inventory continues to respond to climate changes on timescales much longer than the response time of the terrestrial biosphere (on the order of a few centuries), and atmospheric CO₂ varies significantly (Gottschalk et al., 2019a).

Another possible reason for the difference between CO₂ changes during short and long stadials may be related to a stronger reduction of the NADW during long stadials (Henry et al., 2016; Margari et al., 2010), which would cause a stronger upwelling of deep water in the Southern Ocean (Menviel et al., 2008; Schmittner et al., 2007). When large amounts of low-density fresh water are released into the NA, NADW formation can be slowed down. These events may reduce stratification in the Southern Ocean due to an increase in salinity of the surface waters and a relative freshening of the deep water (Schmittner et al., 2007). As a result, atmospheric CO₂ can be increased due to upwelling and outgassing of CO₂ in the Southern Ocean (Schmittner et al., 2007). The co-occurring upwelling in the SO during AIMs for the last termination has been examined (Anderson et al., 2009) but, due to the lack of proxy data with precise age scale for upwelling in the Southern Ocean, this hypothesis cannot be confirmed during MIS 6. During the short stadial in MIS 6 (AIM 6e.1) and the short stadials in MIS 3, the duration and strength of AMOC disruption are similar. This is supported by the marine proxy data for upwelling in the Southern Ocean which do not show strong variations during short stadials for both MIS periods (Anderson et al., 2009).

According to the model results of Margari et al. (2010), the six AIM events of MIS 6 were likely affected by AMOC perturbations of similar strength. During these events, there is no clear evidence for freshwater perturbation in the NA (Figure S12 in SI), and the strength of the associated AMOC perturbations is estimated to be similar to that during non-Heinrich (no ice discharge) or non-classic Heinrich (different source signature for IRD, occurring at transitions in ice volume) AIM events of MIS 3. The durations of NA stadials during early MIS 6 (except for AIM 6e.1) appear to be longer than those during non-Heinrich AIM events of MIS 3, which might be caused by the different boundary conditions during MIS 3 and the increase of hydrological cycle during MIS 6.e. A longer duration of the AMOC disruption may require more time to recover AMOC strength, which may impact the amplitude of CO₂ variations (Bouttes et al., 2012; Menviel et al., 2008). Considering that the timescale of the AMOC recovery may be affected by climate background conditions (Bouttes et al., 2012; Menviel et al., 2008), this observation suggests how different climate background conditions may impact atmospheric CO₂.

The strength of AMOC perturbations appears to also be an important factor in determining the amplitude of CO₂ variations. For example, the duration of the Heinrich events in MIS 3 (AIM 8, 12 and 14) is shorter than any of the MIS 6 events except for 6e.1, but atmospheric CO₂ varied significantly in all three.

The relationship between the amplitude of atmospheric CO₂ variations and the NA stadial duration is explained by the duration of AMOC disruption during the early MIS 6 period. However, the temporal resolution of δ¹⁸O composition of planktonic foraminifera in MD01–2444 and the precision of the age scale were too low to precisely define the duration of stadials during MIS 6. Additional proxy data providing information about climate change in the NH are needed to confirm the relationship between atmospheric CO₂ variations and the NA stadial duration (Figure 3).



The limited available proxy data permit only an exploratory discussion of the mechanisms responsible for CO₂ variability during MIS 6. To compare the behaviour of the bipolar see saw with atmospheric CO₂ variations, additional investigations about AMOC disturbances and their associated climate responses are needed.

4.2 Why did CO₂ lag the abrupt warming in the NH during MIS 6d?

5 Two different lags of CO₂ variations with respect to NH warming are present in the MIS 6 period (Figure 8). CDM 6e.2 is nearly synchronous with the abrupt warming in the NH (no significant lag of 200 ± 360 yrs), while the lags for CDM 6d.2 (1,300±450 yrs) and CDM 6d.1 (1,500±280 yrs) are much longer. Two modes of CO₂ variations are also observed during the last glacial period. As the last glaciation progressed from MIS 5 to MIS 3 (Figure 8), the lag of CO₂ maxima with respect to
10 MIS 5 and MIS 3 (Bereiter et al., 2012). As observed during the last glacial period, a mode change of oceanic circulation between MIS 6e to MIS 6d might be also the cause of the change in the time lags between NH abrupt warming events and CO₂ variations during the early MIS 6.

During MIS 3, the oceanic circulation in the Atlantic was in a “glacial” state, with shallower NADW and carbon-rich AABW extended to the north, while during MIS 5 circulation was similar to the present, in what can be referred to as a “modern-like”
15 state. At the onset of Dansgaard–Oeschger (DO) events, AMOC is thought to accelerate rapidly, delivering heat to the north and resuming the formation of NADW. When the NADW cell expands, AABW is withdrawn and the upwelling of carbon-rich deep water in the Southern Ocean is enhanced. Essentially, over the time scale of ocean overturn, part of the previously expanded carbon-rich southern sourced water is converted to carbon-poor northern sourced water.

Therefore, due to the northward expansion of AABW before the onset of DO events in the MIS 3 period, additional carbon
20 may have been added to the atmosphere by deep water exchange from carbon enriched AABW to the NADW. Thus, CO₂ continues to be released into the atmosphere for another 500 to 1,000 years while NADW formation continues (Bereiter et al., 2012). However, during MIS 5 where AABW was not expanded in the NA, less CO₂ is thought to be released from the ocean to the atmosphere.

The temporal resolution of proxy data related to oceanic circulation during MIS 3 and 5 is unfortunately not sufficient to
25 validate from the marine realm itself whether the two different modes of CO₂ variations reflect the hypothesized mechanism described above. Modelling studies of the carbon stock in AABW and NADW during MIS 5 and 3 have been attempted. However, dependent on the chosen model, the modes of atmospheric CO₂ variation are different (Gottschalk et al., 2019a). Some studies, for example, Menviel et al. (2008) generate an atmospheric CO₂ decrease when the AMOC is reduced. Others, for example Bouttes et al. (2012) confirm an atmospheric CO₂ release from the ocean when the AMOC resumed.

30 It is possible that a similar change in the AMOC may explain the presence of two different lags in MIS 6. The density of the NA surface water was low during 180–168 kyr BP due to the higher intensity of the low latitude hydrological cycle (Ayalon et al., 2002; Bard et al., 2002; Cheng et al., 2016; Mélières et al., 1997). From MIS 6e to MIS 6d, NADW became shallower and the lag of CO₂ maxima with respect to NH millennial-scale warming significantly increased. Moreover, the AMOC cell



appears to have been generally shallower during the early MIS 6 (180–168 kyr BP) relative to MIS 3 (Margari et al., 2010), and so Bereiter et al. (2012)’s hypothesis would indicate an extended CO₂ release from an expanded AABW reservoir, potentially explaining the millennial–scale delays with respect to abrupt NH warming events. Wilson et al. (2015) show that circulation conditions in the Atlantic Ocean during the earliest MIS 6 period were similar to those during MIS 5, with deeper NADW and AMOC than during MIS 6d. This may explain the shorter lag between abrupt NH warming and CDM 6e.2. However, in order to explain the much longer lags of CDM 6d.1 and 6d.2, NADW should be even shallower than in MIS 3. During MIS 6d, the lag between the abrupt warming in NH and CDMs (1,400±375 yrs on average) is longer than the lags of CDMs (770±180 years on average) during MIS 3, which may be related to the observation that the NADW during MIS 6.d might have been shallower than that during MIS 3 (Margari et al., 2010). However, our study has lower temporal resolution compared to CO₂ data set during MIS 3. To confirm the larger lag of CO₂ with respect to abrupt warming in NH, additional CO₂ reconstructions during MIS 6.d are needed.

To further investigate the applicability of the shallow–AMOC glacial hypothesis to the CO₂ variations observed in this study, higher resolution proxy data to estimate variations in the strength of the AMOC during the MIS 6 period are needed. However, as is the case for proxy data during the last glacial period as well, the temporal resolution is unfortunately not sufficient to quantify the lag between AMOC resumption and temperature increase at the surface during DO warmings (Henry et al., 2016). In addition, because of the low accumulation at EDC, the estimation of the exact timing of CDM from the EDC ice core might be less accurate compared to that from the TALDICE ice core, for example, due to the narrower gas age distribution of TALDICE (Bereiter et al., 2012). To further investigate the exact relationship between CDM and abrupt warming in the NH, additional CO₂ measurements from a higher accumulation site could be helpful.

20 5 Conclusion

We measured 150 samples of the EDC ice core to reconstruct atmospheric CO₂ during the MIS 6 period (189–135 kyr BP), with an unprecedented time resolution and analytical precision. Five millennial-scale atmospheric CO₂ changes are revealed during the early part of MIS 6 (189–160 kyr BP), with amplitudes ranging between 15 to 25 ppm, mimicking similar trends in Antarctic δD variations. During the shortest stadials in the NA, atmospheric CO₂ variations are negligible and decoupled with δD in EDC, probably because the duration of upwelling in the Southern Ocean was not sufficient to impact atmospheric CO₂, in line with Ahn and Brook (2014). In the earliest MIS 6 (MIS 6e and 6d, corresponding to 189 to 169 kyr BP), a change of CO₂ lags with respect to NH warming – as deduced from atmospheric CH₄ changes – is revealed. During MIS 6e, CDM 6e.2 (at ~182 kyr BP) is nearly synchronous with the abrupt warming in the NH (non-significant lag of 200±360 yr), while the lags during MIS 6d corresponding to CDM 6d.1 and 6d.2 (at ~171 and ~175 kyr BP, respectively) are much longer, 1,400±375 yrs on average. The change in lag time might be related to a change in the organization of the AMOC from MIS 6e to MIS 6d. Similar observations are drawn for the time period covered by our study in comparison with previous studies on MIS 3 and MIS 5 periods, although the lag of CO₂ with respect to NH warming reaches larger values during MIS 6d. However, the limited



available proxy data from the marine realm only permits an exploratory discussion of the mechanisms responsible for CO₂ variability during MIS 6. Because the boundary conditions of the last glacial period cannot be applied to MIS 6, additional proxy data and multiple modelling studies conducted during MIS 6 period are needed.

Acknowledgement

5 This work is a contribution to the “European Project for Ice Coring in Antarctica” (EPICA), a joint European Science Foundation/European Commission scientific program, funded by the European Union and by national contributions from Belgium, Denmark, France, Germany, Italy, The Netherlands, Norway, Sweden, Switzerland, and the United Kingdom. The main logistic support was provided by IPEV and PNRA. This is EPICA publication no. XX. It also received funding from the European Community’s Seventh Framework Programmes ERC-2011-AdG under grant agreement no. 291062 (ERC
10 ICE&LASERS). As part of the PhD work of J. Shin, it was also supported by the LabEX OSUG@2020 project of the Grenoble Observatory of Sciences of the Universe (OSUG). Swiss authors also acknowledge long-term financial support of the ice core research at the University of Bern by the Swiss National Science Foundation under grants 200020_159563, 200020_172745, 200020_172506 and 20FI21_189533. The authors would like to thank G. Aufresne for assistance in the additional methane
15 discussions. We would like to thank J. Gottschalk for discussions about CO₂ variability and marine sediment core data available for the last two glacial periods.



References

- Ahn, J. and Brook, E. J.: Atmospheric CO₂ and climate on millennial time scales during the last glacial period, *Science*, 322, 83-85, 2008.
- Ahn, J. and Brook, E. J.: Siple Dome ice reveals two modes of millennial CO₂ change during the last ice age, *Nat. Commun.*, 5, 3723, 2014.
- Anderson, R., Ali, S., Bradtmiller, L., Nielsen, S., Fleisher, M., Anderson, B., and Burckle, L.: Wind-driven upwelling in the Southern Ocean and the deglacial rise in atmospheric CO₂, *science*, 323, 1443-1448, 2009.
- 5 Ayalon, A., Bar-Matthews, M., and Kaufman, A.: Climatic conditions during marine oxygen isotope stage 6 in the eastern Mediterranean region from the isotopic composition of speleothems of Soreq Cave, Israel, *Geology*, 30, 303-306, 2002.
- Bard, E., Antonioli, F., and Silenzi, S.: Sea-level during the penultimate interglacial period based on a submerged stalagmite from Argentarola Cave (Italy), *Earth Planet. Sci. Lett.*, 196, 135-146, 2002.
- 10 Barker, S., Knorr, G., Edwards, R. L., Parrenin, F., Putnam, A. E., Skinner, L. C., Wolff, E., and Ziegler, M.: 800,000 years of abrupt climate variability, *science*, 334, 347-351, 2011.
- Baumgartner, M., Kindler, P., Eicher, O., Floch, G., Schilt, A., Schwander, J., Spahni, R., Capron, E., Chappellaz, J., and Leuenberger, M.: NGRIP CH₄ concentration from 120 to 10 kyr before present and its relation to a $\delta^{15}\text{N}$ temperature reconstruction from the same ice core, *Clim. Past*, 10, 903-920, 2014.
- 15 Bazin, L., Landais, A., Lemieux-Dudon, B., Kele, H. T. M., Veres, D., Parrenin, F., Martinerie, P., Ritz, C., Capron, E., and Lipenkov, V.: An optimized multi-proxy, multi-site Antarctic ice and gas orbital chronology (AICC2012): 120-800 ka, *Clim. Past*, 9, 1715-1731, 2013.
- Bereiter, B., Eggleston, S., Schmitt, J., Nehrbass-Ahles, C., Stocker, T. F., Fischer, H., Kipfstuhl, S., and Chappellaz, J.: Revision of the EPICA Dome C CO₂ record from 800 to 600 kyr before present, *Geophys. Res. Lett.*, 42, 542-549, 2015.
- Bereiter, B., Lüthi, D., Siegrist, M., Schüpbach, S., Stocker, T. F., and Fischer, H.: Mode change of millennial CO₂ variability during the last glacial cycle associated with a bipolar marine carbon seesaw, *Proc. Natl. Acad. Sci.*, 109, 9755-9760, 2012.
- 20 Bereiter, B., Schwander, J., Lüthi, D., and Stocker, T. F.: Change in CO₂ concentration and O₂/N₂ ratio in ice cores due to molecular diffusion, *Geophys. Res. Lett.*, 36, 2009.
- Berger, A. L.: Long-Term Variations of Caloric Insolation Resulting from the Earth's Orbital Elements 1, *Quat. Res.*, 9, 139-167, 1978.
- Bond, G., Heinrich, H., Broecker, W., Labeyrie, L., McManus, J., Andrews, J., Huon, S., Jantschik, R., Clasen, S., and Simet, C.: Evidence for massive discharges of icebergs into the North Atlantic ocean during the last glacial period, *Nature*, 360, 245, 1992.
- 25 Bouttes, N., Roche, D., and Paillard, D.: Systematic study of the fresh water fluxes impact on the carbon cycle, *Clim. Past*, 7, 2012.
- Bozbiyik, A., Steinacher, M., Joos, F., Stocker, T., and Menviel, L.: Fingerprints of changes in the terrestrial carbon cycle in response to large reorganizations in ocean circulation, *Clim. Past*, 7, 319-338, 2011.
- Bréant, C., Landais, A., Orsi, A., Martinerie, P., Extier, T., Prié, F., Stenni, B., Jouzel, J., Masson-Delmotte, V., and Leuenberger, M.: Unveiling the anatomy of Termination 3 using water and air isotopes in the Dome C ice core, *East Antarctica, Quaternary Sci. Rev.*, 211, 156-165, 2019.
- Broecker, W., Bond, G., Klas, M., Clark, E., and McManus, J.: Origin of the northern Atlantic's Heinrich events, *Clim. Dyn.*, 6, 265-273, 1992.
- Broecker, W. S.: Paleocean circulation during the last deglaciation: a bipolar seesaw?, *Paleoceanography*, 13(2), 119-121, 1998.
- 35 Brook, E. J., Sowers, T., and Orchardo, J.: Rapid variations in atmospheric methane concentration during the past 110,000 years, *Science*, 273, 1087-1091, 1996.
- Buizert, C., Adrian, B., Ahn, J., Albert, M., Alley, R. B., Baggenstos, D., Bauska, T. K., Bay, R. C., Bencivengo, B. B., and Bentley, C. R.: Precise inter-polar phasing of abrupt climate change during the last ice age, *Nature*, 520, 661, 2015.
- Cheng, H., Edwards, R. L., Sinha, A., Spötl, C., Yi, L., Chen, S., Kelly, M., Kathayat, G., Wang, X., and Li, X.: The Asian monsoon over the past 640,000 years and ice age terminations, *Nature*, 534, 640, 2016.
- 40 Craig, H., Horibe, Y., and Sowers, T.: Gravitational separation of gases and isotopes in polar ice caps, *Science*, 242, 1675-1678, 1988.
- de Abreu, L., Shackleton, N. J., Schönfeld, J., Hall, M., and Chapman, M.: Millennial-scale oceanic climate variability off the Western Iberian margin during the last two glacial periods, *Mar. Geol.*, 196, 1-20, 2003.
- Denton, G. H., Anderson, R. F., Toggweiler, J., Edwards, R., Schaefer, J., and Putnam, A.: The last glacial termination, *Science*, 328, 1652-1656, 2010.
- 45 Dreyfus, G. B., Jouzel, J., Bender, M. L., Landais, A., Masson-Delmotte, V., and Leuenberger, M.: Firn processes and $\delta^{15}\text{N}$: potential for a gas-phase climate proxy, *Quaternary Sci. Rev.*, 29, 28-42, 2010.
- Fischer, H., Schmitt, J., Lüthi, D., Stocker, T. F., Tschumi, T., Parekh, P., Joos, F., Köhler, P., Völker, C., and Gersonde, R.: The role of Southern Ocean processes in orbital and millennial CO₂ variations—A synthesis, *Quaternary Sci. Rev.*, 29, 193-205, 2010.
- 50 Fourteau, K., Faïn, X., Martinerie, P., Landais, A., Ekaykin, A. A., Lipenkov, V. Y., and Chappellaz, J.: Analytical constraints on layered gas trapping and smoothing of atmospheric variability in ice under low-accumulation conditions, *Clim. Past*, 13, 1815-1830, 2017.
- Gottschalk, J., Battaglia, G., Fischer, H., Frölicher, T. L., Jaccard, S. L., Jeltsch-Thömmes, A., Joos, F., Köhler, P., Meissner, K. J., and Menviel, L.: Mechanisms of millennial-scale atmospheric CO₂ change in numerical model simulations, *Quaternary Sci. Rev.*, 220, 30-74, 2019a.



- Gottschalk, J., Skinner, L. C., Jaccard, S. L., Menviel, L., Nehrbass-Ahles, C., and Waelbroeck, C.: Southern Ocean link between changes in atmospheric CO₂ levels and northern-hemisphere climate anomalies during the last two glacial periods, *Quaternary Sci. Rev.*, doi: <https://doi.org/10.1016/j.quascirev.2019.106067>, 2019b. 106067, 2019b.
- 5 Grachev, A. M. and Severinghaus, J. P.: Laboratory determination of thermal diffusion constants for ²⁹N₂/²⁸N₂ in air at temperatures from -60 to 0°C for reconstruction of magnitudes of abrupt climate changes using the ice core fossil-air paleothermometer, *Geochim Cosmochim Acta*, 67, 345-360, 2003.
- Heinrich, H.: Origin and consequences of cyclic ice rafting in the northeast Atlantic Ocean during the past 130,000 years, *Quaternary Res.*, 29, 142-152, 1988.
- Henry, L., McManus, J. F., Curry, W. B., Roberts, N. L., Piotrowski, A. M., and Keigwin, L. D.: North Atlantic ocean circulation and abrupt
10 climate change during the last glaciation, *Science*, 353, 470-474, 2016.
- Higgins, J. A., Kurbatov, A. V., Spaulding, N. E., Brook, E., Introne, D. S., Chimiak, L. M., Yan, Y., Mayewski, P. A., and Bender, M. L.: Atmospheric composition 1 million years ago from blue ice in the Allan Hills, Antarctica, *Proc. Natl. Acad. Sci.*, 2015. 201420232, 2015.
- Huber, C., Leuenberger, M., Spahni, R., Flückiger, J., Schwander, J., Stocker, T. F., Johnsen, S., Landais, A., and Jouzel, J.: Isotope calibrated Greenland temperature record over Marine Isotope Stage 3 and its relation to CH₄, *Earth Planet. Sci. Lett.*, 243, 504-519, 2006.
- 15 Jouzel, J., Masson-Delmotte, V., Cattani, O., Dreyfus, G., Falourd, S., Hoffmann, G., Minster, B., Nouet, J., Barnola, J.-M., and Chappellaz, J.: Orbital and millennial Antarctic climate variability over the past 800,000 years, *Science*, 317, 793-796, 2007.
- Köhler, P., Joos, F., Gerber, S., and Knutti, R.: Simulated changes in vegetation distribution, land carbon storage, and atmospheric CO₂ in response to a collapse of the North Atlantic thermohaline circulation, *Clim. Dyn.*, 25, 689, 2005.
- Krinner, G., Raynaud, D., Doutriaux, C., and Dang, H.: Simulations of the Last Glacial Maximum ice sheet surface climate: Implications
20 for the interpretation of ice core air content, *J. Geophys. Res.*, 105, 2059-2070, 2000.
- Landais, A., Barnola, J.-M., Kawamura, K., Caillon, N., Delmotte, M., Van Ommen, T., Dreyfus, G., Jouzel, J., Masson-Delmotte, V., and Minster, B.: Firn-air δ¹⁵N in modern polar sites and glacial-interglacial ice: a model-data mismatch during glacial periods in Antarctica?, *Quaternary Sci. Rev.*, 25, 49-62, 2006.
- Landais, A., Dreyfus, G., Capron, E., Jouzel, J., Masson-Delmotte, V., Roche, D., Prié, F., Caillon, N., Chappellaz, J., and Leuenberger, M.:
25 Two-phase change in CO₂, Antarctic temperature and global climate during Termination II, *Nat. Geosci.*, 6, 1062, 2013.
- Leduc, G., Vidal, L., Tachikawa, K., Rostek, F., Sonzogni, C., Beaufort, L., and Bard, E.: Moisture transport across Central America as a positive feedback on abrupt climatic changes, *Nature*, 445, 908, 2007.
- Loulergue, L., Parrenin, F., Blunier, T., Barnola, J.-M., Spahni, R., Schilt, A., Raisbeck, G., and Chappellaz, J.: New constraints on the gas age-ice age difference along the EPICA ice cores, 0–50 kyr, *Clim. Past*, 3(3), 527–540, 2007.
- 30 Loulergue, L., Schilt, A., Spahni, R., Masson-Delmotte, V., Blunier, T., Lemieux, B., Barnola, J.-M., Raynaud, D., Stocker, T. F., and Chappellaz, J.: Orbital and millennial-scale features of atmospheric CH₄ over the past 800,000 years, *Nature*, 453, 383, 2008.
- Lourantou, A., Chappellaz, J., Barnola, J.-M., Masson-Delmotte, V., and Raynaud, D.: Changes in atmospheric CO₂ and its carbon isotopic ratio during the penultimate deglaciation, *Quaternary Sci. Rev.*, 29, 1983-1992, 2010.
- Lüthi, D., Le Floch, M., Bereiter, B., Blunier, T., Barnola, J.-M., Siegenthaler, U., Raynaud, D., Jouzel, J., Fischer, H., and Kawamura, K.:
35 High-resolution carbon dioxide concentration record 650,000–800,000 years before present, *Nature*, 453, 379, 2008.
- Marcott, S. A., Bauska, T. K., Buizert, C., Steig, E. J., Rosen, J. L., Cuffey, K. M., Fudge, T., Severinghaus, J. P., Ahn, J., and Kalk, M. L.: Centennial-scale changes in the global carbon cycle during the last deglaciation, *Nature*, 514, 616, 2014.
- Margari, V., Skinner, L., Tzedakis, P., Ganopolski, A., Vautravers, M., and Shackleton, N.: The nature of millennial-scale climate variability during the past two glacial periods, *Nat. Geosci.*, 3, 127, 2010.
- 40 Margari, V., Skinner, L. C., Hodell, D. A., Martrat, B., Toucanne, S., Grimalt, J. O., Gibbard, P. L., Lunkka, J., and Tzedakis, P.: Land-ocean changes on orbital and millennial time scales and the penultimate glaciation, *Geology*, 42, 183-186, 2014.
- McManus, J. F., Anderson, R. F., Broecker, W. S., Fleisher, M. Q., and Higgins, S. M.: Radiometrically determined sedimentary fluxes in the sub-polar North Atlantic during the last 140,000 years, *Earth Planet. Sci. Lett.*, 155, 29-43, 1998.
- Mélières, M. A., Rossignol-Strick, M., and Malaizé, B.: Relation between low latitude insolation and δ¹⁸O change of atmospheric oxygen for the last 200 kyrs, as revealed by Mediterranean sapropels, *Geophys Res Lett*, 24, 1235-1238, 1997.
- 45 Menviel, L., England, M. H., Meissner, K., Mouchet, A., and Yu, J.: Atlantic-Pacific seesaw and its role in outgassing CO₂ during Heinrich events, *Paleoceanography*, 29, 58-70, 2014.
- Menviel, L., Timmermann, A., Mouchet, A., and Timm, O.: Meridional reorganizations of marine and terrestrial productivity during Heinrich events, *Paleoceanography*, 23, PA1203, 2008.
- 50 Nehrbass-Ahles, C., Shin, J., Schmitt, J., Bereiter, B., Joos, F., Schilt, A., Schmidely, L., Silva, L., Teste, G., Grilli, R., Chappellaz, J., Hodell, D., Fischer, H., and Stocker, T. F.: Abrupt CO₂ release to the atmosphere under both glacial and early interglacial conditions, in review. in review.
- Obata, A.: Climate-carbon cycle model response to freshwater discharge into the North Atlantic, *Journal of Climate*, 20, 5962-5976, 2007.
- Parrenin, F., Barker, S., Blunier, T., Chappellaz, J., Jouzel, J., Landais, A., Masson-Delmotte, V., Schwander, J., and Veres, D.: On the gas-ice depth difference (Delta depth) along the EPICA Dome C ice core, *Clim. Past*, 8, 1239-1255, 2012.
- 55



- Parrenin, F., Barnola, J.-M., Beer, J., Blunier, T., Castellano, E., Chappellaz, J., Dreyfus, G., Fischer, H., Fujita, S., and Jouzel, J.: The EDC3 chronology for the EPICA Dome C ice core, *Clim. Past*, 3, 485-497, 2007.
- Parrenin, F., Masson-Delmotte, V., Köhler, P., Raynaud, D., Paillard, D., Schwander, J., Barbante, C., Landais, A., Wegner, A., and Jouzel, J.: Synchronous change of atmospheric CO₂ and Antarctic temperature during the last deglacial warming, *Science*, 339, 1060-1063, 2013.
- 5 Peterson, L. C., Haug, G. H., Hughen, K. A., and Röhl, U.: Rapid changes in the hydrologic cycle of the tropical Atlantic during the last glacial, *Science*, 290, 1947-1951, 2000.
- Petit, J.-R., Jouzel, J., Raynaud, D., Barkov, N. I., Barnola, J.-M., Basile, I., Bender, M., Chappellaz, J., Davis, M., and Delaygue, G.: Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica, *Nature*, 399, 429, 1999.
- Rasmussen, S. O., Bigler, M., Blockley, S. P., Blunier, T., Buchardt, S. L., Clausen, H. B., Cvijanovic, I., Dahl-Jensen, D., Johnsen, S. J., and Fischer, H.: A stratigraphic framework for abrupt climatic changes during the Last Glacial period based on three synchronized Greenland ice-core records: refining and extending the INTIMATE event stratigraphy, *Quaternary Sci. Rev.*, 106, 14-28, 2014.
- 10 Richter, I. and Xie, S.-P.: Moisture transport from the Atlantic to the Pacific basin and its response to North Atlantic cooling and global warming, *Clim. Dyn.*, 35, 551-566, 2010.
- Rosen, J. L., Brook, E. J., Severinghaus, J. P., Blunier, T., Mitchell, L. E., Lee, J. E., Edwards, J. S., and Gkinis, V.: An ice core record of near-synchronous global climate changes at the Bølling transition, *Nat. Geosci.*, 7, 459, 2014.
- 15 Schaefer, H., Laurantou, A., Chappellaz, J., Lüthi, D., Bereiter, B., and Barnola, J.-M.: On the suitability of partially clathrated ice for analysis of concentration and $\delta^{13}\text{C}$ of palaeo-atmospheric CO₂, *Earth Planet. Sci. Lett.*, 307, 334-340, 2011.
- Schmittner, A., Brook, E. J., and Ahn, J.: Impact of the ocean's overturning circulation on atmospheric CO₂. In: *Ocean Circulation: Mechanisms and Impacts*, AGU Geophysical Monograph Series, 173, American Geophysical Union, Washington DC, 315-334, 2007.
- 20 Schmittner, A. and Galbraith, E. D.: Glacial greenhouse-gas fluctuations controlled by ocean circulation changes, *Nature*, 456, 373, 2008.
- Shackleton, N. J., Hall, M. A., and Vincent, E.: Phase relationships between millennial-scale events 64,000-24,000 years ago, *Paleoceanography*, 15(6), 565-569, 2000.
- Sigman, D. M. and Boyle, E. A.: Glacial/interglacial variations in atmospheric carbon dioxide, *Nature*, 407, 859, 2000.
- 25 Sigman, D. M., De Boer, A. M., and Haug, G. H.: Antarctic stratification, atmospheric water vapor, and Heinrich events: A hypothesis for late Pleistocene deglaciations, 173, American Geophysical Union, Washington DC, 315-334, 2007.
- Sowers, T., Bender, M., and Raynaud, D.: Elemental and isotopic composition of occluded O₂ and N₂ in polar ice, *J. Geophys. Res.*, 94, 5137-5150, 1989.
- Spahni, R., Chappellaz, J., Stocker, T. F., Loulergue, L., Hausammann, G., Kawamura, K., Flückiger, J., Schwander, J., Raynaud, D., and Masson-Delmotte, V.: Atmospheric methane and nitrous oxide of the late Pleistocene from Antarctic ice cores, *Science*, 310, 1317-1321, 30 2005.
- Stocker, T. F. and Johnsen, S. J.: A minimum thermodynamic model for the bipolar seesaw, *Paleoceanography*, 18, PA000920, 2003.
- Toggweiler, J. R., Russell, J. L., and Carson, S. R.: Midlatitude westerlies, atmospheric CO₂, and climate change during the ice ages, *Paleoceanography*, 21, PA2005, 2006.
- 35 Wang, Y., Cheng, H., Edwards, R. L., Kong, X., Shao, X., Chen, S., Wu, J., Jiang, X., Wang, X., and An, Z.: Millennial-and orbital-scale changes in the East Asian monsoon over the past 224,000 years, *Nature*, 451, 1090, 2008.
- Wilson, D. J., Piotrowski, A. M., Galy, A., and Banakar, V. K.: Interhemispheric controls on deep ocean circulation and carbon chemistry during the last two glacial cycles, *Paleoceanography*, 30, 621-641, 2015.



Table 1: Ages of the tie points of MD01–2444 record (Margari et al., 2010) on the EDC03 (Parrenin et al., 2007) and AICC2012 ice age scales (Bazin et al., 2013).

MD01–2444 (m)	Previous age tie points by Margari et al. (2010) (EDC03 age scale, kyr BP)	New age tie points (AICC2012, kyr BP)
22.02	136.100	135.761
22.50	141.686	141.677
23.70	149.586	150.287
24.48	159.105	160.327
24.72	162.476	163.878
25.32	168.273	170.349
25.71	172.009	174.649
25.95	175.461	178.423
26.01	177.065	180.033
27.03	188.009	190.229
27.30	192.231	194.186



Table 2: The minima and maxima defining the duration of each NA stadial using the AICC2012 chronology.

	MIS 6c.1		MIS 6c.2		MIS 6d.1		MIS 6d.2		MIS 6e.1		MIS 6e.2	
	Min.	Max.	Min.	Max.	Min.	Max.	Min.	Max.	Min.	Max.	Min.	Max.
Age (kyr BP)	160.1	161.9	163.2	166.0	170.2	173.5	174.2	176.8	177.2	178.3	179.9	182.2
Uncertainty (kyr)	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.1	0.1	0.1




Figure 1: Proxy data during 250 kyr BP. A: $\delta^{18}\text{O}_{\text{calcite}}$ from Sanbao cave, corresponding with the strength of the East Asian monsoon (Cheng et al., 2016). B: 21 June insolation for 65°N (Berger, 1978). C: Dust flux in EDC (Lambert et al., 2012). D: Atmospheric CH_4 in EDC (green dots) (Loulergue et al., 2008) and Atmospheric CH_4 in EDC in this study (light yellow dots).
5 E: Atmospheric CO_2 from EDC in this study (light blue dots) and composite CO_2 from Antarctic ice cores (dark blue dots) (Bereiter et al., 2015). F: δD composition in EDC, Antarctica (Jouzel et al., 2007). Vertical grey bars indicate the timing of Heinrich events.

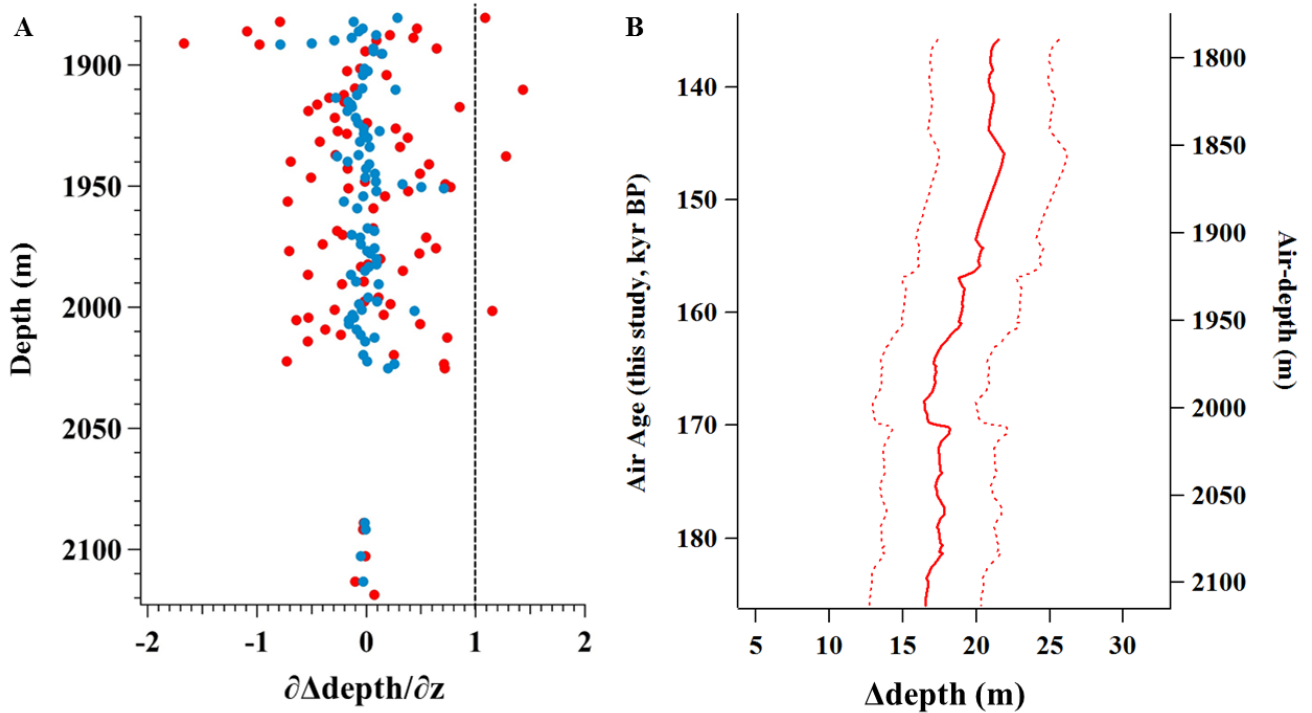


Figure 2: A: $\partial\Delta\text{depth}/\partial z$ as a function of depth. Red dots from the raw $\delta^{15}\text{N}$ measurements. Blue dots from a 3-point running mean weighted by $500/dT$. Vertical dashed line indicates when $\partial\Delta\text{depth}/\partial z$ function is 1. B: Δdepth (bold line) for EDC from 1787.5 to 1870.2 m below the surface, deduced from $\delta^{15}\text{N}$ and the thinning function calculated in this depth range. The two dash lines correspond to the analytical uncertainties.

10

15

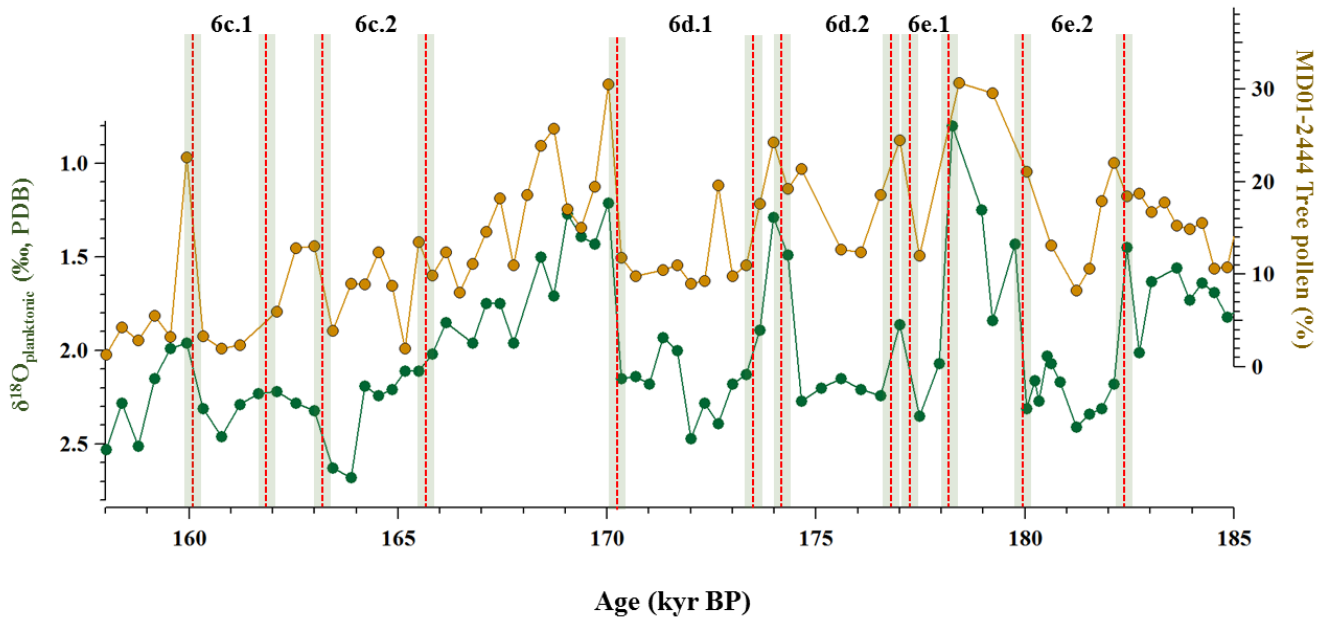


Figure 3: Six variations on millennial time scales of tree pollen percentage (top) and $\delta^{18}\text{O}$ of planktonic foraminifera (bottom) in the MD01-2444 selected by Margari et al. (2010).




Figure 4: Atmospheric CO₂ from EDC and Vostok ice cores, compared to the δD of water at EDC (temperature proxy) during 190–135 kyr BP. Blue dots: Atmospheric CO₂ from EDC (this study). Yellow dots: Atmospheric CO₂ from EDC (Lourantou et al., 2010). Grey dots: Atmospheric CO₂ from the Vostok ice core (Petit et al., 1999). Grey line: δD of water at EDC (Jouzel et al., 2007).

5

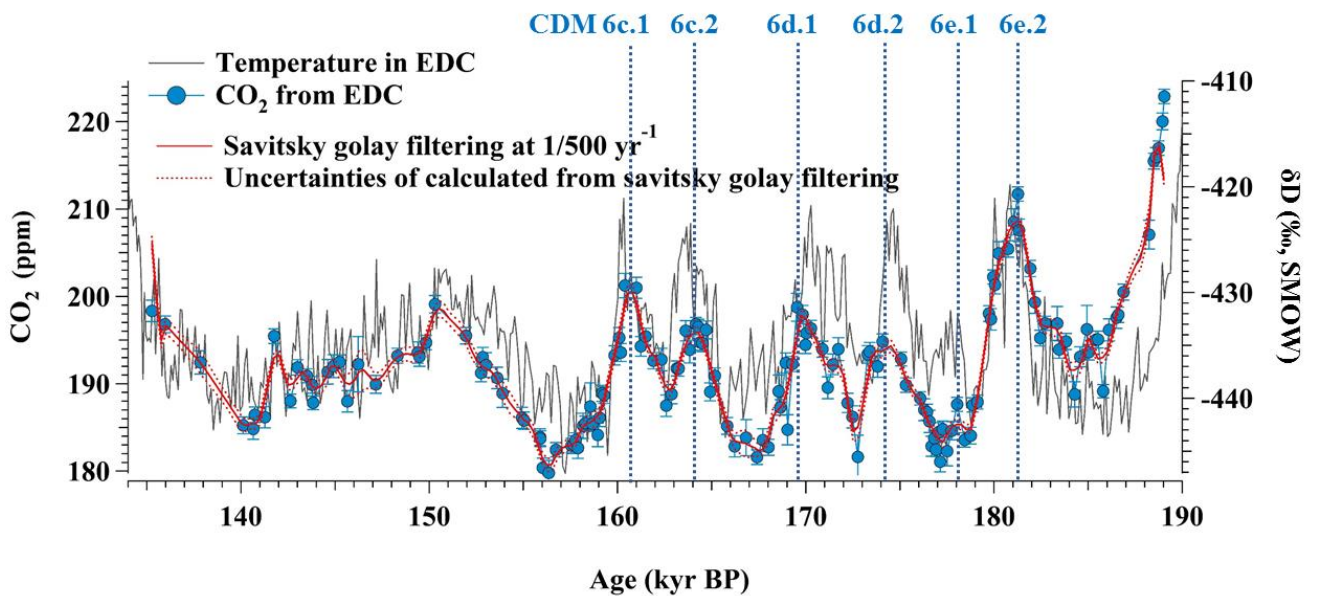


Figure 5: Atmospheric CO₂ from EDC (this study) and EDC water isotopic record (Jouzel et al., 2007). Red line indicates Savitsky Golay filtering curve made with a 500 yr cut-off period (red dotted line). Vertical blue dotted lines indicate the six CDM events that we identify during the early MIS 6.

5

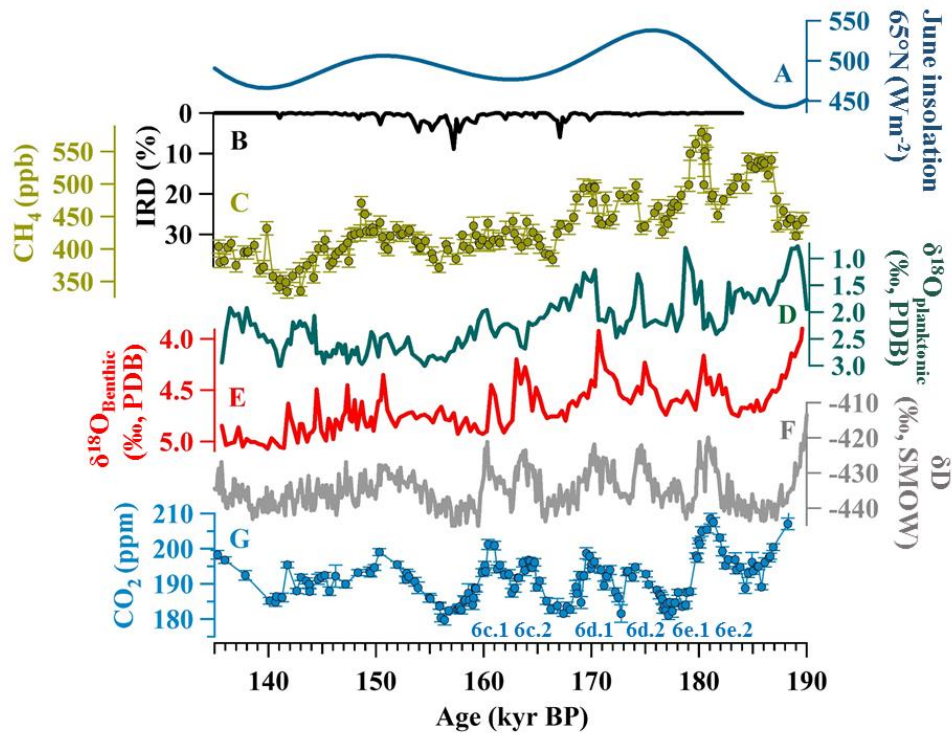


Figure 6: A comparison of climate with atmospheric CO₂ during MIS 6 period. A: 21 June insolation at 65°N (Berger, 1978). B: Ice-rafted debris (IRD) input in the Iberian margin core MD95–2040 (de Abreu et al., 2003). C: Atmospheric CH₄ in EDC (Loulergue et al., 2008; this study). D: δ¹⁸O of planktonic foraminifera in the Iberian margin marine Core MD01–2444 (Margari et al., 2010). E: δ¹⁸O of Benthic foraminifera in the Iberian margin marine Core MD01–2444 (Margari et al., 2010). F: Temperature in Antarctica from δD composition of the EDC ice core (Jouzel et al., 2007). G: Atmospheric CO₂ in EDC during MIS 6 period (in this study). The numbers of CDM events are written at the bottom.




Figure 7: The relationship between NA Stadial duration and magnitude of CO₂ increase during the early MIS 6 period. Green dots indicate non Heinrich CDM events during MIS 3, green dots with a white dot in the middle indicate classic Heinrich CDM events during MIS 3, and green dots with a red cross in the middle indicate non classic Heinrich CDM events during the MIS 3 period. Blue dots indicate CDM events during MIS 6 respectively.

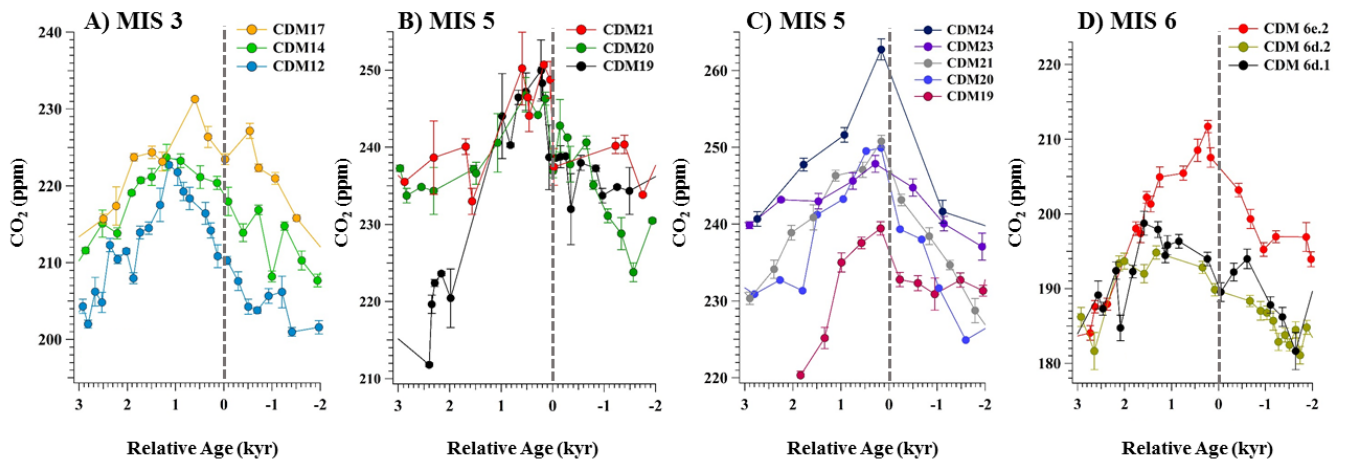


Figure 8: CDM lags relative to abrupt temperature increase in NH. Grey dotted lines indicate when climate changes abruptly in NH. A: Atmospheric CO₂ was recorded from TALDICE (MIS 3), B: Atmospheric CO₂ was recorded from Byrd (MIS 5), C: Atmospheric CO₂ was recorded from EDML (MIS 5), D: Atmospheric CO₂ was recorded from EDC (MIS 6). For MIS 6, we selected the 3 CDMs that correspond to an abrupt methane increase; the other CDMs do not correspond to an abrupt change. The scale of the y-axis is not the same for the four panels

Supplement Information of
Millennial-scale atmospheric CO₂ variations during the Marine Isotope Stage 6 period (190-135 kyr BP)

Jinhwa Shin et al.,

5 Correspondence to: Jérôme Chappellaz (jerome.chappellaz@univ-grenoble-alpes.fr)




Figure S1: Comparison of climate related signals during MIS 3 (left) and 6 (right) period. A: 21 June insolation at 65°N (Berger, 1978). B: Ice-rafted debris (IRD) input in the Iberian margin core MD95–2040 (de Abreu et al., 2003). C: $\delta^{18}\text{O}$ of planktonic foraminifera in the Iberian margin marine Core MD01–2444 (Margari et al., 2010). D: $\delta^{18}\text{O}$ of Benthic foraminifera in the Iberian margin marine Core MD01–2444 (Margari et al., 2010). E: Temperature in Antarctica from δD composition of the EDC ice core (Jouzel et al., 2007).

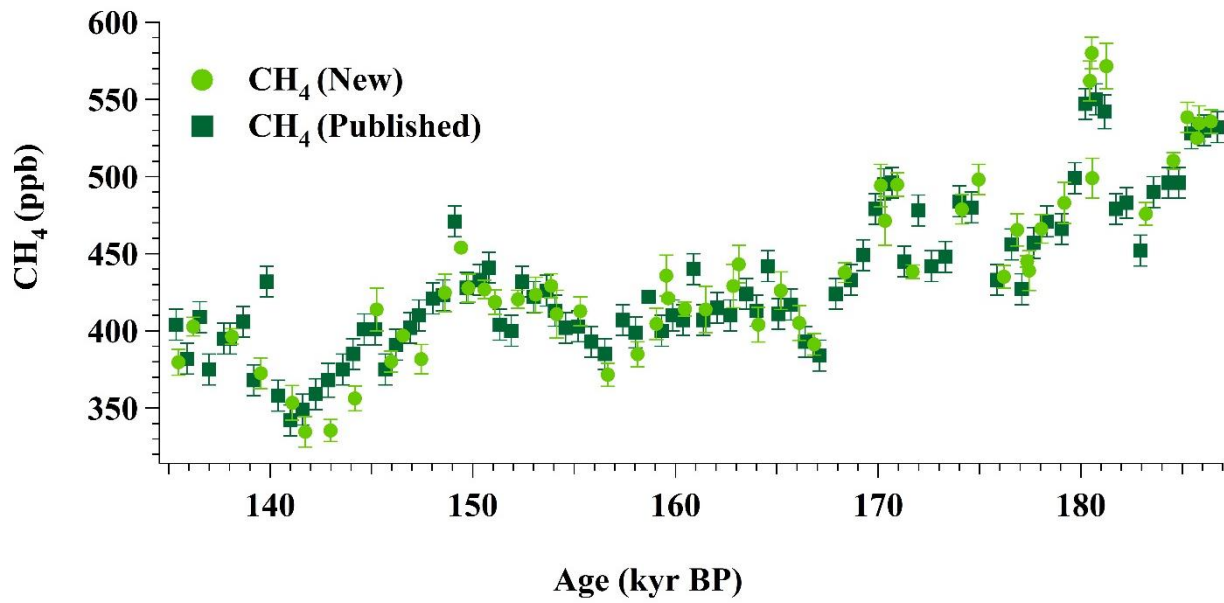


Figure S2: Atmospheric CH₄ concentrations from EDC ice core during MIS 6. Both the dark green squares (Loulergue et al., 2008) and the light green dots (this study) were measured at IGE.

5

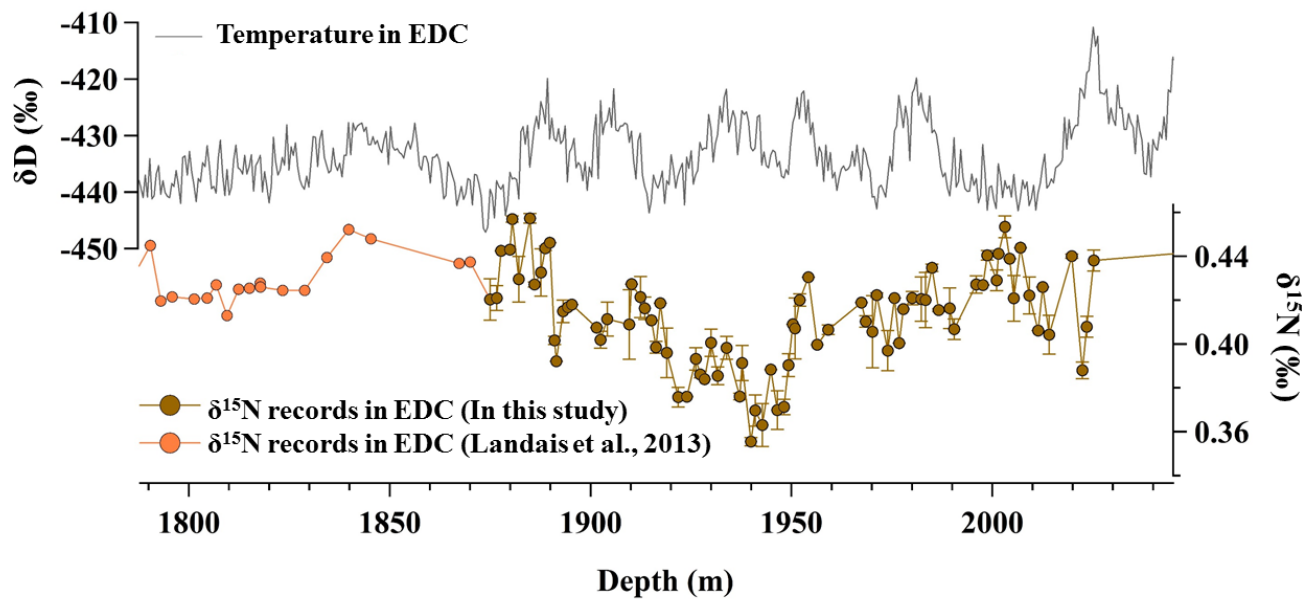


Figure S3: δD and $\delta^{15}N$ from EDC ice core plotted as a function of depth. For $\delta^{15}N$, 88 new data points are added to the previous measurements (Landais et al., 2013). The error bar indicates the standard deviation of replicate measurements.

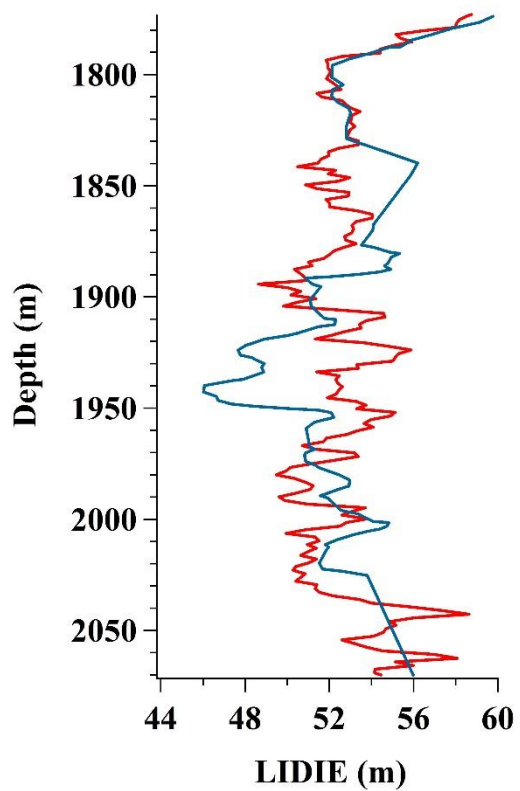


Figure S4: The Lock-In Depth in Ice Equivalent (LIDIE) calculated in AICC2012 (Bazin et al., 2013) and the LIDIE deduced from $\delta^{15}\text{N}$ in this study.




Figure S5: Blue dots: Atmospheric CO₂ from EDC on the revised AICC 2012 age scale. Green dots: Atmospheric CO₂ from EDC on the AICC 2012 age scale. Grey line: δD of water from EDC (temperature proxy) (Jouzel et al., 2007).

5

10

15

20

5

Definition of the onset of abrupt climate change in the NH

Over the last glacial period, rapid CH₄ jumps were always synchronous with abrupt temperature increases in Greenland within ± 50 ppb (Huber et al., 2006). We pick intervals when CH₄ increases rapidly by at least 50 ppb over a time period of less than 1 kyr that correspond with Antarctic isotope maxima (Louergue et al., 2008). The timing of abrupt CH₄ increases was defined as the midpoint between the beginning of the increase of CH₄ and its maximum. The age uncertainty of the midpoint is defined by the time difference between the midpoint and either of the two endpoints.

We found three abrupt CH₄ increases during MIS 6 period at 171.1 ± 0.2 , 175.4 ± 0.4 and 181.5 ± 0.3 kyr BP (Figure S6). Due to the low accumulation rate and low temperature at the site during glacial periods, abrupt changes of CH₄ concentration might be smoothed, and identifying abrupt changes of CH₄ is more difficult than for interglacial periods. The climate change at 175.4 kyr BP does not seem to occur as abruptly as the other two, since CH₄ varied slowly over ~ 800 years. However, we include this event because corresponding data of $\delta^{18}\text{O}$ composition of planktonic foraminifera (Shackleton et al., 2000) indicate a rapid warming, and therefore an abrupt climate change in NH. Rapid increases during the last glacial period (MIS 3 and 5) are also calculated using this method to identify the onset of abrupt warming in NH. In total, eight changes are selected during this period (Figure S7–S9).




Figure S6: Atmospheric CH₄ records from EDC during MIS 6 period. Three boxes show CH₄ jumps at 171.1 ± 0.2 , 175.4 ± 0.4 and 181.5 ± 0.3 kyr BP.

15

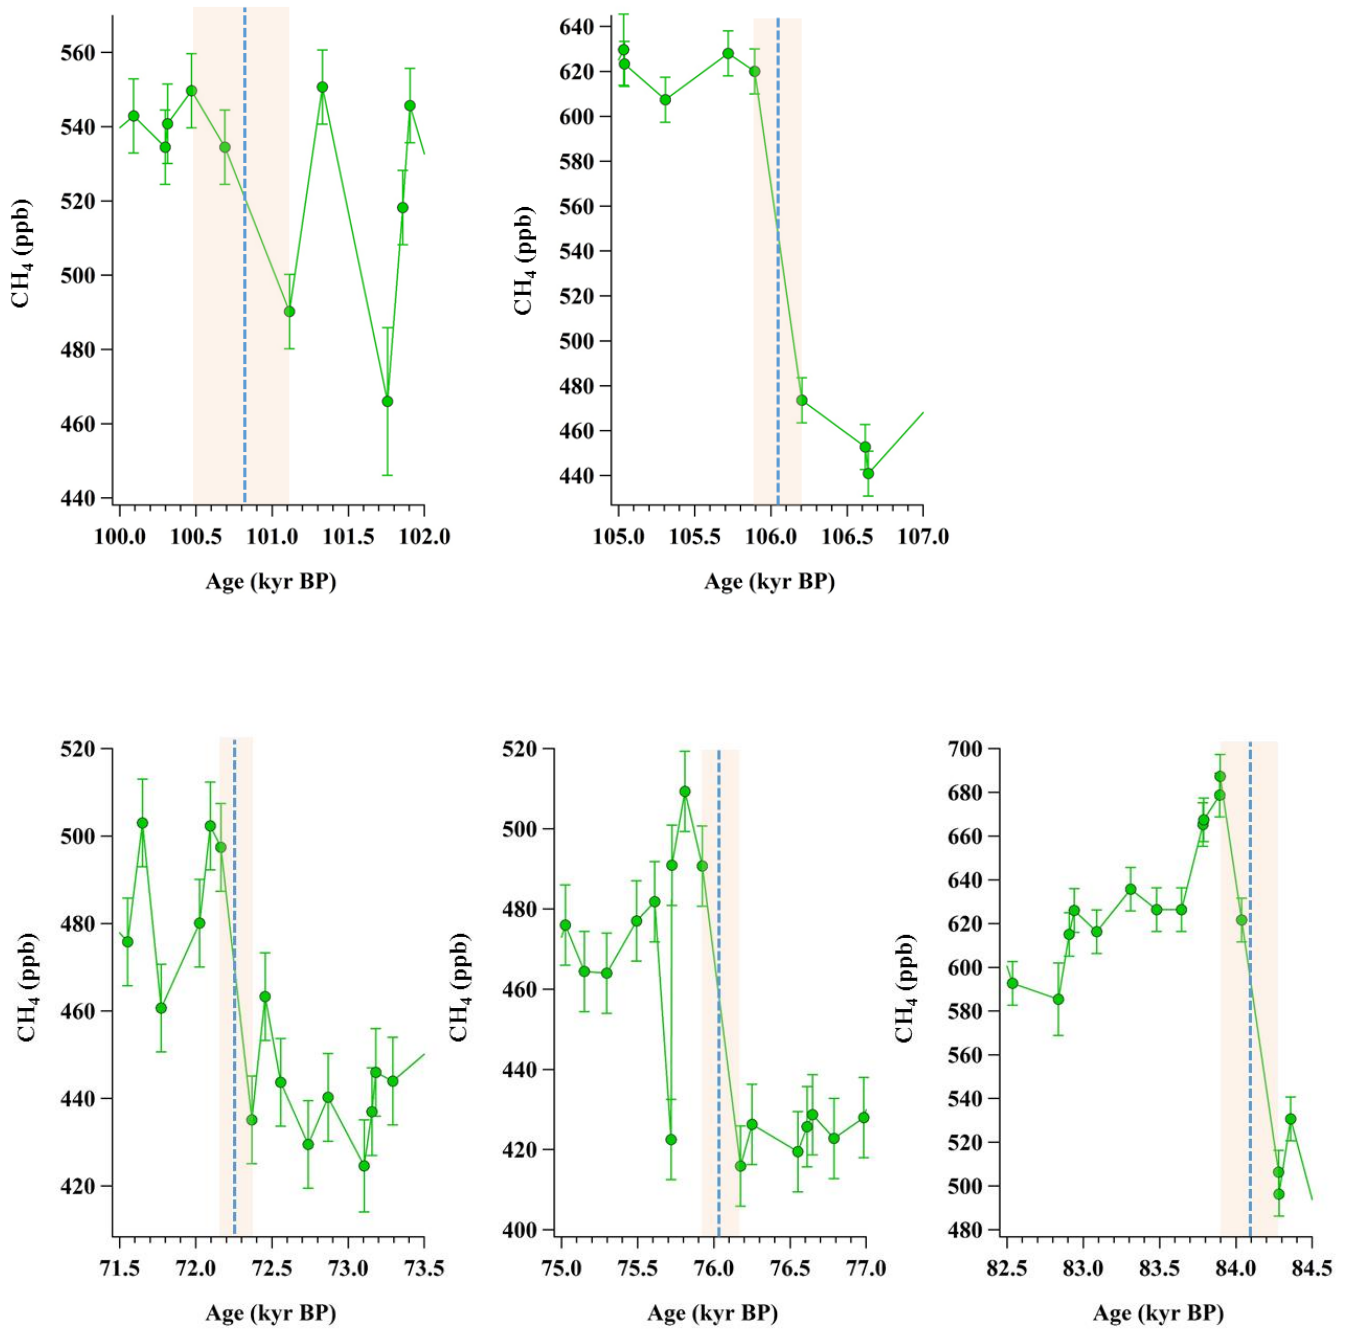


Figure S7: Atmospheric CH₄ records from EDML during the MIS 5 period. Two boxes show CH₄ jumps at 72.3±0.1, 76.0±0.1, 84.1±0.2, 100.8±0.5 and 106.0±0.2 kyr BP.




Figure S8: Atmospheric CH₄ records from the Byrd ice core during MIS 5 period. Three boxes show CH₄ jumps at 72.2 ± 0.1 , 76.0 ± 0.2 and 84.1 ± 0.04 kyr BP.

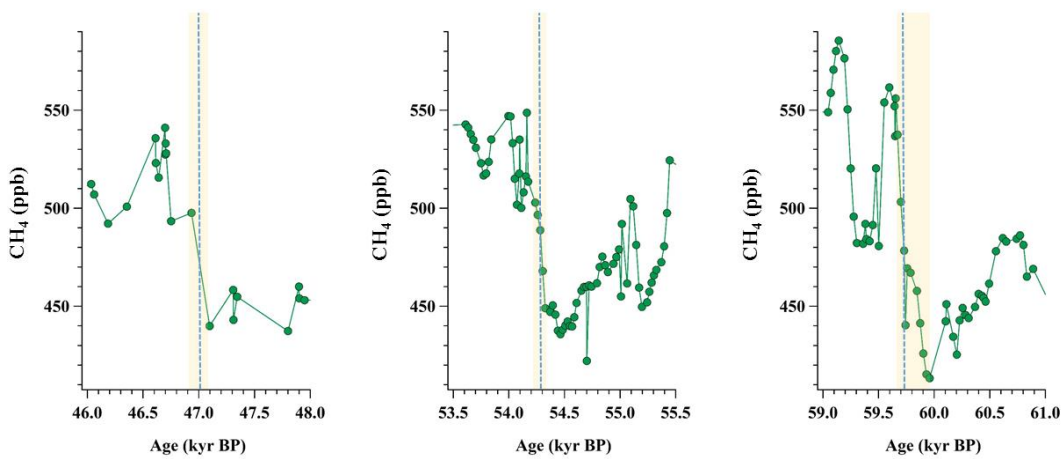


Figure S9: Atmospheric CH₄ records from TALDICE during MIS 3 period. Three boxes show CH₄ jumps at 46.7 ± 0.2 , 54.2 ± 0.1 and 59.7 ± 0.1 kyr BP.

Definition of minima and maxima of atmospheric CO₂

A two-steps procedure was used in order to select the maxima and minima of CO₂ concentrations during the penultimate glacial periods, and calculate the associated age uncertainty (Figure S10 and Table S1). First, inflection points were selected by finding zero values in the second Savitsky–Golay filtered derivative of the data. The parameters of the Savitsky–Golay filters were chosen in order to remove sub–millennial scale variations. We use the same parameters for MIS 3, MIS 5 and 6. Second, a Monte Carlo simulation was conducted, in which the original data were resampled within their uncertainty, and the absolute minima and maxima between pairs of inflection points were selected. This allows us to assign an approximate uncertainty to the timing of each minimum/maximum. The square of the age uncertainty associated with sampling (taken to be the mean sampling resolution) was added to the squared uncertainty calculated in the Monte Carlo procedure to calculate a total uncertainty value.

Table S1: The minima and maxima locations of atmospheric CO₂ during the MIS 6.

	CDM 6c.1			CDM 6c.2			CDM 6d.1			CDM 6d.2			CDM 6e.2		
	Min	Max	Min	Min	Max	Min	Min	Max	Min	Min	Max	Min	Min	Max	Min
Age (kyr BP)	156.3	160.6	162.6	162.6	164.2	167.2	167.2	169.6	172.7	172.7	174.1	177.2	177.2	181.3	184.9
2 δ (kyr BP)	0.3	0.3	0.2	0.2	0.3	0.5	0.5	0.2	0.2	0.2	0.2	0.3	0.3	0.2	0.8

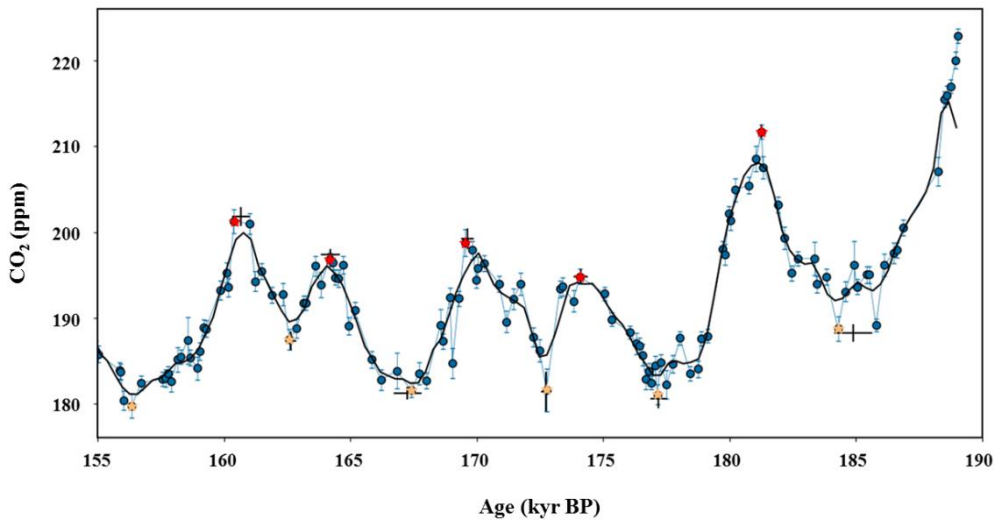


Figure S10: The red and yellow points are the minimum/maximum measured points of atmospheric CO₂ during MIS 6 respectively. Blue dots indicate atmospheric CO₂. The bars indicate the timing and CO₂ uncertainty for each minimum/maximum.

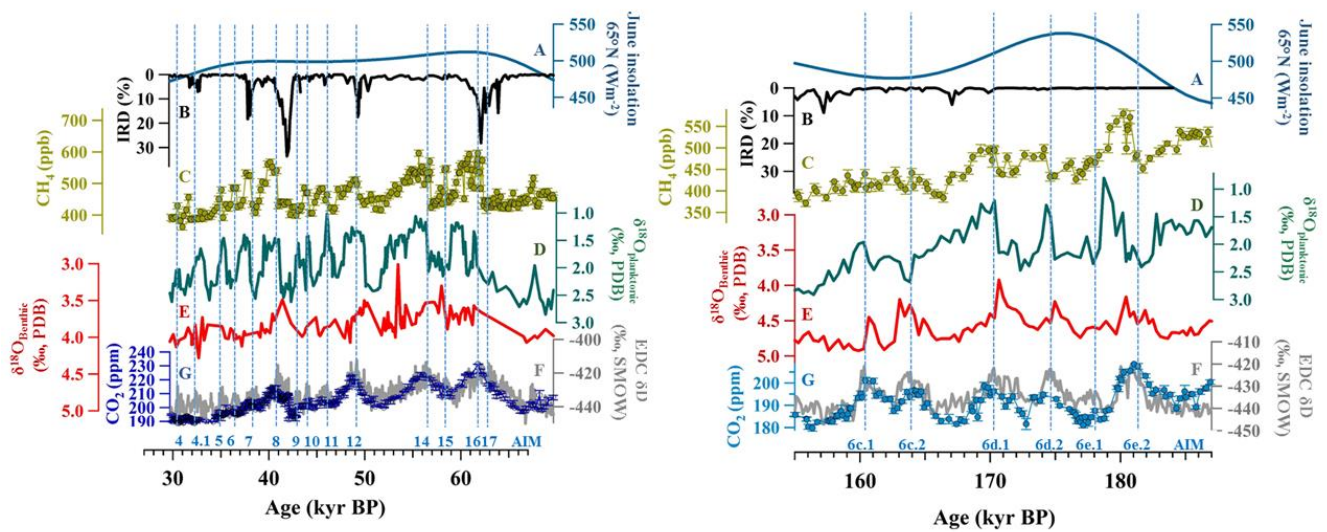


Figure S11: Comparison of climate with atmospheric CO₂ during MIS 3 (left) and 6 (right) period. A: 21 June insolation at 65°N (Berger, 1978). B: Ice-rafted debris (IRD) input in the Iberian margin core MD95–2040 (de Abreu et al., 2003). C: Atmospheric CH₄ in EDC during MIS 3 (Loulergue et al., 2008) and composite data of atmospheric CH₄ in EDC during MIS 6 (Loulergue et al., 2008; this study). D: δ¹⁸O of planktonic foraminifera in the Iberian margin marine Core MD01–2444 (Margari et al., 2010). E: δ¹⁸O of Benthic foraminifera in the Iberian margin marine Core MD01–2444 (Margari et al., 2010). F: Temperature in Antarctica from δD composition of the EDC ice core (Jouzel et al., 2007). G: Composite data of atmospheric CO₂ in Antarctic ice cores during MIS 3 (Bereiter et al., 2015) and atmospheric CO₂ in EDC during MIS 6 (this study). Dashed lines indicates the timing of AIM events. The numbers of AIM events are written at the bottom of the dashed lines.

5

10

15

10

References

- Bazin, L., Landais, A., Lemieux-Dudon, B., Kele, H. T. M., Veres, D., Parrenin, F., Martinerie, P., Ritz, C., Capron, E., and Lipenkov, V.: An optimized multi-proxy, multi-site Antarctic ice and gas orbital chronology (AICC2012): 120-800 ka, *Clim. Past*, 9, 1715-1731, 2013.
- 5 Bereiter, B., Eggleston, S., Schmitt, J., Nehrbass-Ahles, C., Stocker, T. F., Fischer, H., Kipfstuhl, S., and Chappellaz, J.: Revision of the EPICA Dome C CO₂ record from 800 to 600 kyr before present, *Geophys. Res. Lett.*, 42, 542-549, 2015.
- Berger, A. L.: Long-Term Variations of Caloric Insolation Resulting from the Earth's Orbital Elements 1, *Quat. Res.*, 9, 139-167, 1978.
- 10 de Abreu, L., Shackleton, N. J., Schönfeld, J., Hall, M., and Chapman, M.: Millennial-scale oceanic climate variability off the Western Iberian margin during the last two glacial periods, *Mar. Geol.*, 196, 1-20, 2003.
- Jouzel, J., Masson-Delmotte, V., Cattani, O., Dreyfus, G., Falourd, S., Hoffmann, G., Minster, B., Nouet, J., Barnola, J.-M., and Chappellaz, J.: Orbital and millennial Antarctic climate variability over the past 800,000 years, *Science*, 317, 793-796, 2007.
- 15 Landais, A., Dreyfus, G., Capron, E., Jouzel, J., Masson-Delmotte, V., Roche, D., Prié, F., Caillon, N., Chappellaz, J., and Leuenberger, M.: Two-phase change in CO₂, Antarctic temperature and global climate during Termination II, *Nat. Geosci.*, 6, 1062, 2013.
- Loulergue, L., Schilt, A., Spahni, R., Masson-Delmotte, V., Blunier, T., Lemieux, B., Barnola, J.-M., Raynaud, D., Stocker, T. F., and Chappellaz, J.: Orbital and millennial-scale features of atmospheric CH₄ over the past 800,000 years, *Nature*, 453, 20 383, 2008.
- Margari, V., Skinner, L., Tzedakis, P., Ganopolski, A., Vautravers, M., and Shackleton, N.: The nature of millennial-scale climate variability during the past two glacial periods, *Nat. Geosci.*, 3, 127, 2010.
- Shackleton, N. J., Hall, M. A., and Vincent, E.: Phase relationships between millennial-scale events 64,000–24,000 years ago, *Paleoceanography*, 15(6), 565-569, 2000.

25