Title

Exploring Measurements, Model Acceleration, and Model Benchmarking Towards Informing Improved Air Quality and Climate Policy

ISBN

9798189276439

Author

Quevedo, Duncan Conrado

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Exploring Measurements, Model Acceleration, and Model Benchmarking Towards Informing Improved Air Quality and Climate Policy

by

Duncan Conrado Quevedo

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Civil and Environmental Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Cesunica E. Ivey, Chair
Professor Robert A. Harley
Professor Fotini Katopodes Chow

Spring 2026

Abstract

Exploring Measurements, Model Acceleration, and Model Benchmarking Towards Informing Improved Air Quality and Climate Policy

by

Duncan Conrado Quevedo

Doctor of Philosophy in Engineering – Civil and Environmental Engineering

University of California, Berkeley

Professor Cesunica E. Ivey, Chair

This dissertation explores a variety of modeling approaches for improving air quality policy through data analysis with direct policy recommendations and through model development with policy implications.

Much of California's population lives in regions that do not meet the annual National Ambient Air Quality Standards (NAAQS) for $PM_{2.5}$. Generalized additive modeling (GAM) can characterize the $PM_{2.5}$ response to predictor covariates, while sensitivity analysis enables variable importance ranking for covariates. These methods can improve understanding of chemical and meteorological influences on $PM_{2.5}$ in California. In Chapter 2, we construct GAMs for $PM_{2.5}$ and particulate NO_3^-, SO_4^{2-}, NH_4^+, elemental carbon, and organic carbon in Bakersfield, Fresno, Los Angeles, Riverside, Sacramento, and San Jose. We compute performance metrics for all models and further analyze our Bakersfield, Riverside, and San Jose models for $PM_{2.5}$ and selected speciated components by performing sensitivity analysis and calculating marginal effects. We find chemistry drives $PM_{2.5}$ in Bakersfield (71% of model variance) and San Jose (66%) while meteorology balances chemistry's influence in Riverside (48% meteorology). In Bakersfield, NO_X is the dominant influence whereas HCHO dominates in San Jose. Riverside is influenced nearly equally by NO_X, HCHO, and relative humidity. Each of these covariates displays a nonlinear, monotonic, positive association with $PM_{2.5}$. As analysis sites are in California Assembly Bill 617 communities, our results can inform local control strategies to alleviate exposure concerns in environmental justice communities, and we make recommendations to this end.

In 3, we extend the analytical techniques of Chapter 2 to hourly data measured during four field campaigns in California, two located in Riverside and one each in Wilmington and Bakersfield. We train GAMs to predict organic aerosol (OA) at these study sites against covarying particulate and gaseous species as well as meteorological covariates. For each site, we train a distinct model for daytime OA and nighttime OA in order to investigate

temporally different regimes. Initially, we train each model with the same set of covariates, but the models undergo a selection algorithm such that each final model's predictor set consists of only the most statistically significant covariates. We perform sensitivity analysis for each model and analyze the covariates' marginal effects to conduct variable importance ranking. Drawing from the literature, we contextualize modeled relationships based on first-principles mechanistic explanations or correlations that might act as proxies for uncontrolled variables. We find mobile emissions are strongly associated with OA in all three study sites, with acidic SO_4^{2-} contributions in Riverside, O_3 control co-benefits in Wilmington, and strong biogenic influences in Bakersfield. We make policy recommendations based on these results with the aim of targeting local control strategies to be most efficacious.

Turning from data-driven to first-principles modeling, the Community Multiscale Air Quality (CMAQ) model simulates atmospheric phenomena including advection, diffusion, gas-phase chemistry, aerosol physics and chemistry, and cloud processes. Gas-phase chemistry is often a major computational bottleneck due to its representation as large systems of coupled nonlinear stiff differential equations. In Chapter 4, we leverage the parallel computational power of graphics processing unit (GPU) hardware to accelerate the numerical integration of these systems in CMAQ's CHEM module. Our implementation, CMAQ-CUDA, migrates CMAQ's Rosenbrock solver from Fortran to CUDA Fortran. CMAQ-CUDA accelerates the Rosenbrock solver such that simulations using the chemical mechanisms RACM2, CB6R5, and SAPRC07 require only 51%, 50%, or 35% as much time, respectively, as CMAQv5.4 to complete a chemistry time step. Our results demonstrate that CMAQ is amenable to GPU acceleration and highlight a novel Rosenbrock solver implementation for reducing the computational burden imposed by the CHEM module. We discuss the policy implications of model acceleration.

Chapter 5 focuses on a different dimension of model development: error quantification. There remains wide variability in model predictions of aerosol and its impacts in regional air quality models and Earth system models. This variability stems from the diverse approaches taken in the design of aerosol models in terms of their representations of particle populations, microphysics, and chemistry. The Aerosol Model Benchmark Repository and Standards (AMBRS) Project seeks to develop a platform to benchmark aerosol models towards characterizing model variability. AMBRS uses the Particle Monte Carlo (PartMC) model as its particle-resolved benchmark model and the 4-Mode Modal Aerosol Module (MAM4) as a representative reduced model. PartMC is configurable in a highly flexible manner whereas MAM4, like other reduced aerosol modules implemented in large-scale atmosphere models, is less flexible and features extensive hard coding. To benchmark a reduced model with PartMC, PartMC should be configured to behave similarly to the reduced model. We achieve this by developing an interface between AMBRS and the Chemistry Across Multiple Phases (CAMP) model, which features configurable gas- and particle-phase chemical kinetics. This interface facilitates configuring CAMP for scenarios driven by AMBRS. In addition, we develop an interface between MAM4 and CAMP that enables extended chemistry relative to

MAM4's native treatment. The AMBRS CAMP interface streamlines CAMP configuration so that PartMC can be configured for like-to-like comparisons with reduced aerosol models like MAM4. The MAM4 CAMP interface enables varying the structure of MAM4's chemistry so that the impacts of its simplifying assumptions can be quantified. We discuss policy implications for this type of analysis.

To Dr. Gabrielle Miranda Hobson

Acknowledgments

Science and engineering are collaborative endeavors; pursuing a doctorate doubly so. It takes a village, and I would not be where I am today without the support of my mentors, peers, colleagues, friends, and family.

I would first like to thank my research advisor, Professor Cesunica Ivey, who welcomed me into her group based on a few early conversations about shared research interests. It was this alignment that convinced me to stay on for my Ph.D., and Professor Ivey's unwavering trust and encouragement propelled me forward through my degree. She afforded me the freedom to bring my skills to bear on a variety of problems and to explore diverse methods to tackle them, and for this I am greatly appreciative; I am a better researcher for it.

I also want to thank my other committee members: Rob Harley, whose rigorous lines of questioning deepened my understanding of air quality engineering; and Tina Chow, who demonstrated that modeling need not be abstracted from the reality of people's lived experience. My training as an environmental engineer was further supported by the other world-class professors whose classes I took: David Sedlak, Josh Apte, Mark Stacey, Inez Fung, Bill Boos, Kathy Hammond, Kristie Boering, and Brian White. My mentors extended beyond Berkeley, and I have learned much from the AMBRS team: Ben Murphy, Laura Fierce, Nicole Riemer, Jeffrey Johnson, and Matt Dawson. Thank you to my mentors at UNC Chapel Hill who introduced me to research: Wesley Hamilton and Jason West.

Alongside my mentors, my peers and colleagues made my degree possible. Thank you to all students in the ENV program, with special thanks to the 2021 cohort, and thank you to the Airheads. My cohort helped me navigate the vagaries of academia and injected joy into the experience, and the Airheads provided a venue to hone my skills as a science communicator and have fun doing so.

The ENV community is a friendly and welcoming one. I cannot list every name that has made my degree so enjoyable to obtain. Special thanks to Libby Hurtado Koolik, Cam Phelan, Yi Ji, Sumu Prasad, Chirag Manchanda, Paul Seibert, Anthony DeSalvo, Aydin Eskafi, Haley Byrne, Will McNeil, Caroline McCormack, Hannah Wharton, Camille Brule, Jeremy Lowe, Matt Falcone, Alex Georges, Jono Uhler, Minerva Teli, Alma Bartholow, and honorary ENVies Kelly Clevenson, Vincent Laroche, Ian Kolaja, and David Li.

Last but not least, none of this would have been possible without my family. Thank you to Renzo, Nina, Kat, and Conrado, and the most special of thanks to Gabrielle Hobson, my favorite person in the world and my partner in all things.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background

What is air quality?

Air quality refers to the cleanliness of the air we breathe. Poor air quality is characterized by contamination with pollution. When we think of pollution, we might immediately imagine oil spills, wastewater flowing into surface water, chemicals like DDT and GenX, or thick black plumes emitting from smokestacks, big rig trucks, and wildfires. While all of these are examples of pollution that threatens environmental and human health, the listed examples that are most directly relevant for air quality are the latter examples: industrial activity, vehicle exhaust, and, increasingly, wildfires. These are sources of air pollution. Broadly, air pollution can be separated into two categories: primary and secondary.

Primary air pollutants are emitted directly into the atmosphere while secondary air pollutants are formed through chemical reactions in the atmosphere. Industrial activity is responsible for an array of primary pollutants, including soot, also known as black carbon (BC); heavy metals like lead (Pb) and arsenic (As); hazardous air pollutants (HAPs) like benzene, toluene, ethylbenzene, and xylene (BTEX); and per- and polyfluorinated alkyl substances (PFAS), like PFOS, PFOA, and GenX. BC is generally emitted by combustion activity, like burning coal for power generation or running gasoline and diesel engines, and is associated with adverse health outcomes due to its tendency to sorb toxic chemicals (Janssen et al. 2011). Heavy metals are emitted by heavy industry, like smelting, refining, and mining, and are associated with increased morbidity and mortality from direct toxicity (Li, Qian, and Q. Wang 2013) and may sorb to BC (Hu et al. 2009). HAPs like BTEX are potent carcinogens and endocrine disruptors emitted by fossil fuel combustion, industrial solvent use, and chemical synthesis (Bolden, Kwiatkowski, and Colborn 2015). PFAS are of growing interest, with increasing evidence of adverse health effects like immune, thyroid, and endocrine disruption, reproductive and developmental disorders, and cancer (Fenton et al. 2026), and are emitted by a wide array of consumer products and industrial activities such that they are distributed ubiquitously in the atmosphere, even in the most remote locations

(Faust 2023). Vehicle exhaust emits BC and HAPs just like industrial activity (Hu et al. 2009; Bolden, Kwiatkowski, and Colborn 2015), while brake and tire wear contribute metal emissions alongside additional particulate carbon as well as silicon, phosphorus, sulfur, and chlorine (Garg et al. 2000). Vehicles additionally emit nitrogen oxide (NO) and hydrocarbons, which are important precursors for a range of secondary pollutants (Gentner et al. 2017). Wildfire emissions also emit NO and hydrocarbons, as well as BC, and when the fire spreads to infrastructure, the smoke can contain a complex, highly reactive, highly toxic mixture of particulate matter and HAPs that can participate in secondary formation (Fent et al. 2018).

Primary pollutants like NO and hydrocarbons can react to form secondary nitrogen dioxide (NO_2), which is associated with adverse health effects and can go on to produce ground-level ozone (O_3). While O_3 is protective in the stratosphere, blocking harmful ultraviolet radiation, it is hazardous at the surface (Lippmann 1991). Hydrocarbons can continue to oxidize to form carbon monoxide (CO) and eventually carbon dioxide (CO_2), the predominant greenhouse gas (GHG). Coal power generation also emits sulfur compounds, which can react from sulfur dioxide (SO_2) to sulfuric acid (H_2SO_4) and partition into the particle phase to produce acidic aerosol. Moreover, HAPs like BTEX can be formed by chemical reactions in the atmosphere (Kelly et al. 1994), wildfire emissions can enhance downwind O_3 (D. Jaffe et al. 2008; D. A. Jaffe et al. 2013; Ji et al. 2025), and primary volatile organic compounds (VOCs) and particulate matter (PM) provide precursors and seed particles for secondary particle mass accumulation (Seinfeld and Pandis 2016).

Whether primary or secondary, the phase of matter of the pollution is highly important. Air pollution can be gaseous, liquid, or solid, and can be dissolved or not. Liquids and solids form aerosol particles. One might think of aerosol as being solid particles, but liquid particles are possible, and it is often the case that aerosol particles consist of both liquids and solids: a solid core covered by a liquid layer. Indeed, this is how cloud droplets form. The phase, gas or aerosol, is important for a variety of physical and chemical reasons that ultimately define a pollutant's toxicity and environmental impact.

In 1948 in Donora, Pennsylvania, and 1952 in London, England, major smog events took place. The smog was so thick that daytime seemed like twilight and epidemiological studies report spikes in morbidity and premature mortality associated with these events: 20 people died in Donora with another 5,900 who experienced respiratory distress during the incident, while cardiovascular disease and cancer rates in the subsequent decade were nearly double expectations (Jacobs, Burgess, and Abbott 2018). In the two years following the London event, up to 12,000 incidents of premature mortality can be associated with smog exposure (Bell and D. L. Davis 2001). In Donora, the sources of the smog were the American Steel and Wire and the Donora Zinc Works plants, which suggest a singularly toxic smog laden with heavy metals from smelting and refining activity. In London, in addition to heavy industry, sulfur emissions from extensive coal burning contributed to the smog's toxicity (D. Davis 2002).

Ambient air quality in the United States has seen significant improvement since the Donora event due to regulatory policy like the United States National Ambient Air Quality

Standards (NAAQS), passed in 1971, which regulate ambient levels of O_3, NO_X, SO_2, CO, particulate matter (PM), and lead (Pb). The NAAQS have led to emissions reductions of 71%, 93%, and 88%, respectively, for CO, SO_2, and Pb between 1990 and 2024, leading ambient concentrations to drop by over 80% each (US EPA 2026). PM and O_3, however, have proven more stubborn, with ambient concentrations dropping less than 30% each between 1990 and 2024 (US EPA 2026). Given the health impacts of PM and O_3, their persistence in ambient air emphasizes the need for continued improvement. Indeed, although air quality improvements have been significant at a national scale, the benefits have not been equitably distributed. In California, despite a 65% reduction in statewide exposure to vehicle-emitted PM with diameter less than 2.5 μm ($PM_{2.5}$), the relative exposure disparity between racial groups has largely stagnated, even increasing slightly for some groups since 2015 (Koolik et al. 2024).

This disproportionate exposure burden experienced by marginalized communities is an important dimension to consider when discussing the health effects of air pollution. Environmental justice was expounded by Bullard (1996) as follows:

> Environmental justice embraces the principle that all people and communities are entitled to equal protection of environmental and public health laws and regulations.

Environmental justice was defined as the remedy for environmental racism, which Bullard (1996) defines as:

> Environmental racism refers to any policy, practice, or directive that differentially affects or disadvantages (whether intended or unintended) individuals, groups, or communities based on race or color.

Environmental racism, though long experienced and resisted by marginalized communities, entered the academic and political discourse in 1982 when communities in Warren County, North Carolina organized against planned dumping of soil contaminated with polychlorinated biphenyls (PCBs) into the water table that supplied the communities' wells (Mohai, Pellow, and Roberts 2009). PCBs are highly toxic, associated with cancer and adverse endocrine and neurological effects (Quinete et al. 2014). Environmental racism persists, with marginalized communities often located near emissions sources due to historical redlining, a discriminatory mortgage appraisal practice, leading to the siting of polluting activity in close proximity to marginalized communities (Lane et al. 2022). Therefore, the problem of air quality is a problem of environmental stewardship, public health, and social equity.

How do we quantify air quality?

We quantify air quality by two primary means: measurement and modeling. Measurement techniques leverage sophisticated instrumentation to quantify the composition of real-world air. Measurements are made for ambient air, indoor air, and personal microenvironments

during field campaigns, laboratory experiments, and monitoring initiatives. Measurements are highly labor intensive, leading to temporally sparse data, measured with coarse time resolution, or sporadic data, with fine time resolution but short duration. Monitoring data tends to be sparse while field campaigns and laboratory experiments are more sporadic. Moreover, measurement data are generally sparse in space, representing point measurements made at specific study sites. Spatially dense measurements can be made by aircraft and mobile laboratories, but the extent of their spatial coverage is typically restricted. Despite these caveats, measurements represent the gold standard for quantifying atmospheric composition. Measurements reflect the real world, and often probe the limits of our knowledge, expanding our understanding of the atmosphere.

Complementary to measurement is modeling. Air quality modeling is more spatiotemporally complete than measurement, with fine time resolution and complete coverage of a spatial domain. The caveats here tend to relate to the density of spatial coverage, the computational expense of simulations, and the necessarily idealized representations of physics and chemistry that models use to make simulation tractable. While complete, the spatial resolution of large-scale models tends to be coarse, generally at tens of kilometers for regional-scale models or up to over one hundred kilometers for global-scale models. This coarseness is a direct result of computational restrictions: high-resolution spatial domains are more compute intensive to simulate. Similarly, the complexity of model representations of physics and chemistry is limited by compute intensity, as well as by the limits of our knowledge. Chemical mechanisms, the representations of coupled systems of chemical reactions, scale compute intensity with the number of reactions in the system, and reactions can only be included in a mechanism once their reaction sequences have been characterized by experimental or computational methods. Despite these caveats, models provide the means for predicting air quality where measurements are not available; forecasting air quality; and projecting air quality forward in time to explore potential emissions scenarios.

Why does quantification matter?

Without quantification, policymakers lack the information required to design effective regulations. Communities burdened by poor air quality recognize the problem and advocate for themselves, but policymakers do not act without data. People understand from their lived experiences that air pollution is a problem, but policymakers do not often live in the most affected communities. The findings of Tessum et al. (2021), who report that people of color systematically experience disproportionate $PM_{2.5}$ exposure across virtually all source categories, and Koolik et al. (2024), who show that the trends in that disproportionate exposure remain systematically biased, call attention to the persistence of the problem, even when the national picture shows marked improvement. Y. Wang et al. (2022) shows that a focus on national-scale strategies cannot solve the issue; the most effective regulation is local. To design such targeted controls, we need to quantify air pollution emissions and air quality to identify sources and design strategies for maximal efficacy.

1.2 Dissertation Structure

In this dissertation, we explore a variety of modeling techniques and developments, each of which can directly or indirectly influence regulatory policy. We first develop a workflow to apply statistical techniques for characterizing regulatory monitor PM measurements and adapt the workflow to field campaign measurements, ultimately producing results that can inform local policy controls. Next, we demonstrate an accelerated implementation of an air quality model's chemistry module with implications for improving air quality forecasting and increasing model grid resolutions, which can support improved air quality policy that is informed by finer scale dynamics. Finally, we present our contributions to an aerosol model benchmarking initiative, which provide functionality for quantifying structural errors that contribute to inter-model variability. This is a necessary step towards constraining that variability to reduce uncertainty in climate projections, which can increase confidence in model predictions of climate change under different emissions scenarios.

In Chapter 2, we develop a workflow for training statistical models to predict 24-hour average $PM_{2.5}$ and quantify the response of $PM_{2.5}$ to predictor covariates and how influential the different covariates are (Quevedo, Z. Gao, et al. 2024). The models, which are generalized additive models (GAMs), learn relationships between $PM_{2.5}$ and predictors like NO_X, formaldehyde (HCHO), O_3, wind speed and direction, temperature, and relative humidity, driven by many years of regulatory monitor data. At each distinct study site in California, we train a distinct model for $PM_{2.5}$ and speciated components like SO_4^{2-}, NO_3^-, organic carbon (OC), elemental carbon (EC), and ammonium (NH_4^+). The models yield $PM_{2.5}$ responses as marginal effects, and sensitivity analysis of the models ranks each one's covariates to identify the most explanatory associations. This site-by-site approach enables me to make location-specific policy recommendations.

Chapter 3 extends the workflow of Chapter 2 to predict hourly average organic aerosol (OA), as opposed to the 24-hour average $PM_{2.5}$ predicted in Chapter 2. We take a similar site-by-site approach to training GAMs and again compute marginal effects as well as variable importance rankings. The higher temporal resolution of this data, measured during a field campaign that took place across three study sites in California, enables accounting for diurnal variation, which we use to control for noise and trends in the data set. With an additional layer of model selection applied during training on a more chemically-resolved set of predictor covariates, the models make predictions that we use to recommend local policy strategies for controlling OA.

We demonstrate improved computational efficiency for solving chemistry in the Community Multiscale Air Quality (CMAQ) model in Chapter 4. To achieve accelerated chemistry solves, we port CMAQ's Rosenbrock-3 solver Fortran source code to CUDA Fortran, which enables computation on NVIDIA graphics processing unit (GPU) hardware (Quevedo, Do, et al. 2025). Our implementation achieves notable acceleration, demonstrating the viability of this approach, which saves compute time while preserving the underlying mathematical representation of the system. This work has implications for air quality forecasting, which is generally unable to simulate complex chemistry since non-accelerated chemistry requires

too much compute time to provide forecast results on a timely schedule. Another benefit of accelerated chemistry is increased horizontal grid resolutions, since this translates to a corresponding increase in compute time. Improved forecast fidelity and spatial resolutions support air quality policy informed by up-to-date and finer-scale air pollution dynamics, corresponding more closely with what communities actively experience.

Our final project is encapsulated in Chapter 5, where we present our contributions to the Aerosol Model Benchmark Repository and Standards (AMBRS) Project. AMBRS is a platform that facilitates aerosol model benchmarking, with the goal of engaging the aerosol modeling community to adopt model performance standards that reduce inter-model variability. Inter-model aerosol variability remains the largest source of uncertainty in climate projections, emphasizing the need to constrain variability to increase confidence in those projections. Our contributions help facilitate inter-model benchmarking and comparison, which can be used to quantify structural errors in aerosol models that contribute to inter-model variability, helping identify sources of bias that can be targeted for correction.

This dissertation collects a number of novel contributions. Chapters 2 and 3 introduce a global variance-based sensitivity analysis technique to the field of air quality modeling and leverage this new method to provide the state of California with new evidence to support its air quality policy. Chapter 4 introduces the first implementation of a GPU-accelerated chemistry solver in CMAQ and it is the first work to couple GPU-accelerated chemistry back to the full model system to produce an operational CTM deployed on heterogeneous architecture. Chapter 5 expands the chemical complexity of an aerosol module commonly found in major Earth system models to enable robust quantitative interrogation of model design choices. We conclude in Chapter 6 with a summary of this dissertation and we outline future research directions.

Chapter 2

Analyzing Particulate Matter in California

This chapter is adapted from:

Duncan Quevedo, Ziqi Gao, et al. (Jan. 2024). "Multidecadal Analysis of Meteorological and Emissions Regimes for PM2.5 Across California". In: *ACS ES&T Air* 1.1, pp. 33–42. DOI: 10.1021/acsestair.3c00019

The dissertation author was the primary investigator and author of this paper.

2.1 Introduction

As of 2025, large swathes of California remain in nonattainment of the Environmental Protection Agency (EPA)'s National Ambient Air Quality Standards (NAAQS) for particulate matter (PM) (*Maps of State and Federal Area Designations — California Air Resources Board* 2022). $PM_{2.5}$ nonattainment is concentrated around the state's most populous regions (*Maps of State and Federal Area Designations — California Air Resources Board* 2022). This puts many millions of California residents at risk of adverse health effects from $PM_{2.5}$ exposure. $PM_{2.5}$ is associated with adverse health effects including morbidity and mortality from respiratory and cardiovascular diseases (Douglas W. Dockery et al. 1993; Pope et al. 1995; Atkinson et al. 2001; Halonen et al. 2008; Zanobetti, Schwartz, and D W Dockery 2000; Raizenne et al. 1996; Ackermann-Liebrich et al. 1997). Furthermore, the EPA has recently revised the $PM_{2.5}$ annual standard to 9 μg m^{-3} down from its previous value of 12 μg m^{-3} (EPA 2023). For California to meet either of these standards and safeguard public health, the forces driving nonattainment must be well characterized so that effective control strategies can be designed.

Factors influencing high ambient $PM_{2.5}$ in California include emissions and meteorology. Emissions of nitrogen oxides (NO_X) and sulfur oxides (SO_X) are oxidized to nitrate (NO_3^-)

and sulfate (SO_4^{2-}), respectively, while ammonia (NH_3) emissions form ammonium (NH_4^+). In addition to these inorganic aerosol species, emissions of reactive organic gases (ROG) provide reactants for secondary organic aerosol (SOA) formation. Further, different aerosol species dominate $PM_{2.5}$ pollution in different regions of California. In addition to the geographic variability in $PM_{2.5}$ mass profiles, meteorology varies throughout the state and plays a role in high ambient $PM_{2.5}$ levels (Zhu et al. 2019). Wind, humidity, and temperature patterns differ, for instance, between coastal regions like the South Coast Air Basin (SoCAB) and inland regions like the San Joaquin Valley (SJV), leading to differential meteorological modulation of $PM_{2.5}$ formation regimes. As such, $PM_{2.5}$ levels depend on a complex interplay between numerous environmental factors, precluding an immediate understanding of the interaction of meteorological and emissions drivers and their role in $PM_{2.5}$ nonattainment in California.

The association between $PM_{2.5}$ and a set of predictor covariates can be inferred through the use of generalized additive models (GAMs) (Z. Gao et al. 2022; Pearce, Beringer, Nicholls, Hyndman, and Tapper 2011). GAMs enable us to capture complexities in the statistical relationships between $PM_{2.5}$ and covariates by fitting nonlinear functions to the covariates of interest. Corresponding to each covariate is a nonlinear fit that can be analyzed independently of others' fits, elucidating the distinct effects of each individual covariate. Moreover, with variable importance ranking we can identify which covariates are most important for modeling $PM_{2.5}$, thereby identifying potential drivers of high ambient $PM_{2.5}$ levels. This makes the problem of understanding high ambient $PM_{2.5}$ levels in California more tractable.

Knowledge of how different covariates are associated with $PM_{2.5}$ throughout California and specifically which covariates drive high ambient $PM_{2.5}$ levels can inform regulatory policy design. By taking geographic variability into account and focusing on driving variables, policymakers can better target and tailor regulations to control $PM_{2.5}$.

This regional approach is well suited to informing programs like the California Air Resources Board (CARB)'s Community Air Protection Program. Implemented in response to Assembly Bill (AB) 617, which seeks to alleviate air pollution impacts on environmental justice communities, the Community Air Protection Program supports community-level efforts to monitor air quality and reduce emissions. As the locations we consider in this paper are all located in or near AB 617 communities, our results can help inform Community Air Protection Program strategies.

Our main goal is to develop models that characterize the historical $PM_{2.5}$ response to different chemical and meteorological influences in California. Exactly how $PM_{2.5}$ responds to different variables and how that response differs across the state are not well understood. While Vutukuru, Griffin, and Dabdub (2006) have characterized simulated SOA responses as functions of certain covariates, their analysis examines just one meteorological and two emissions covariates and is limited to the Los Angeles area. More recent work by Nussbaumer and Cohen (2021) produced temperature trends for $PM_{2.5}$ in the Los Angeles Basin. Numerical modeling by Zhu et al. (2019), however, shows that emissions and meteorological impacts on $PM_{2.5}$ vary by air basin, indicating that the response curves of Vutukuru, Griffin, and Dabdub (2006) and Nussbaumer and Cohen (2021) may not be generalizable to the rest of California. Our work extends the characterization of $PM_{2.5}$ response curves to more of

California and a greater number of covariates. Moreover, our work incorporates measured rather than simulated data, leveraging California's extensive measurement network. In this chapter, we present a series of GAMs for total and speciated $PM_{2.5}$ at Chemical Speciation Network (CSN) sites in the state of California. These models characterize $PM_{2.5}$ responses to chemical and meteorological variables throughout the state and how that response differs from site to site.

2.2 Materials and Methods

Study Area

We construct GAMs at 6 of the 7 CSN sites in the state of California over many years of data (Figure 2.1). These sites are located in Bakersfield (2004-2019), Fresno (2012-2019), Los Angeles (2002-2019), Riverside (2001-2019), Sacramento (2010-2019), and San Jose (2011-2019). The El Cajon site is omitted because of its limited data archive at the time of this study. From these sites we obtain $PM_{2.5}$ mass and speciation data. We also utilize data from meteorological stations in the proximity of these CSN sites, as well as from the nearest radiosonde stations or the North American Mesoscale (NAM) historical analysis product.

Bakersfield and Fresno are located in the San Joaquin Valley, which has topography characterized by mountain ranges to the east, west, and south that inhibit the transport of air pollutants out of the valley. Climate in the San Joaquin Valley is characterized by high temperatures, low humidity, and sparse rainfall (Marjollet et al. 2015). NO_X and VOC emissions in the San Joaquin valley are dominated by mobile sources and oil and gas production, respectively, while primary $PM_{2.5}$ emissions come mainly from road and agricultural dust (*ARB Almanac 2013 - Chapter 4: Regional Trends and Forecasts* 2023).

Los Angeles and Riverside are found in the South Coast Air Basin, which is bounded by the Pacific Ocean on the west and mountains to the north, south, and east. The regional climate is characterized by warm, sunny days and stagnant air conditions exacerbated by frequent inversions (*ARB Almanac 2013 - Chapter 4: Regional Trends and Forecasts* 2023). NO_X and VOC emissions are dominated by on-road motor vehicles while primary $PM_{2.5}$ emissions come predominantly from commercial cooking (*ARB Almanac 2013 - Chapter 4: Regional Trends and Forecasts* 2023; Cheung et al. 2023).

Sacramento is located in the northern Central Valley with topography characterized by mountains to the east and west, the San Francisco Bay to the southwest, and the rest of the Central Valley to the north and southeast. Its climate sees cool, wet winters and hot, dry summers, with emissions dominated by on-road motor vehicles (*ARB Almanac 2013 - Chapter 4: Regional Trends and Forecasts* 2023).

San Jose is in the San Francisco Bay Area, which is a coastal area dominated by the San Francisco Bay, leading to year-round mild temperatures and good ventilation. Emissions are dominated by on-road motor vehicles for NO_X and VOC and wood burning for primary

$PM_{2.5}$ (*ARB Almanac 2013 - Chapter 4: Regional Trends and Forecasts* 2023; *Understanding Particulate Matter: Protecting Public Health in the San Francisco Bay Area* 2012).

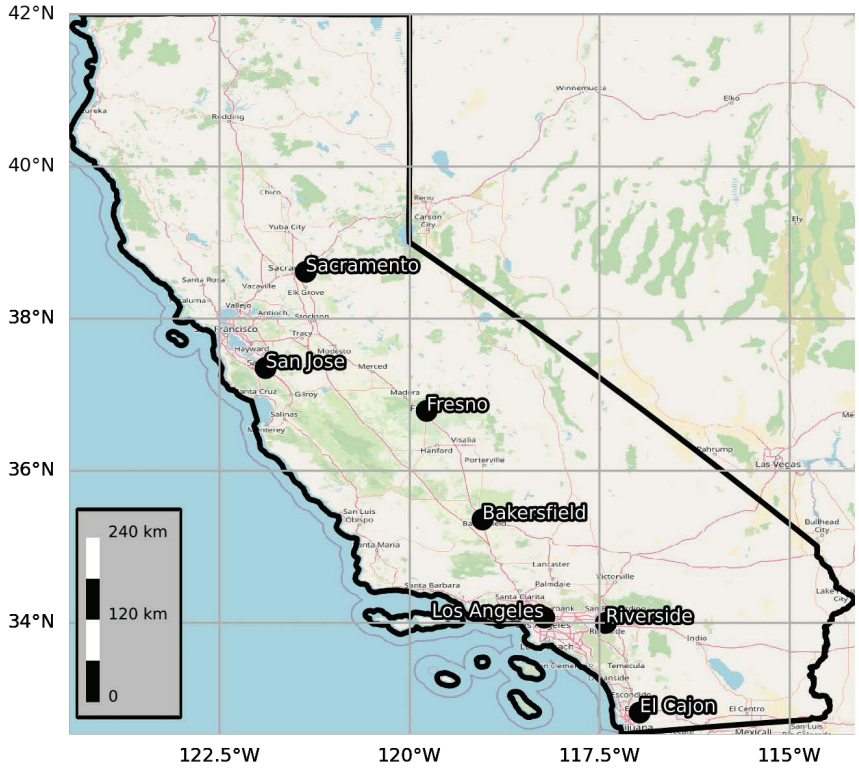

Figure 2.1: Map of study area showing locations of CSN sites in California.

Choice of Covariates

We investigate a set of covariates chosen based on existing work as well as first principles (Z. Gao et al. 2022; Pearce, Beringer, Nicholls, Hyndman, and Tapper 2011; Blanchard et al. 2019; Ivey et al. 2022; Seinfeld and Pandis 2016). Our models include meteorological and gaseous covariates. Meteorological covariates include surface and upper-air measurements. Gaseous covariates included are those measured at collocated photochemical assessment monitoring stations (PAMS) and play a role in secondary PM formation.

Air Pollution Data

Air pollution data are obtained from the EPA Air Quality System (AQS). We retrieve daily summary data for all species, although species measurements are not available every day. Species of interest include total $PM_{2.5}$ mass ($PM_{2.5}$), organic carbon (OC), elemental carbon (EC), nitrate (NO_3^-), sulfate (SO_4^{2-}), and ammonium (NH_4^+). We use $PM_{2.5}$ as measured at selected $PM_{2.5}$ Monitoring Network sites and its speciated components as measured at co-located CSN sites, which all sample for 24 hours at a collection frequency of 1-in-3 days (*40 CFR 58.12 – Operating schedules.* 2022). The $PM_{2.5}$ Monitoring Network sites under consideration sample $PM_{2.5}$ either with Federal Reference Method (FRM) R&P Model 2025 $PM_{2.5}$ Sequential Air Samplers or Federal Equivalent Method (FEM) MetOne BAM-1020 samplers. CSN sites sample 33 aerosol species on Teflon filters and ions on nylon filters in MetOne SASS/SuperSASS samplers. Quartz filters in URG3000N samplers are used to sample carbon. Both carbon species, EC and OC, are measured using the thermal optical transmittance (TOT) method. OC is corrected for sampling artifacts according to Enayati Ahangar et al. (2021): we linearly regress OC against $PM_{2.5}$ and calculate corrected OC by subtracting the model's intercept and omitting negative values.

In addition to particulate matter data, gaseous species are also retrieved from AQS. These include formaldehyde (HCHO), nitrogen oxides (NO_X), and ozone (O_3) as representatives of reactants in secondary $PM_{2.5}$ formation or, in the case of HCHO, as a tracer of those reactants. These species are measured at PAMS via automated gas chromatography coupled with a mass spectrometer or flame ionization detector (GC-MS or GD-FID), an automated chemiluminescence detector, or a UV absorbance monitor, respectively (Hafner and Penfold 2018).

Meteorological Data

All surface meteorological data are retrieved from California Air Resources Board (CARB)'s Air Quality and Meteorological Information System (AQMIS) or meteorological monitors available at CSN co-located PAMS (*Air Quality Data (PST) Query Tool* 2022). From AQMIS we retrieve daily maximum solar radiation from the meteorological monitors located nearest to the CSN site under consideration. Daily average relative humidity, daily maximum surface temperature, and daily average 10-meter wind speed and direction are retrieved from PAMS.

Upper-air meteorology observations, defined as meteorological data measured at the 850 mb air pressure level, are retrieved from the NOAA Earth System Research Laboratory (NOAA/ESRL) Radiosonde Database (*NOAA/ESRL/GSL - RAOB* 2022). We use upper-air meteorology at this pressure level following Blanchard et al. (2019) to quantify potential transport effects (Blanchard et al. 2019). The radiosonde station chosen for each CSN site is the nearest one with a data archive that does not restrict our timespan. From NOAA/ESRL, we retrieve wind speed, wind direction, temperature, and dew point depression data. Dew point is recovered from the depression and combined with temperature to calculate upper-air

relative humidity using the August-Roche-Magnus equation (Equation 2.1) (Alduchov and Eskridge 1997):

$$e_s(T) = 6.1094 \ \exp\left(\frac{17.625\ T}{T + 243.04}\right) \tag{2.1}$$

where e_s is saturation vapor pressure and T is temperature in Celsius. Based on Equation 2.1, relative humidity is calculated as

$$RH = 100\% \times \frac{e_s(T_d)}{e_s(T)} \tag{2.2}$$

where T and T_d are temperature and dew point, respectively (Seinfeld and Pandis 2016).

In addition to observed upper-air meteorology, we incorporate modeled upper-air data because between the Fresno and Bakersfield CSN sites and their nearest active radiosonde launch sites are mountain ranges that may change the upper-air weather patterns between the CSN sites and the radiosonde launch sites. See Table 2.2 for a model performance comparison between observed and modeled upper-air data at sites whose final models used radiosonde data.

For our modeled upper-air data we use the North American Mesoscale (NAM) 12-kilometer historical analysis product for the same covariates at the same pressure level as our radiosonde data, namely wind speed and direction (recovered from the U and V components of wind) as well as relative humidity at the 850 mb pressure level.

Generalized Additive Models

We employ the generalized additive model (GAM) in this study, introduced by Hastie and Tibshirani (1986). We use GAMs because of their ability to capture nonlinear relationships between predictor covariates and the target variable for how they facilitate a covariate-by-covariate analysis of the resulting models (Z. Gao et al. 2022; Pearce, Beringer, Nicholls, Hyndman, and Tapper 2011). A GAM is analogous to a generalized linear model (GLM) of the form

$$g(\mu) = \beta_0 + \sum_i \beta_i x_i + \epsilon \tag{2.3}$$

where g is the link function that represents the relationship between the covariates x_i and the expected value μ of the target variable, where the β_i are the fitted model coefficients, β_0 is the intercept, and ϵ are the residuals. When g is the identity function, we have multiple linear regression, whereas we have logistic regression in the case where g is the logit link $g(p) = \log\left(\frac{p}{1-p}\right)$. Many choices for the link function exist.

In contrast to the form of a GLM, which fits constant coefficients of the covariates x_i, a GAM fits functions of x_i. GAMs therefore have the form

$$g(\mu) = \beta_0 + \sum_i f_i(x_i) + \epsilon \tag{2.4}$$

where the functions f_i are fit with smoothing algorithms and typically take the form of splines. The fits are therefore data driven and capture nonlinear relationships naturally.

We fit our GAMs using the package `mgcv` in the R environment for statistical computing (Team 2021; S. N. Wood 2000; Simon N. Wood 2011). We use `mgcv`'s `gam` function to fit our models. For full algorithm details see Simon N. Wood (2011) and Simon N. Wood, Pya, and Säfken (2016). Briefly, `gam` constructs and iteratively penalizes basis functions for each smoothing spline using one of several smoothing parameter selection criteria such as generalized cross validation (GCV) or restricted maximum likelihood (REML) scores (*R: Generalized additive models with integrated smoothness...* 2022). For the model specification detailed next, we choose REML as our scoring criterion because it is less prone than GCV to selecting local minima rather than the desired global minimum in the score space (*R: Generalized additive models with integrated smoothness...* 2022).

Model Description

Our base model is

$$\begin{aligned}\hat{y} = \log(\mu) = \beta_0 + \; & s(\texttt{TMAX}) \; + \; s(\texttt{WINDS}) \; + \; s(\texttt{WINDD}) \\ & + \; s(\texttt{RH}) \; + \; s(\texttt{SR}) \; + \; s(\texttt{RH850}) \\ & + \; s(\texttt{WS850}) \; + \; s(\texttt{WD850}) \; + \; s(\texttt{NOX}) \\ & + \; s(\texttt{HCHO}) \; + \; s(\texttt{O3}) \; + \; \epsilon\end{aligned} \tag{2.5}$$

where $\hat{y}$ is the model fit and the argument of the log link, μ, is the expected value of the target variable, which is 24-hour average $PM_{2.5}$ mass or one of its speciated components. We use the log link in our models following precedent (Z. Gao et al. 2022; Pearce, Beringer, Nicholls, Hyndman, and Tapper 2011; Aldrin and Haff 2005). Moreover, we found the log link to reduce bias in our models' fits against annual 98^{th} percentile values as compared with the identity link. As before, β_0 is the model intercept and ϵ is the residual. We express spline fits as $s(\cdot)$. Table 2.1 summarizes the covariates.

Each fit $s(\cdot)$ is a cubic regression spline (cyclic for wind direction) with an extra shrinkage penalty that enables the spline to be penalized to zero at sufficiently high smoothing parameters (*R: Generalized additive models with integrated smoothness...* 2022). This manifests as automatic model selection, zeroing out and essentially dropping terms that the smoothing process heavily penalizes, which are the terms that do not contribute explanatory power to the model. This facilitates a data driven approach to distinguish on a site-by-site basis which covariates drive $PM_{2.5}$ levels and which have less impact.

A further important detail about our model specification is the number of knots chosen for the spline fits. The choice of knot quantity has consequences for fit, with too many knots often resulting in overfitting and too few knots leading to bias from underfitting (Perperoglou et al. 2019). In all of our models, we fit splines with 4 knots. While our use of penalization would help control overfitting if we used many knots, we found no improvement in models fit with more than 4 knots per spline.

Table 2.1: Summary of terms in Equation 2.5.

Covariate	Description	Units
`TMAX`	Daily maximum surface temperature	°C
`WINDS`	Daily average 10 m wind speed	m s^{-1}
`WINDD`	Daily average 10 m wind direction	Degrees from north
`RH`	Daily average surface relative humidity	%
`SR`	Daily maximum solar radiation	W m^{-2}
`RH850`	850 mb daily average relative humidity	%
`WS850`	850 mb daily average wind speed	m s^{-1}
`WD850`	850 mb daily average wind direction	Degrees from north
`NOX`	Daily average nitrogen oxides concentration	ppb
`HCHO`	Daily average formaldehyde concentration	ppb of carbon (ppbC)
`O3`	Daily average ozone concentration	ppb

Finally, we need to specify an error distribution. Defaulting to a Gaussian distribution produces heteroscedastic residuals with predictions that are biased low at large observations of the target variable. The resulting error distribution suggests a skewed distribution is most appropriate. Therefore, we choose a Gamma distribution, and the resulting models produce homoscedastic residuals.

Model Performance Evaluation

Model performance metrics reported are calculated by 10-fold cross validation. Performance metrics include mean bias (MB), root-mean-square error ($RMSE$), and coefficient of determination (R^2):

$$MB = \frac{1}{N}\sum_{i=1}^{N}(\hat{y}_i - y_i) \tag{2.6}$$

$$RMSE = \sqrt{\frac{\sum_{i=1}^{N}(\hat{y}_i - y_i)^2}{N}} \tag{2.7}$$

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(\hat{y}_i - y_i)^2}{\sum_{i=1}^{N}(y_i - \overline{y})^2} \tag{2.8}$$

where $\hat{y}_i$ are the modeled values, y_i are the observed values, N is the number of data points, and $\overline{y}$ is the mean of the observed data. We use the standard R^2 over the adjusted R^2 because our models are penalized and zero out non-explanatory covariates.

Variable Importance

While penalization in the model fitting algorithms helps identify driving covariates, we additionally employ a variable importance ranking method to further distinguish between impactful covariates and those that do not hold significant explanatory power. We use a variance-based method of global sensitivity analysis for our variable importance determination. Global sensitivity analysis, as opposed to local sensitivity analysis, better characterizes model sensitivity to covariates in the presence of non-additivity, typically manifesting as covariate interactions (Ferretti, Saltelli, and Tarantola 2016).

The measure of importance we calculate is the Sobol' index, which apportions the variance in model output attributable to different covariates. The index of any given covariate is the proportion of total model variance attributable to that covariate (Puy et al. 2021).

To find the first order Sobol' index of a covariate x_i, we first calculate the model output $\hat{y}$ after varying all other covariates with fixed x_i. Borrowing notation from sensitivity analysis, we denote by $\mathbf{x}_{\sim i}$ all covariates except the ith, where "$\sim i$" indicates all except i. We do this for each value of x_i and then take the mean of the variance in model output,

$$E\left[V(\hat{y}|x_i)\right], \tag{2.9}$$

where $E(\cdot)$ and $V(\cdot)$ are the mean and variance operators, respectively. We note that

$$V(\hat{y}) = V\left[E(\hat{y}|x_i)\right] + E\left[V(\hat{y}|x_i)\right] \tag{2.10}$$

where $V(\hat{y})$ is the model variance. Let

$$V_i = V\left[E(\hat{y}|x_i)\right] \tag{2.11}$$

Then the first order Sobol' index for covariate x_i is

$$S_i = \frac{V_i}{V(\hat{y})} \tag{2.12}$$

S_i measures the individual variance contribution of x_i. We can additionally calculate the total index

$$T_i = 1 - \frac{V\left[E(\hat{y}|\mathbf{x}_{\sim i})\right]}{V(\hat{y})} \tag{2.13}$$

The total index T_i measures the sensitivity of the model to x_i and any interactions. If $S_i < T_i$, this suggests the presence of significant covariate interactions. On the other hand, if $S_i \approx T_i$ and $\sum_i S_i \approx 1$, this suggests the model is additive and has no significant interactions. In our models, we impose additivity because interaction terms were not found to improve model fits. Sobol' indices are estimated with the `sensobol` package in R via Monte Carlo methods with 95% confidence intervals bootstrapped (Puy et al. 2021).

Marginal Effects

Additionally, we report the response of $PM_{2.5}$ to individual covariates as marginal effects. We calculate marginal effects as follows. We hold each covariate except that which is under consideration, that is, $\mathbf{x}_{\sim i}$, constant at the observation corresponding to the mean fitted value $\overline{\hat{y}}$. We then make predictions as we vary only the covariate of interest x_i over its observed range. We denote this by

$$ME = 100\% \times \left[\exp\left\{s_i\left(x_i\right) - s_i(x_i|\overline{\hat{y}})\right\} - 1\right] \tag{2.14}$$

where ME is the marginal effect, $s_i(x_i)$ is the spline fit for covariate x_i, and $s_i(x_i|\overline{\hat{y}})$ is the spline fit for covariate x_i at the value of x_i that corresponds to the mean fitted value $\overline{\hat{y}}$. The marginal effect is interpreted as the percent change from the mean fitted value as x_i varies (Pearce, Beringer, Nicholls, Hyndman, Uotila, et al. 2011). To see this, note that the full model would include the exponential of the sum of every spline, but we would subtract each spline evaluated at the value corresponding to the mean fitted value $\overline{\hat{y}}$. Since we are holding $\mathbf{x}_{\sim i}$ constant at that same value, all but two terms cancel, leaving us with Equation 2.14.

2.3 Results and Discussion

Model Performance Metrics

We report metrics for the 24-hour average, 3-year average annual mean, and 3-year average 98^{th} percentile values, the latter two of which are the values regulated by the $PM_{2.5}$ NAAQS. All models predict 24-hour average values, so the metrics reported are for 24-hour values and 24-hour values aggregated to annual and 98^{th} percentile values. Annual and 98^{th} percentile values include an additional rolling average over the preceding two years to produce 3-year backward rolling averages.

We report MB and RMSE to characterize the distribution of fitted values around the observed value, with MB representing a measure of central tendency and RMSE measuring spread (Figure 2.2). Generally, the models perform reasonably well with respect to MB and RMSE, although 24-hour RMSE exhibits a wider range.

The most salient metric for our objective, however, is the coefficient of determination, R^2 (Figure 2.2). R^2 measures the proportion of observed variability captured by the model and therefore measures how well the model characterizes the behavior of the data, which is our primary goal. From this perspective, our models are generally performing well. The worst-performing site has average 24-hour R^2 across species of only 24% but average annual and 98^{th} percentile R^2 of 47% and 44%, respectively, while the best performing site has 24-hour R^2 across species of 70% with annual and 98^{th} percentile R^2 of 76% and 73%, respectively.

In general, our Fresno models ($57\% \leq R^2 \leq 85\%$) perform best while our Sacramento models ($32\% \leq R^2 \leq 59\%$) perform worst (Figure 2.2). Based on R^2, one potential reason for this trend is that our models are missing important sources of variability, such as organic

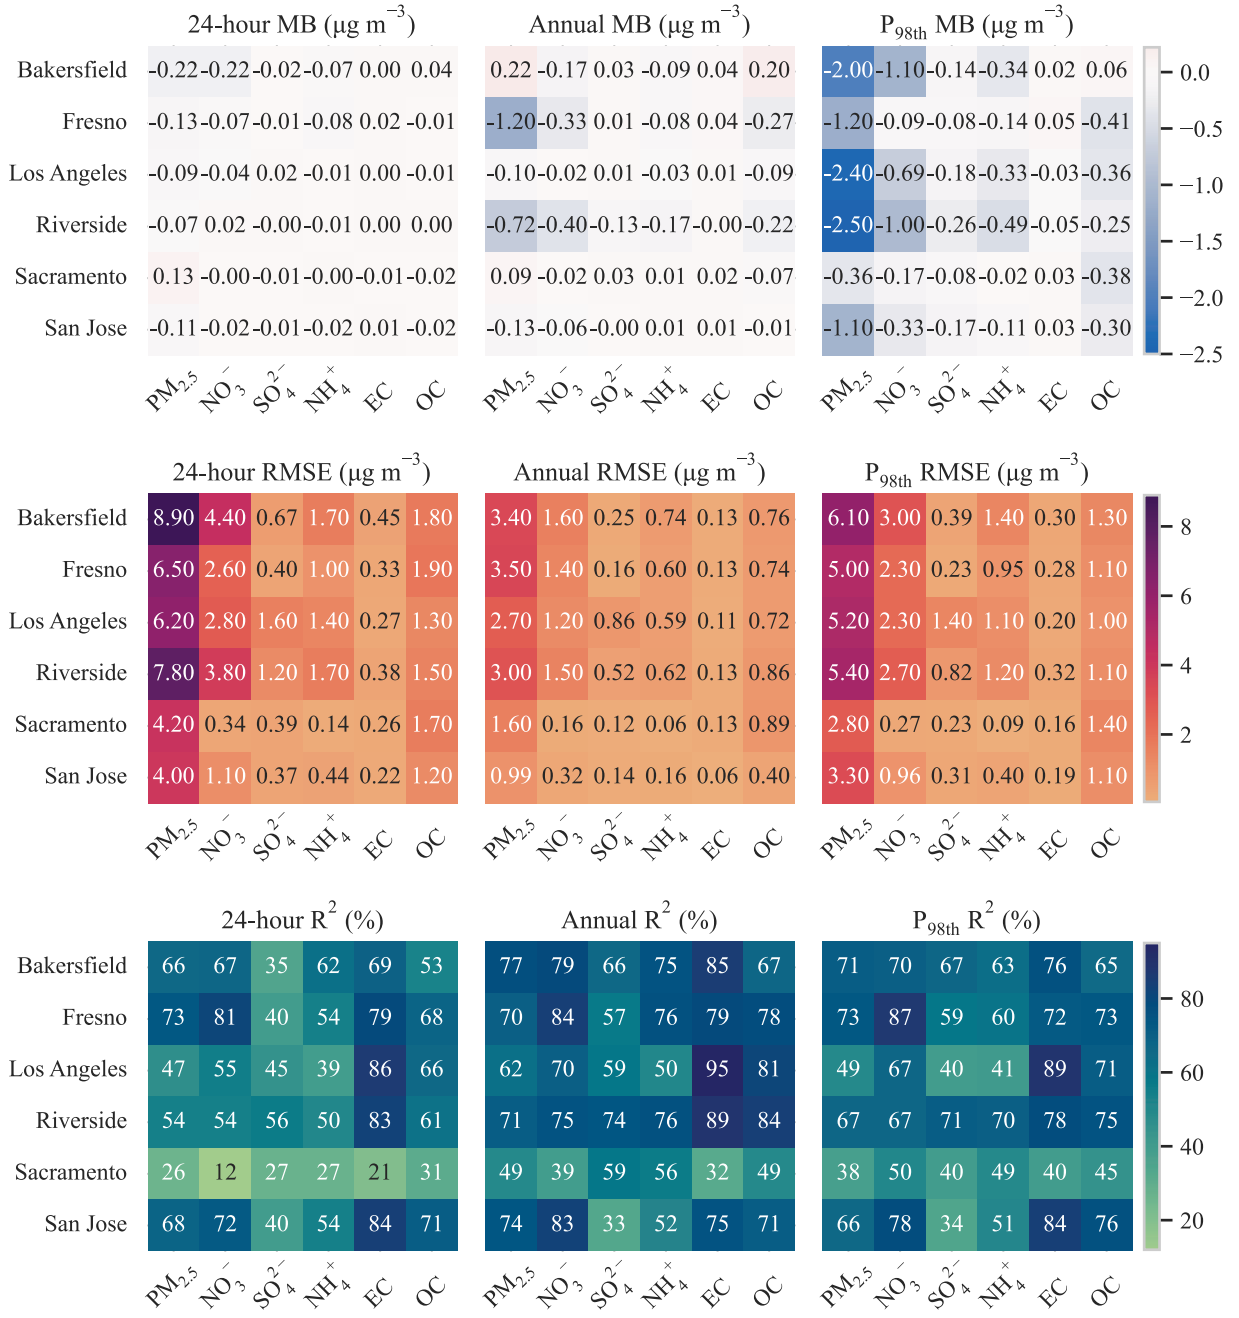

Figure 2.2: Heatmaps of model performance metrics across all sites and species. "24-hour" refers to 24-hour average predictions. "Annual" refers to 3-year average annual means. "P_{98th}" refers to 3-year average 98th percentiles.

vapors or mixing layer height, and the relative importance of these missing sources is greatest in Sacramento.

While our models perform satisfactorily with respect to 24-hour metrics, their performance improves when aggregated to the annual averaging period. This suggests our models may be useful for investigating factors that influence annual NAAQS nonattainment. The 98^{th} percentile model performance is worse than the annual model. Therefore, our models are best suited for characterizing annual NAAQS (non)attainment, although some sites' 98^{th} percentile models, notably Bakersfield, Fresno, and Riverside, perform nearly as well or even better in some cases than the annual models (Figure 2.2).

Variable Importance

We report variable importance from our sensitivity analysis for $PM_{2.5}$ and the dominant speciated component at Bakersfield, Riverside, and San Jose. These sites are highlighted as representatives of three distinct airsheds and geographic regions in California.

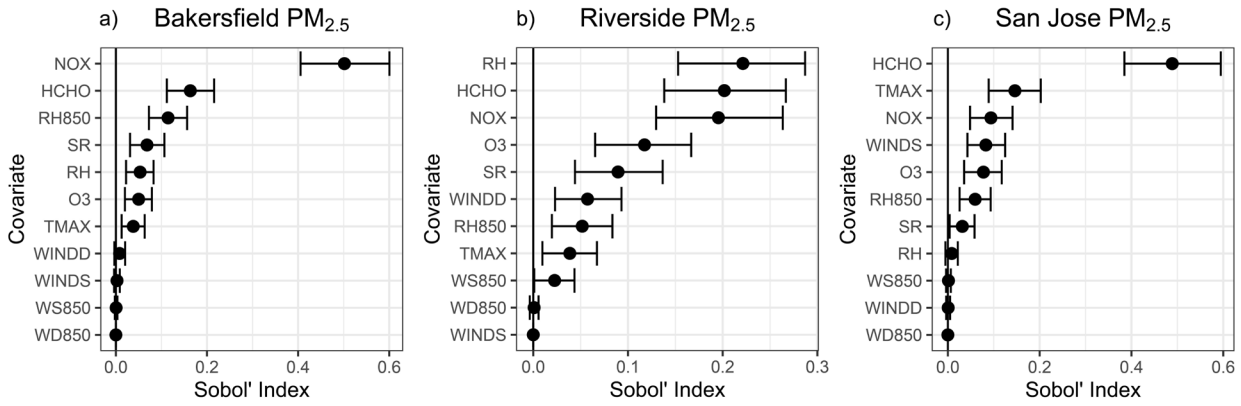

Figure 2.3: Sensitivity analysis results for a) Bakersfield $PM_{2.5}$, b) Riverside $PM_{2.5}$, and c) San Jose $PM_{2.5}$.

We see that Bakersfield $PM_{2.5}$ is influenced most strongly by gaseous precursors, suggesting a strong secondary component to Bakersfield $PM_{2.5}$. The only meteorological covariates that contribute variance significantly different from zero are solar radiation as well as upper-air and surface relative humidity, altogether contributing 24% of model variance. NO_X and HCHO, on the other hand, account for 66% of model variance, with NO_X alone contributing 50%. The remainder of the model variance is contributed by O_3 and temperature. Given the importance of NO_X, we might expect nitrate aerosol to be a prominent speciated component of Bakersfield $PM_{2.5}$. In fact, it is the dominant speciated component. Distinguishing between meteorological and gaseous precursor covariates, we find that direct meteorological

effects contribute 29% of model variance with the remaining 71% attributable to chemistry (Figure 2.3a).

In Riverside, we see that meteorology plays a more significant role than in Bakersfield, with relative humidity appearing as the most important covariate, contributing 22% of model variance. Altogether, the significant meteorological covariates, i.e., those whose 95% confidence intervals do not cross zero, contribute 48% of model variance, nearly twice as much as in Bakersfield. Gaseous precursors account for the remaining 52%. Therefore, we see a differential in meteorological importance between Bakersfield and Riverside. In the former, less than a third of model variance is attributable to meteorology, while one half is contributed by meteorology in Riverside (Figure 2.3b).

San Jose $PM_{2.5}$ is dominated by the influence of HCHO, indicating a strong association with photochemistry and potentially the presence of significant SOA. Indeed, OC, which includes SOA, constitutes the dominant speciated component of San Jose $PM_{2.5}$. Including NO_X and O_3, gaseous covariates contribute 66% of model variance, with the remaining 34% attributable to direct meteorology. Thus, San Jose falls between Bakersfield and Riverside on a spectrum from meteorology-dominance to precursor-dominance with respect to $PM_{2.5}$ (Figure 2.3c).

While meteorology influences chemistry, our additive model structure excludes covariate interactions because they were found to contribute negligible explanatory power at the modeled timescales. Therefore, the effects we see are the main effects attributable to the direct or proxy influences of precursors and meteorology.

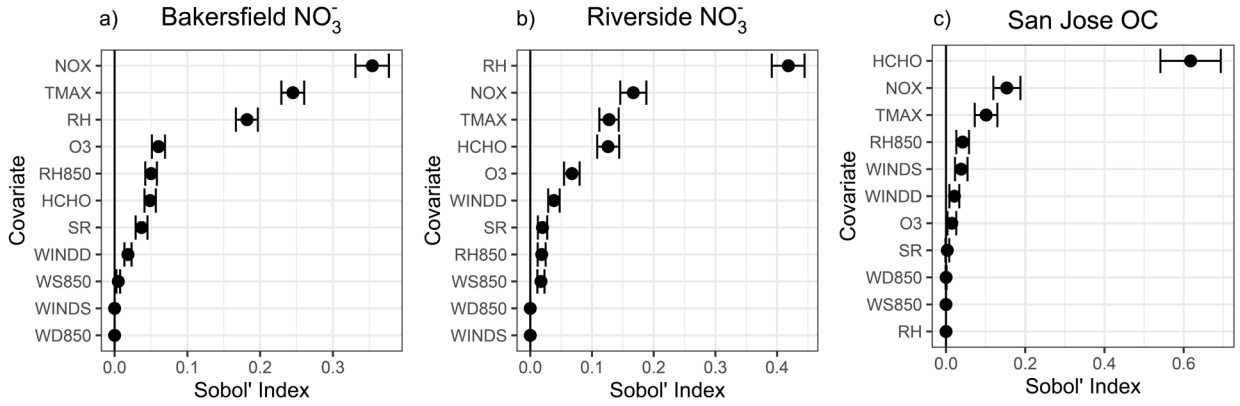

Figure 2.4: Sensitivity analysis results for a) Bakersfield NO_3^-, b) Riverside NO_3^-, and c) San Jose OC.

Gaseous precursors account for 46% of Bakersfield NO_3^- model variance, with NO_X contributing 35%. This leaves 54% of model variance attributable to meteorology, with temperature being the dominant covariate at 25% followed by relative humidity at 18% (Figure 2.4a).

The Riverside NO_3^- model's variance is dominated by relative humidity, which contributes 42% of model variance. Altogether, meteorology contributes 64% of model variance, leaving 36% attributable to gaseous precursors. The dominant gaseous precursor is NO_X at 17% (Figure 2.4b). The importance of relative humidity for both Riverside and Bakersfield NO_3^- may be at least partly explained by the ion's hygroscopicity and equilibrium partitioning of nitric acid (HNO_3), which is greater at higher relative humidity and leads to aqueous NO_3^- production (Seinfeld and Pandis 2016).

San Jose OC is dominated by HCHO, to which 62% of model variance is attributable. This is likely a result of HCHO acting as a tracer for ROG chemistry that produces secondary OC PM while OC contributes the majority of PM mass. Altogether, gaseous precursors contribute 79% of model variance, leaving 21% attributable to meteorology, with daily maximum temperature dominating meteorological contributions at 10% (Figure 2.4c).

The preceding discussion presents our sensitivity analysis results. However, it must be noted that Sobol' indices measure variable importance with respect to the model and not with respect to observations. A simple back-of-the-envelope calculation might qualify the numbers as follows. With 52% of Riverside $PM_{2.5}$ model variance attributable to gaseous precursors and considering that the model explains 54% of observed daily variability, we might conclude that the gaseous precursors we consider in our model explain $52\% \times 54\% = 28\%$ of observed variability. Similarly, 26% of observed daily variability is explained by the particular set of meteorological variables we consider, leaving 46% of observed daily variability to be explained by some set of covariates for which we have not accounted.

For example, some additional covariates that could contribute explanatory power to our models might include boundary layer height for meteorology and more VOC species for chemistry. Boundary layer height plays a role in concentrating or diluting near-surface air pollution and HCHO alone only traces some of the many possible pathways to secondary PM production involving VOCs. Unfortunately, we are unable to include these covariates in our models due to a lack of measurements over the entire period of study. Temporal resolution of input data may be yet another source of missing variability in our models, as our 24-hour averages obfuscate diurnal patterns that influence $PM_{2.5}$ behavior.

Marginal Effects

We highlight marginal effects plots for the top two most important covariates identified from our sensitivity analysis for each site's $PM_{2.5}$ and speciated component models. Marginal effects curves are plotted along with their 95% confidence bands. Rug plots are included to indicate observation density.

For both NO_X and HCHO in our Bakersfield $PM_{2.5}$ model, the marginal effects curves exhibit nonlinear monotonic positive associations. The positive association exhibited by NO_X (Figure 2.5a) is consistent with a modeling study by Chen et al. (2014) that suggests reduced NO_X levels would reduce $PM_{2.5}$ in the San Joaquin valley. Although Chen et al. (2014) indicates ROG does not strongly influence $PM_{2.5}$ levels while our HCHO response curve indicates otherwise (Figure 2.5b), a separate measurement study by Zhao et al. (2013)

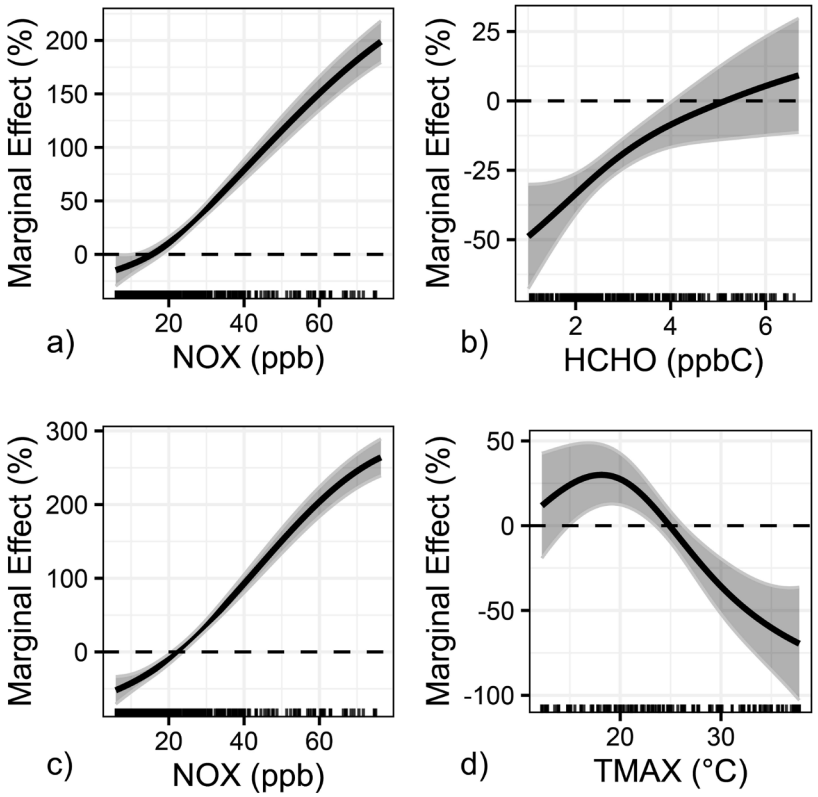

Figure 2.5: Marginal effects plots for the top two most important covariates for the Bakersfield $PM_{2.5}$ model (a and b) and the Bakersfield NO_3^- model (c and d).

suggests that reduced gaseous organic precursors would reduce SOA levels in Bakersfield (Zhao et al. 2013). This discrepancy arises because Chen et al. (2014) modeled wintertime $PM_{2.5}$ in Bakersfield, which is dominated by NO_3^-, whereas Zhao et al. (2013) measured summertime organic aerosol, which is the dominant speciated component of $PM_{2.5}$ during the summer season. Therefore, we have identified the top two most influential drivers of $PM_{2.5}$ in Bakersfield, with the most influential being the wintertime and the second most influential being the summertime driver. This ranking makes sense given that $PM_{2.5}$ levels are higher in the winter (Chen et al. 2014). This may also explain the negative marginal effects we see with HCHO. As a tracer for ROG, the influence of HCHO is strongest when the organic PM mass fraction is greatest, which occurs in the summer. Overall, $PM_{2.5}$ is greatest during the winter, however, so HCHO exhibits a negative marginal effect over most of its range because it is correlated with relatively lower summertime $PM_{2.5}$.

We observe a similar story for NO_3^- and NO_X (Figure 2.5c) as for $PM_{2.5}$ and NO_X, however, temperature exhibits non-monotonic behavior (Figure 2.5d). This suggests cooler

temperatures may be optimal for NO_3^- formation, as we see a correlation to that effect. This effect may be due to the higher likelihood of ammonium nitrate formation during cooler weather, given its potential to evaporate back to the gas phase in warmer weather (Nussbaumer and Cohen 2021).

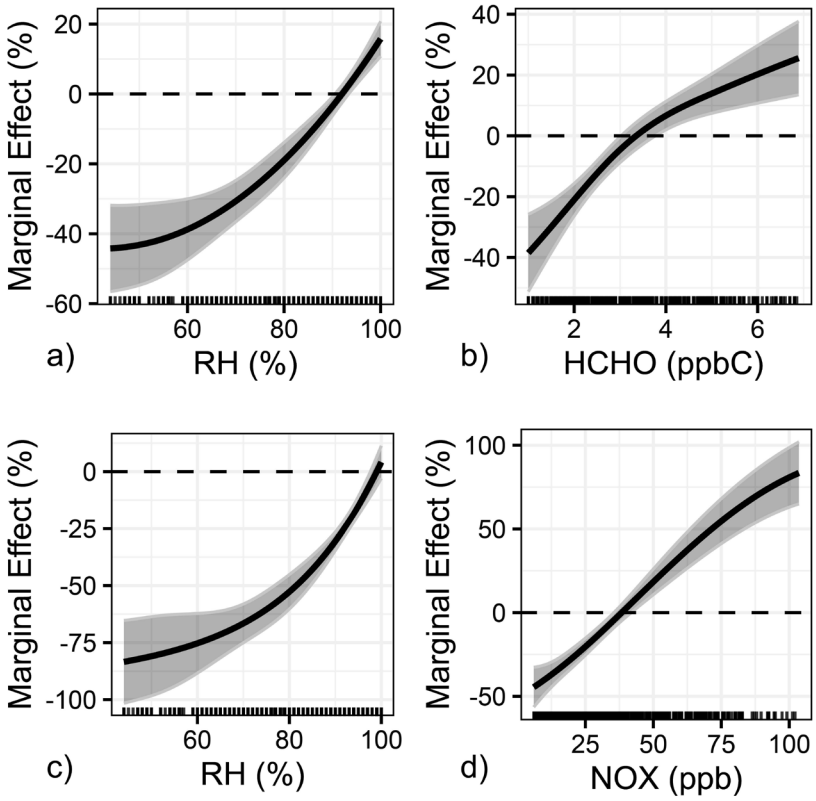

Figure 2.6: Marginal effects plots for the top two most important covariates for the Riverside $PM_{2.5}$ model (a and b) and the Riverside NO_3^- model (c and d).

In Riverside, relative humidity is correlated with reductions in total $PM_{2.5}$ (Figure 2.6a) as well as speciated NO_3^- (Figure 2.6c) over most of its range, but at sufficiently high values relative humidity is correlated with enhancements in $PM_{2.5}$ concentrations. The positive association is likely due to HNO_3 conversion to NO_3^- in aqueous aerosol, which explains the NO_3^-–NO_X relationship (Figure 2.6d) given NO_X conversion to HNO_3 (Seinfeld and Pandis 2016). Like in Bakersfield, HCHO in Riverside is positively associated with $PM_{2.5}$ (Figure 2.6b). This is an expected outcome due to the prevalence of SOA formation in the South Coast Air Basin (Woody et al. 2016). While the OC mass fraction in Bakersfield grows from winter (0.13) to summer (0.17), the Riverside OC mass fraction, to a lesser extent, decreases from winter (0.17) to summer (0.16). Therefore, in Riverside we see $PM_{2.5}$ enhancement over

a greater range of the marginal effect of HCHO than we did for Bakersfield because HCHO is not correlated with relatively less ambient $PM_{2.5}$ in Riverside.

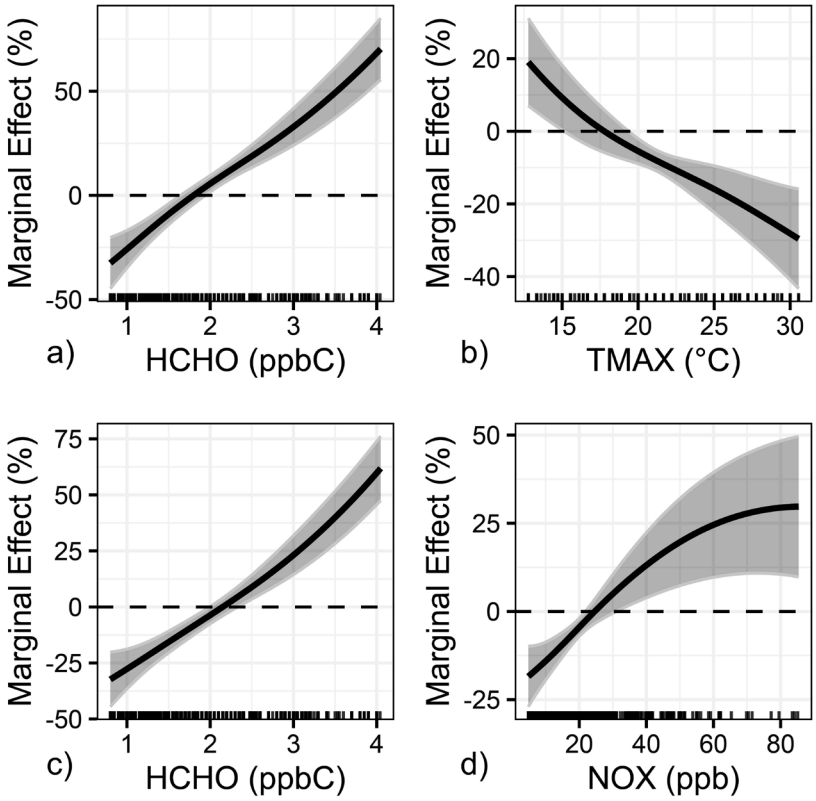

Figure 2.7: Marginal effects plots for the top two most important covariates for the San Jose $PM_{2.5}$ model (a and b) and the San Jose OC model (c and d).

San Jose HCHO is correlated with enhancements in both $PM_{2.5}$ (Figure 2.7a) and OC (Figure 2.7c) over most of the covariate's range. We see daily maximum temperature playing an important role, with cooler temperatures correlated with $PM_{2.5}$ enhancement (Figure 2.7b). This is because $PM_{2.5}$ levels are greatest during cool weather in the wintertime (*Understanding Particulate Matter: Protecting Public Health in the San Francisco Bay Area* 2012). Like HCHO, NO_X is also correlated with enhanced OC, suggesting contributions from organic NO_3^- aerosol (Figure 2.7d). Indeed, secondary $PM_{2.5}$ in the San Francisco Bay Area is dominated by NO_3^- (*Understanding Particulate Matter: Protecting Public Health in the San Francisco Bay Area* 2012). OC appears as the dominant speciated component overall, likely in the form of primary $PM_{2.5}$, because of significant wood burning emissions (*Understanding Particulate Matter: Protecting Public Health in the San Francisco Bay Area* 2012). Such

emissions also explain the importance of HCHO in its role as an ROG tracer, because ROG are important precursors for SOA and are coemitted with primary OC from wood burning.

Radiosonde-NAM Performance Comparison

To build confidence in the use of NAM upper-air data in place of radiosonde observations, we reconstruct the Los Angeles, Riverside, Sacramento, and San Jose $PM_{2.5}$ models using NAM output in place of their final versions' radiosonde observations. We calculate performance metrics for these alternate models and compare with the radiosonde versions' metrics.

Table 2.2: 24-hour performance comparison of Los Angeles, Riverside, Sacramento, and San Jose $PM_{2.5}$ models. Comparison is between models constructed using upper-air data from radiosonde observations or NAM output.

Site	Metric	Radiosonde	NAM	Difference
Los Angeles	24-hour MB (μg m^{-3})	-0.085	-0.070	0.015
	24-hour $RMSE$ (μg m^{-3})	6.2	5.9	-0.30
	24-hour R^2 (%)	47	50	3.0
Riverside	24-hour MB (μg m^{-3})	-0.065	-0.11	-0.045
	24-hour $RMSE$ (μg m^{-3})	7.8	6.9	-0.90
	24-hour R^2 (%)	54	56	2.0
Sacramento	24-hour MB (μg m^{-3})	0.13	-0.019	-0.15
	24-hour $RMSE$ (μg m^{-3})	4.2	3.8	-0.40
	24-hour R^2 (%)	26	28	2.0
San Jose	24-hour MB (μg m^{-3})	-0.11	-0.11	0.00
	24-hour $RMSE$ (μg m^{-3})	4.0	4.1	0.10
	24-hour R^2 (%)	68	67	-1.0

2.4 Conclusions

We acknowledge the following limitations in our work. Our models do not resolve seasonality, which can lead to different dominant speciated components and emissions sources between seasons. Future work should seek to construct models that account for seasonality, whether through factor interactions or separate models per season. Moreover, our models use point data from stationary observation stations, and therefore are not necessarily generalizable to other locations within their regions or air basins. Future work should test this generalizability by making predictions with the models described in this work on new data sets within the models' respective regions. We are further limited by data quality and availability through the EPA's AQS. While we have taken steps to correct OC in particular, subsequent work should consider a more careful treatment of blank corrections and may conduct new measurements for greater control over data quality and density. New measurements may

especially benefit from the use of higher resolution measurement techniques, such as aerosol mass spectrometry (AMS).

In Bakersfield, the marginal effect of HCHO is negative over a substantial portion of its range, whereas NO_X remains positive over almost its entire range, only turning negative for very low NO_X concentrations. This indicates that HCHO precursors, like alkenes and other hydrocarbons, may be more efficacious targets for control, since it would take a smaller reduction in HCHO to reduce $PM_{2.5}$ than it would take in NO_X to achieve the same result.

Compared to Bakersfield, NO_X controls may be more effective in Riverside because more of the range of the marginal effect of NO_X is negative. HCHO controls may be similarly as effective as in Bakersfield.

In San Jose, both NO_X and HCHO marginal effects are positive over most of their ranges, indicating that strong controls for both may be necessary as neither is significantly more efficacious than the other.

Chapter 3

Characterizing Organic Aerosol in California

This chapter is adapted from:

Don Collins, Roya Bahreini, and Cesunica Ivey. "Understanding the Sources and Formation Regimes of Present-Day $PM_{2.5}$ to Mitigate Particulate Pollution in California". Technical Report, *California Air Resources Board*, Award #21RD010, 2026.

The dissertation author was the primary investigator and author of the material adapted here.

3.1 Introduction

Chapter 2 provides an overview of $PM_{2.5}$ at distinct locations in California, aggregated over multiple decades. This long-term aggregation smooths out short-term variability, like seasonality and diurnal patterns. To enrich this analysis, we turn to higher resolution PM measurements made during intensive field campaigns in California. The study sites were in Riverside from February 15 - April 7, 2022 and September 23 - October 27, 2022, Wilmington from February 25 - March 28, 2023, and Bakersfield from March 30 - April 22, 2024. Measurements were made with a variety of instruments, including an aerodyne compact time-of-flight mini aerosol mass spectrometer (CToF-mAMS, hereafter referred to as AMS), which we analyze here. The AMS resolves the composition of non-refractory particulate matter in the submicron range (NR-PM_1). For these field campaigns, the measured species were particulate ammonium (NH_4^+), sulfate (SO_4^{2-}), nitrate (NO_3^-), chloride (Cl^-), and organic aerosol (OA). The AMS measures at a time resolution of 3 min, which we aggregate to hourly averages. Hourly average data is the highest time resolution available for meteorological and gaseous covariates, which are sourced from nearby regulatory monitors and airports.

At hourly resolution, these data enable modeling that captures diurnal patterns. We

train GAMs on smoothed and detrended OA data to control and account for short-term variability and measurement noise, which is important at hourly time scales. These models provide additional insights into the behavior of OA at the study sites, in finer detail than the decadal models. The models are localized temporally and spatially, meaning they are not well-suited to extrapolation or other generalization beyond the seasons and locations in which they are trained. They do perform well in their respective study sites and seasons, lending confidence to the modeled relationships. As in Chapter 2, all model relationships are contextualized in the broader literature to help explain observations and make inferences. These results can inform the design of local control strategies for OA in Riverside, Bakersfield, and Wilmington, and we provide recommendations to do so.

3.2 Materials and Methods

Training Data

The training data are processed to investigate trends and short-term fluctuations separately. Trends are computed first as rolling 24-hour means, then diurnal cycles are calculated from the trends as the average of each hour of the day over the full data set. The denoised training data are the sum of trends and diurnal cycles. Detrended data are defined as the difference between the unprocessed data and the denoised data.

Model Description

The models for each campaign and training set (i.e., the denoised and short-term fluctuations) are initialized as GAMs that predict OA trained against the covariates detailed in Table 3.1. We model OA because it is the major particulate species that exhibits the most complex behavior. The covariates are chosen as the widest pool of predictors for which we have collocated data coincident with AMS measurements.

Analysis of variance (ANOVA) is calculated for the initial model, then the covariate with the smallest F statistic (i.e., the least significant covariate) is dropped, and the model retrained. We repeat this process until dropping a covariate decreases the model R^2 by 0.01 or more in a single iteration.

After denoising or detrending, the training distributions become markedly less skewed. We therefore train GAMs assuming Gaussian distributions with identity link functions instead of the log link Gamma distributions used in Chapter 2. Model performance is summarized in Table 3.2 with R^2 defined as in Equation 2.8. The denoised data are the sum of rolling 24-hour trends and diurnal cycles while the detrended data are the difference between the unprocessed hourly data and the denoised data. Models are trained separately for the denoised and detrended data sets.

Table 3.1: Description of covariates.

Covariate	Description	Units	Data Source
NH_4^+	Particulate ammonium	μg m^3	AMS (ambient)
SO_4^{2-}	Particulate sulfate	μg m^3	AMS (ambient)
NO_3^-	Particulate nitrate	μg m^3	AMS (ambient)
Cl^-	Particulate chloride	μg m^3	AMS (ambient)
NO_X	Nitrogen oxides	ppm	ARB monitor
CO	Carbon monoxide	ppm	ARB monitor
O_3	Ozone	ppm	ARB monitor
TEMP	Temperature	°C	Airport monitor
RH	Relative humidity	%	Airport monitor
WS	Wind speed	m s^{-1}	Airport monitor
WD	Wind direction	°	Airport monitor
f_{43}	m/z 43 fraction	Unitless	AMS (ambient)
f_{44}	m/z 44 fraction	Unitless	AMS (ambient)

Marginal Effects

Because of the identity link used here, marginal effects are more directly obtained than in Chapter 2. They are simply the spline fit for each covariate, without the need for transformation. Because we have smaller sets of covariates after variable selection, we present more than the most influential two marginal effects like we showed in Chapter 2. This helps to contextualize the hourly models in more detail given the richer dynamics at this time scale.

3.3 Results and Discussion

Table 3.2: Model performance for each campaign and training set.

Campaign	Denoised R^2	Detrended R^2
Riverside (Mar)	0.96	0.80
Riverside (Oct)	0.88	0.50
Wilmington (Mar)	0.91	0.66
Bakersfield (Apr)	0.91	0.42

Riverside (Mar)

Focusing on rolling 24-hour trends and diurnal periodicity, ambient OA during this campaign is most parsimoniously characterized by NO_3^-, NO_X, CO, O_3, and SO_4^{2-} (Figure 3.1). Because CO and NO_X are strongly correlated ($r = 0.94$), they are likely related to mobile sources,

which is further supported by their diurnal profiles that follow standard traffic patterns. This explains the positive associations of the marginal effects for CO (Figure 3.2), which is co-emitted with tailpipe primary OA (POA) and secondary OA (SOA) precursors. This also helps explain the significance of O_3 and the positive associations of its marginal effects (Figure 3.2), as ozonolysis of mobile source SOA precursors like alkanes can enhance OA through OOA formation (X. Zhang et al. 2014).

SO_4^{2-}, which likely comes from upwind sources (Mysliwiec and Kleeman 2002), also displays positive associations in its marginal effects (Figure 3.2). This could be due to co-transport of upwind OA or acidic SO_4^{2-} seed particles catalyzing SOA formation (S. Gao et al. 2004; Jang et al. 2002), possibly a mix of both. The marginal effects of NO_3^- (Figure 3.2) are likely because most of the NO_3^- is organic, averaging 67%.

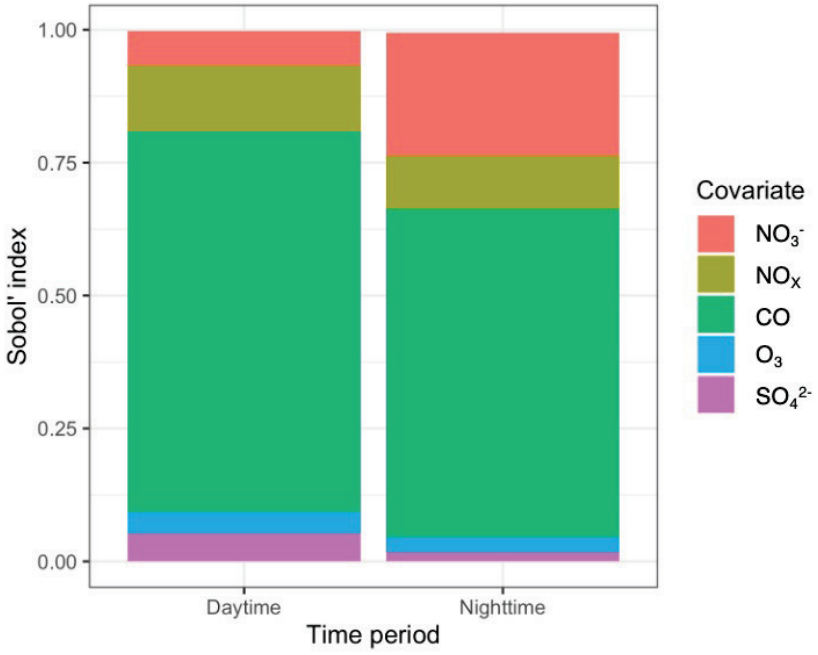

Figure 3.1: Sensitivity analysis for the Riverside (Mar) denoised model.

To investigate short-term fluctuations, we remove the rolling 24-hour and diurnal signals. Immediately, we see that these fluctuations require a wider set of covariates to balance explanatory power with parsimony. As before, SO_4^{2-}, NO_3^-, NO_X, CO, and O_3 all contribute significant explanatory power, but now the mass fractions of m/z 43 (f_{43}) and m/z 44 (f_{44}) are included (Figure 3.3). The covariates common to both models are explained by the same logic, while f_{43} and f_{44} shed additional light on the chemistry taking place. We see f_{43} display a positive association with its nighttime marginal effect and no effect during the day. This indicates that the f_{43} effect we see is not related to m/z 43 associated with hydrocarbon-like OA (HOA), as HOA is related to mobile emissions and would therefore confer daytime

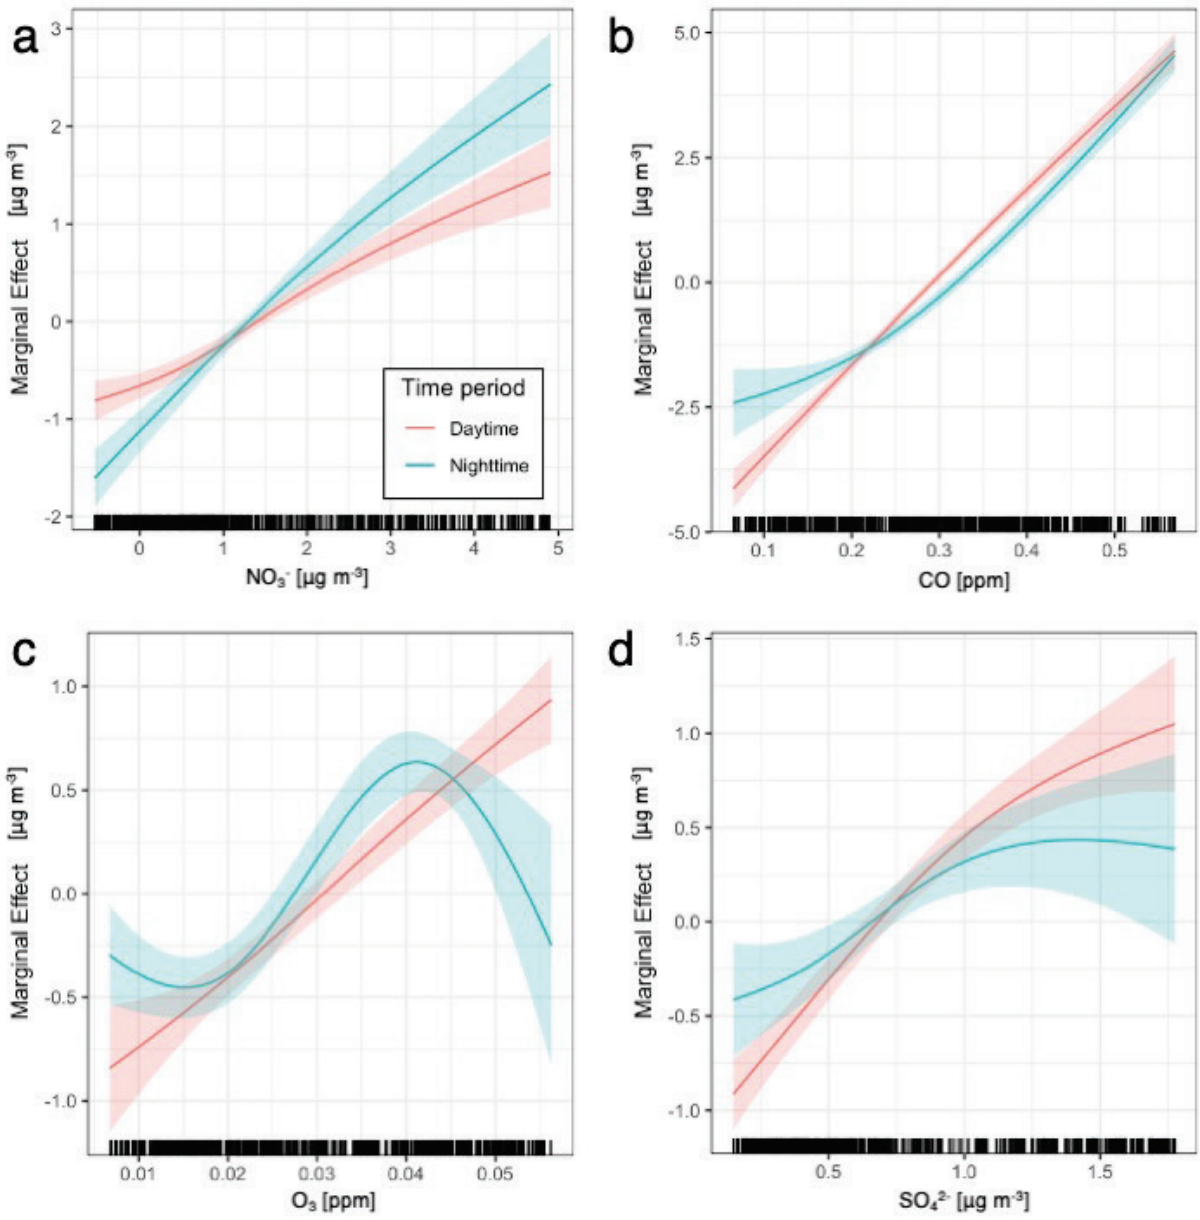

Figure 3.2: Marginal effects of (a) NO_3^-, (b) CO, (c) O_3, and (d) SO_4^{2-} for the Riverside (Mar) denoised model with 95% confidence intervals shaded.

importance, but may instead be related to less-oxidized oxygenated OA (LO-OOA), which also has a strong m/z 43 peak in its mass spectrum. This is indicative of particle aging, which also explains why f_{44} has significant explanatory power, since the m/z 44 peak is also present in LO-OOA and is especially strong in more-oxidized oxygenated OA (MO-OOA).

The negative association we see for the marginal effects of f_{44} are likely because oxygenated OA (OOA) is anticorrelated with OA, attaining its maximum in the afternoon between rush hours.

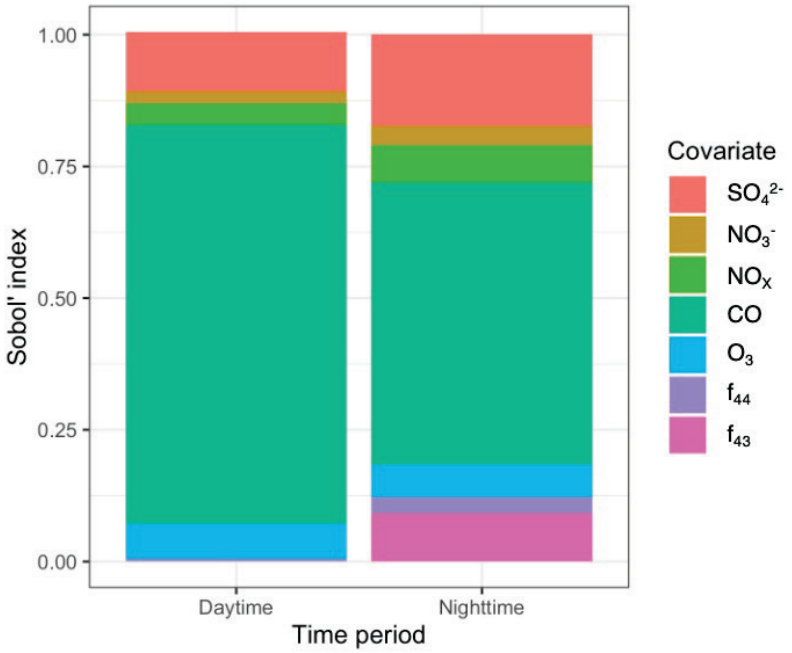

Figure 3.3: Sensitivity analysis for the Riverside (Mar) detrended model.

Riverside (Oct)

Like during Riverside (Mar), SO_4^{2-} and CO contribute explanatory power to the model trained on rolling 24-hour trends and diurnal periodicity, but unlike that campaign, meteorological covariates (wind speed, wind direction, and RH) provide additional explanatory power; we also see f_{43} selected (Figure 3.5). SO_4^{2-}, CO, and f_{43} all display positive associations in their marginal effects (3.6). SO_4^{2-} is likely due to a mixture of inflow and acid-catalyzed SOA formation (S. Gao et al. 2004; Jang et al. 2002; Mysliwiec and Kleeman 2002). Although NO_X was not a statistically significant explanatory covariate, CO is still strongly correlated with it ($r = 0.85$), indicating mobile sources for CO. HOA is also strongly influenced by mobile sources and its strong m/z 43 peak helps explain the positive association in the marginal effects for f_{43}, but biogenic SOA (BSOA) is also characterized by a strong mass spectrum peak at m/z 43 and constitutes a larger fraction of OA mass throughout the campaign. This suggests that high OA loadings are associated with high BSOA and/or HOA mass fractions.

RH displays a negative daytime association and a concave down nighttime association in its marginal effects (3.6). This is likely because RH and OA are anticorrelated, with OA

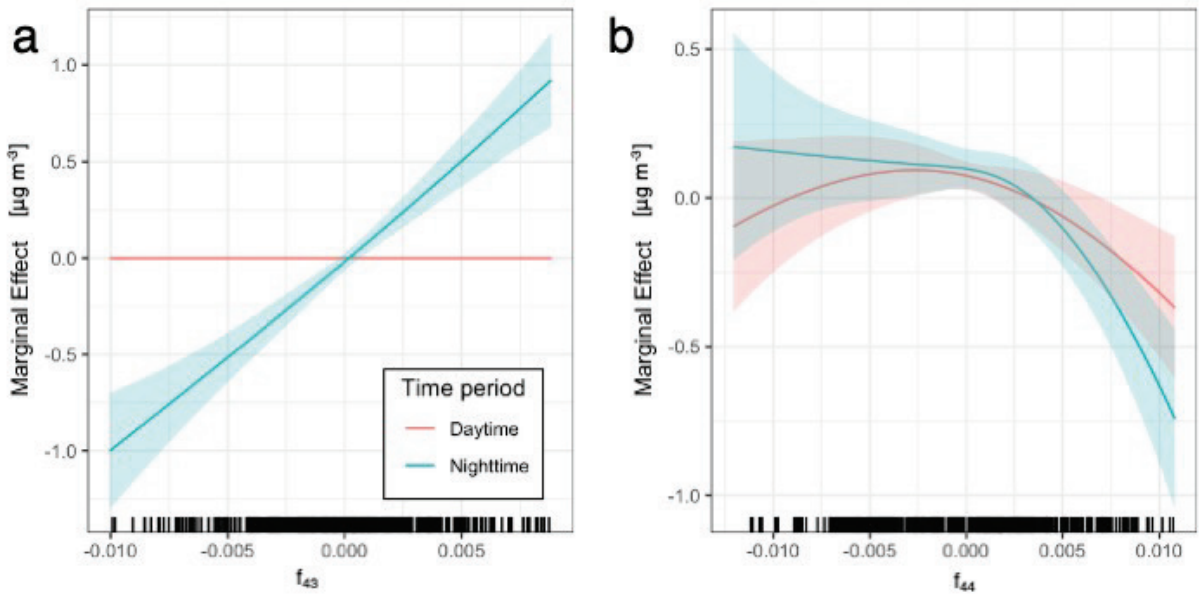

Figure 3.4: Marginal effects of (a) f_{43} and (b) f_{44} for the Riverside (Mar) detrended model with 95% confidence intervals shaded. Note: negative covariate values due to detrending.

higher during the day than overnight and the inverse for RH. Wind speed displays a positive association in its marginal effects (3.7), indicating the importance of transport as high wind speeds would display a negative association if ventilation were the only role they played. The marginal effects of wind speed indicate that westerly to east-north-easterly winds suppress OA while southerly to westerly winds enhance OA, with westerly winds, as the predominant wind direction, producing no effect relative to the mean.

The covariates at our disposal for this campaign are insufficient to adequately explain short-term fluctuations ($R^2 = 0.50$), so we focus on the preceding analysis of rolling 24-hour and diurnal trends.

Wilmington (Mar)

As in the Riverside campaigns, OA in Wilmington during this campaign is best characterized by a covariate set that includes NO_3^- and CO. Alongside O_3 and f_{43}, which were also important in the Riverside campaigns, temperature appears as a significant explanatory variable. With the exception of f_{43}, each selected covariate displays positive associations in its marginal effects (Figures 3.8 and 3.9). NO_3^-, as before, is due to the large fraction of NO_3^- that is organic (0.67), while CO is a tracer for mobile POA and SOA precursor emissions and O_3 is a tracer for photochemical activity and OA oxidation through ozonolysis. Temperature most likely displays positive associations in its marginal effects because it is a proxy

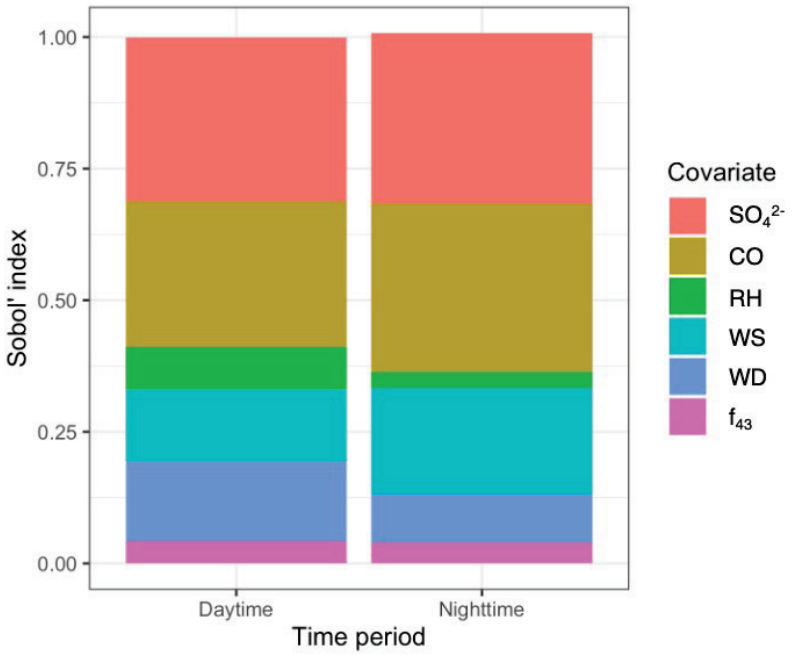

Figure 3.5: Sensitivity analysis for the Riverside (Oct) denoised model.

for insolation, which drives photochemical activity that enhances ozonolysis and thus SOA. In contrast, the marginal effects for f_{43} display negative associations. This is most likely because f_{43} anticorrelates with OA, since the mass fractions for HOA and OOA, which have substantial m/z 43 peaks, grow when total OA decreases.

As in Riverside in March, short-term fluctuations are not as well explained by our data, but most of the variance ($R^2 = 0.66$) is accounted for, this time by SO_4^{2-}, NO_3^-, Cl^-, CO, O_3, f_{44}, and wind direction (Figure 3.11). The relative importances of NO_3^- and CO flip compared to the denoised model, with NO_3^- now more important than CO, suggesting secondary NO_3^- processes are more important than mobile source POA and SOA precursors for short-term fluctuations. The inclusion of Cl^- is indicative of fresh sea salt influence, which together with its marginal effects (Figure 3.12) suggests Cl^- is acting to enhance OA. This can be explained by the mechanisms described by Cai and Griffin (2003) or Laskin et al. (2012), which describe how organic coatings on sea salt particles can enhance VOC uptake and therefore OA mass. The importance of wind direction corroborates this, as onshore versus offshore winds would modulate the influence of Cl^-. Since SO_4^{2-} in the Los Angeles basin is largely marine SO_4^{2-} (Mysliwiec and Kleeman (2002)), wind direction modulates the influence of SO_4^{2-} in the same way. SO_4^{2-} itself, in this case, likely exerts its influence through acid-catalyzed SOA formation rather than co-transport as is the case in Riverside, since the main species that would be co-transported into Wilmington with SO_4^{2-} is Cl^- from sea salt aerosol rather than upwind OA, although there could feasibly be some amount of marine OA co-

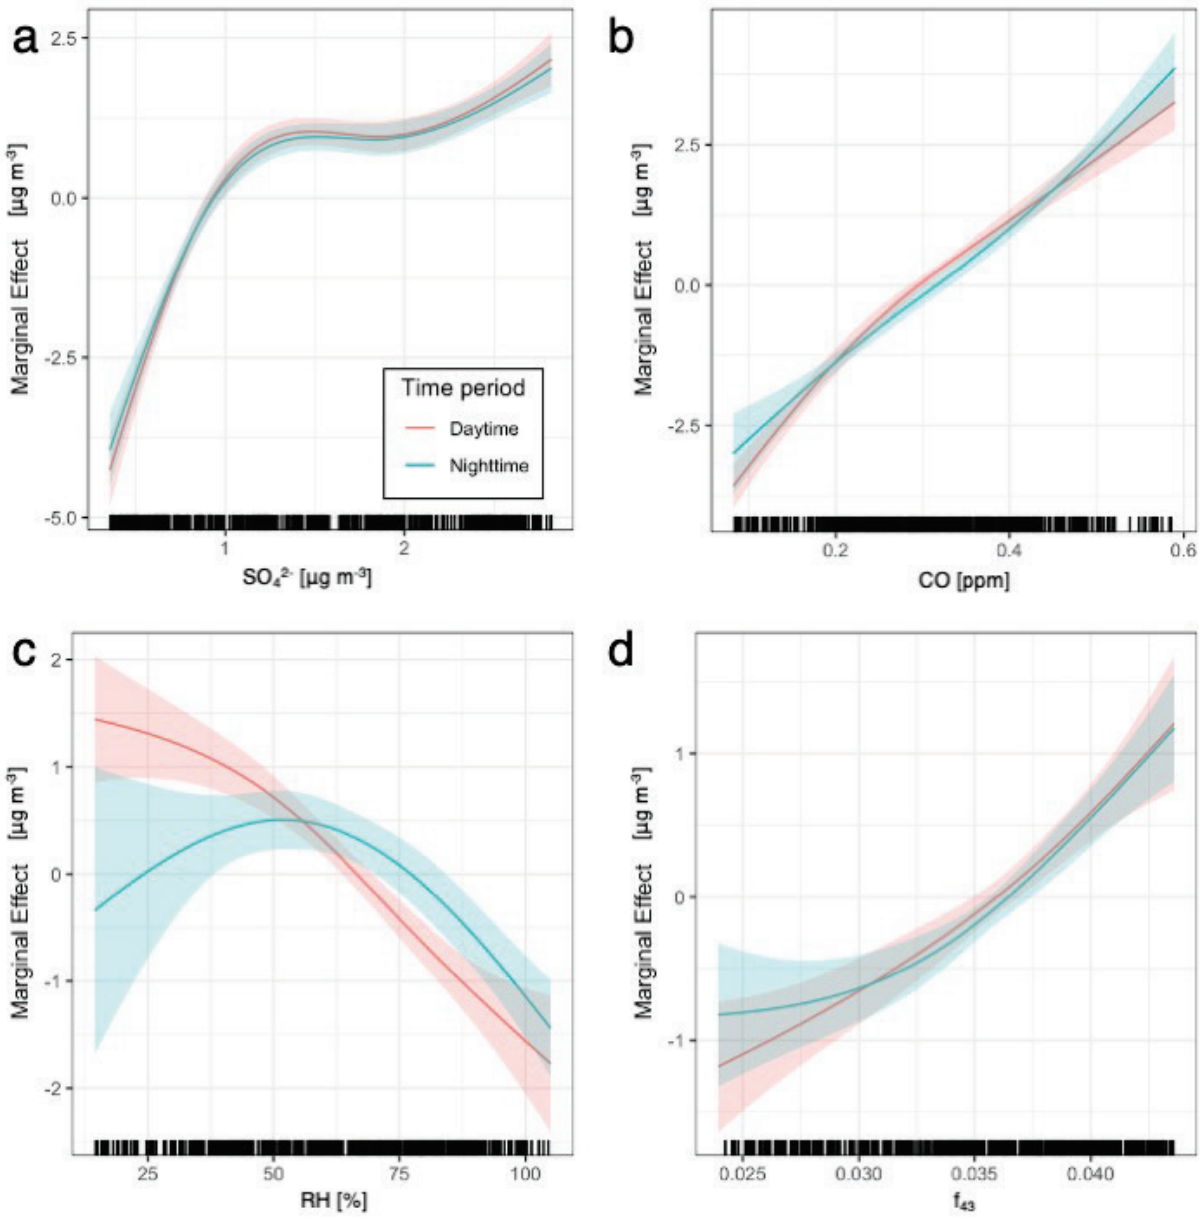

Figure 3.6: Marginal effects of (a) SO_4^{2-}, (b) CO, (c) RH, and (d) f_{43} for the Riverside (Oct) denoised model with 95% confidence intervals shaded.

transported with SO_4^{2-}. O_3 only displays a significant marginal effect during the daytime and exhibits a plateau before a positive association forms at higher O_3 concentrations. Because f_{44} displays negative associations in its marginal effects, the positive association of O_3 is less likely related to SOA formation through ozonolysis since this would form more OOA with

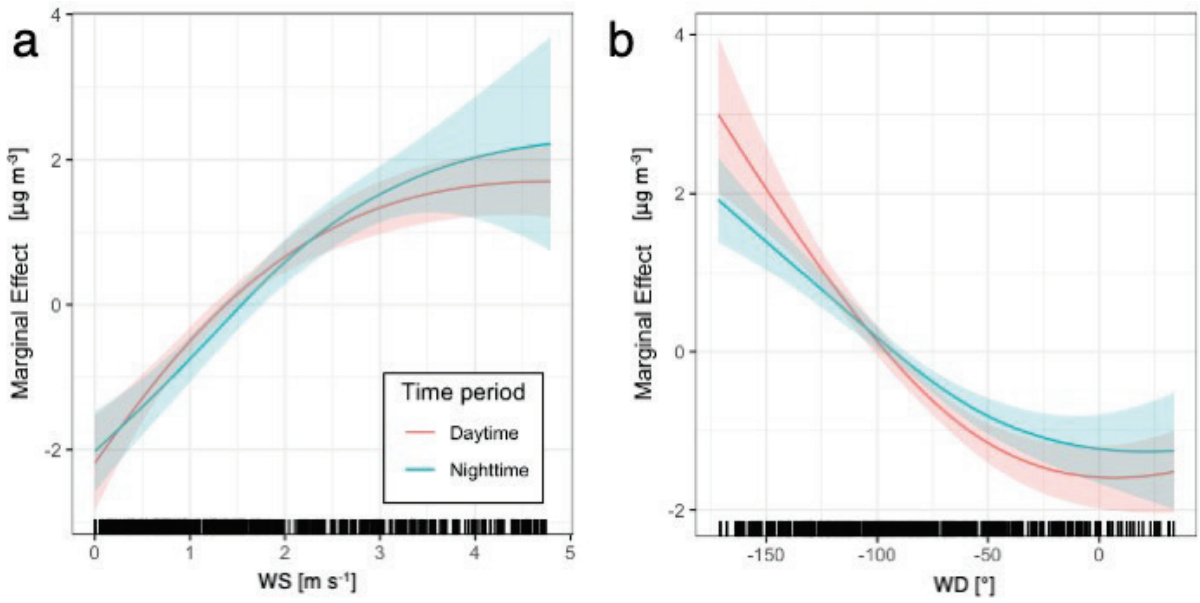

Figure 3.7: Marginal effects of (a) wind speed and (b) wind direction for the Riverside (Oct) denoised model with 95% confidence intervals shaded.

m/z 44 peaks and thus produce positive associations for f_{44} marginal effects. Instead, given the simultaneous importance of Cl^- and NO_3^-, the O_3 marginal effect may be related to gas-phase chlorine-organic reactions. Cl^- displacement by NO_3^- generates Cl gas that can react with organic gases to generate alkyl and alkylperoxy radicals that form O_3 in the presence of NO_X (Finlayson-Pitts 2003). In this scenario, as Cl^- and NO_3^- form Cl gas, concentrations of organic radicals increase. These radicals produce NO_2that reacts with VOCs to form O_3 and can continue reacting to form SOA. In this way, O_3 becomes positively associated with OA.

Bakersfield (Apr)

In this campaign, as in the others, Figure 3.13 shows that NO_3^-, SO_4^{2-}, and CO are selected for their explanatory power, but although this is not the first campaign to include temperature as a significant covariate, OA sensitivity to temperature is greater than in Wilmington (Mar). In contrast to Riverside and Wilmington, where roughly two-thirds of NO_3^- was organic during Riverside (Mar) through Wilmington (Mar), only 49% of NO_3^- is organic in Bakersfield during this campaign. The positive associations in the marginal effects of both NO_3^- and SO_4^{2-} (Figure 3.14) may be related. Rollins et al. (2013) find that the nitrate functional group alone accounted for 4.8% of OA mass during the 2010 CalNex campaign and that isoprene and monoterpene oxidation in the presence of acidic SO_4^{2-} seed formed nitrooxy organosulfates.

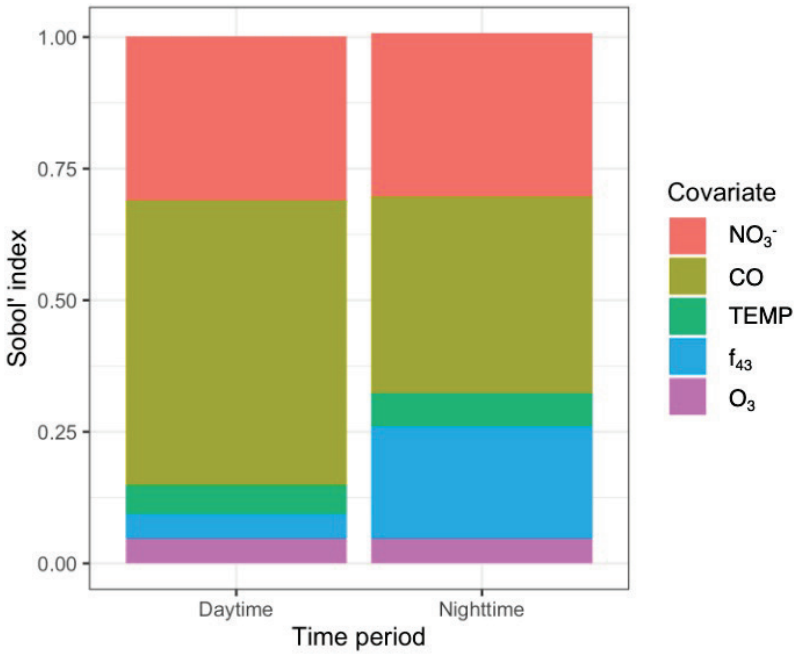

Figure 3.8: Sensitivity analysis for the Wilmington (Mar) denoised model.

The correlation between NO_3^- and SO_4^{2-} ($r = 0.60$) supports this, and could conceivably be higher between specifically organic NO_3^- and organic SO_4^{2-} but organic and inorganic SO_4^{2-} is indistinguishable by AMS (Farmer et al. 2010).

The diurnal profile of CO shows strong rush hour peaks, as does NO_X, suggesting CO is a tracer for mobile source POA and SOA precursors, which explains the positive associations in its marginal effects (Figure 3.14). The importance of Cl^- is surprising since it accounts for just 1% of total AMS mass on average during the campaign. It is not immediately clear why Cl^- exhibits explanatory power for OA during this campaign. One possible explanation is related to biomass burning, but biomass burning aerosol is not a major component of OA during this campaign. Mechanistically, wildland and agricultural fires can release potassium chloride (KCl), then NO_3^- and SO_4^{2-} can deplete Cl^- by displacing it from K^+ and forcing it into gas-phase chlorine species (Schlosser et al. 2017). This would explain the negative associations in the marginal effects of Cl^- (Figure 3.14), since biomass burning would enhance OA as NO_3^- and SO_4^{2-} deplete Cl^-.

Temperature displays positive associations in its marginal effects (Figure 3.14). Physically, higher temperatures generally reduce OA mass by enhancing VOC evaporation out of the particle phase. This discrepancy suggests an indirect relationship with temperature that overrides the importance of evaporation. This may be explained by the temperature dependence of emissions of BVOCs like isoprene (Tingey et al. 1979). Since BVOC emissions increase with temperature and BSOA makes up a large fraction of OA during this campaign,

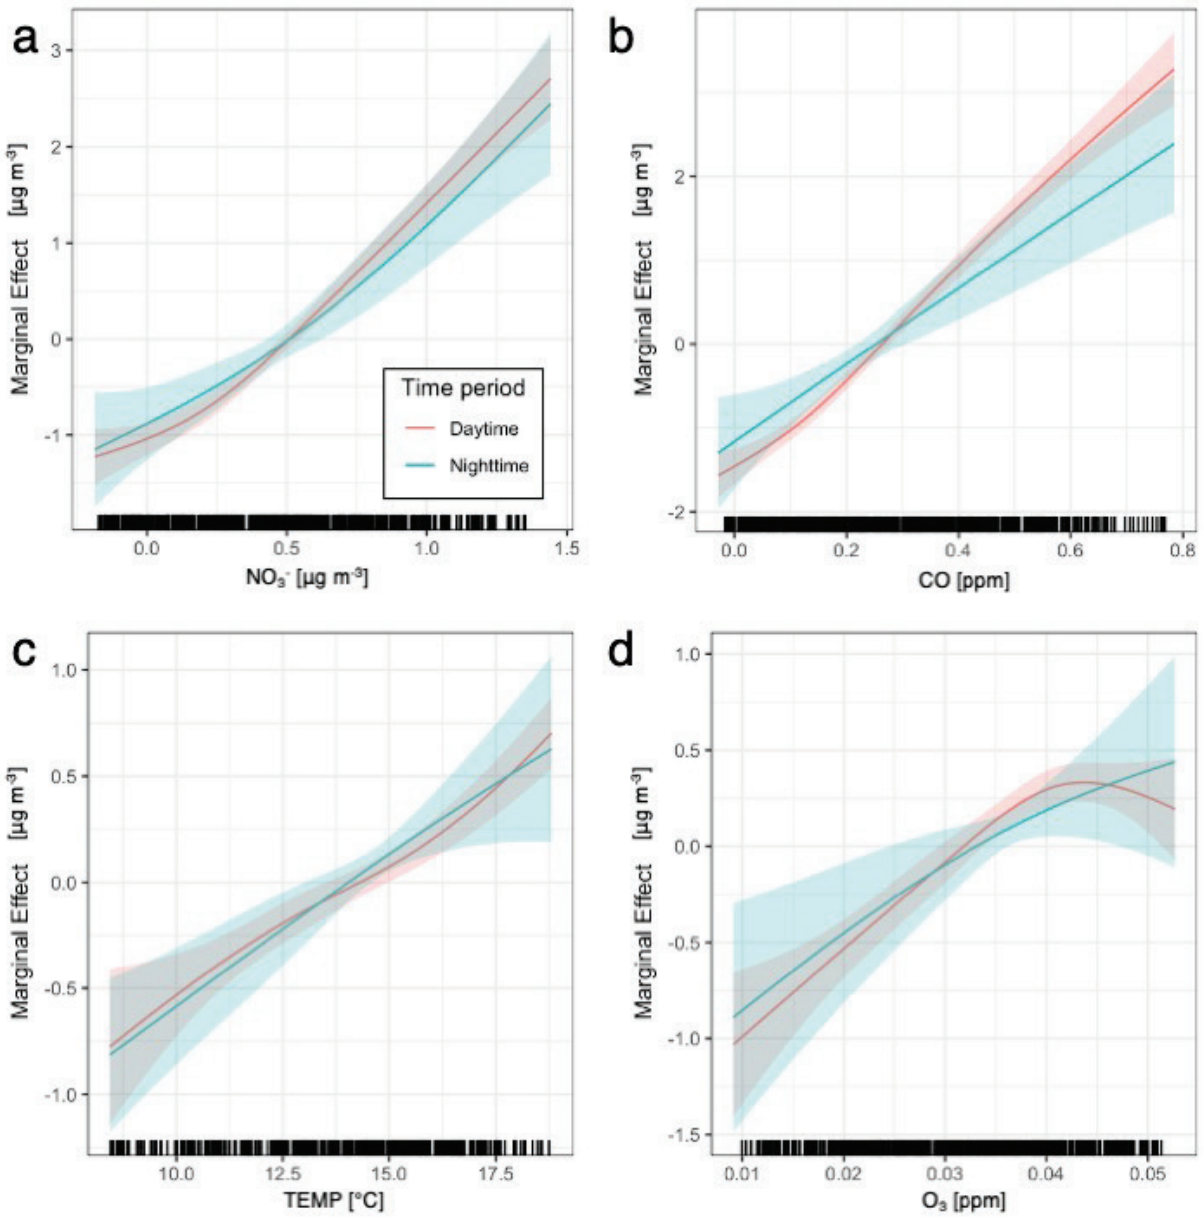

Figure 3.9: Marginal effects of (a) NO_3^-, (b) CO, (c) temperature, and (d) O_3 for the Wilmington (Mar) denoised model with 95% confidence intervals shaded.

the positive associations in the marginal effects of temperature are likely related to BSOA formation. Wind direction, ranging from $-100°$ to $+100°$ from north in our data, suggests that westerly winds (near $-90°$) reduce OA while easterly winds (near $+90°$) enhance OA. This is surprising since the actual measurement site is located east of Bakersfield, leading

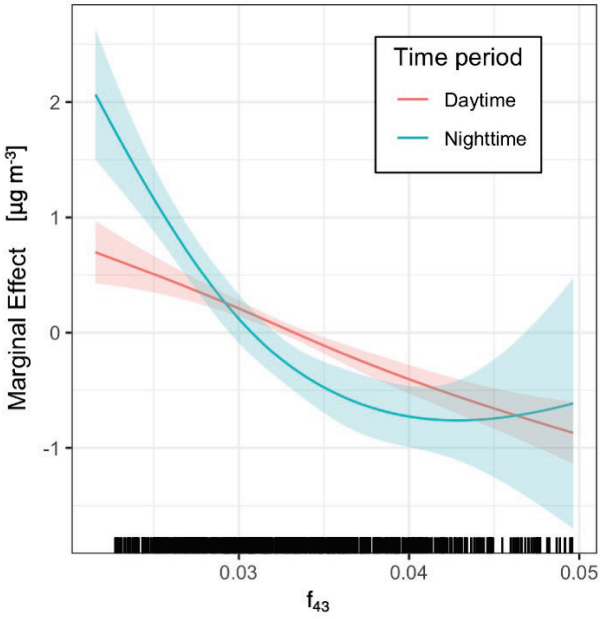

Figure 3.10: Marginal effects of f_{43} for the Wilmington (Mar) denoised model with 95% confidence intervals shaded.

to the expectation that westerly winds would bring polluted air from the city. One possible explanation is that the Bena Landfill lies due east of the measurement site, which may also help explain the role of chloride since landfills are significant chlorine emitters as well as VOC emitters (Bannan et al. 2019).

As in Riverside in October, our data for this campaign do not adequately explain short-term fluctuations ($R^2 = 0.42$), so we focus on the preceding analysis of rolling 24-hour and diurnal trends.

3.4 Conclusions

Generalized additive modeling combined with sensitivity analysis elucidates the relationships between OA covariates and their relative importances for driving OA levels. Marginal effects derived from GAMs enable predicting the individual impacts of changes in individual covariates, while sensitivity analysis allows for the ranking of covariates by the strength of their influence on OA variability. Together, these results help us understand the formation regimes of OA in California, which can help inform how to target controllable covariates for efficient OA reductions.

By contextualizing marginal effects within the literature, this analysis helps attribute marginal effects to physical mechanisms or infer likely correlations where marginal effects do

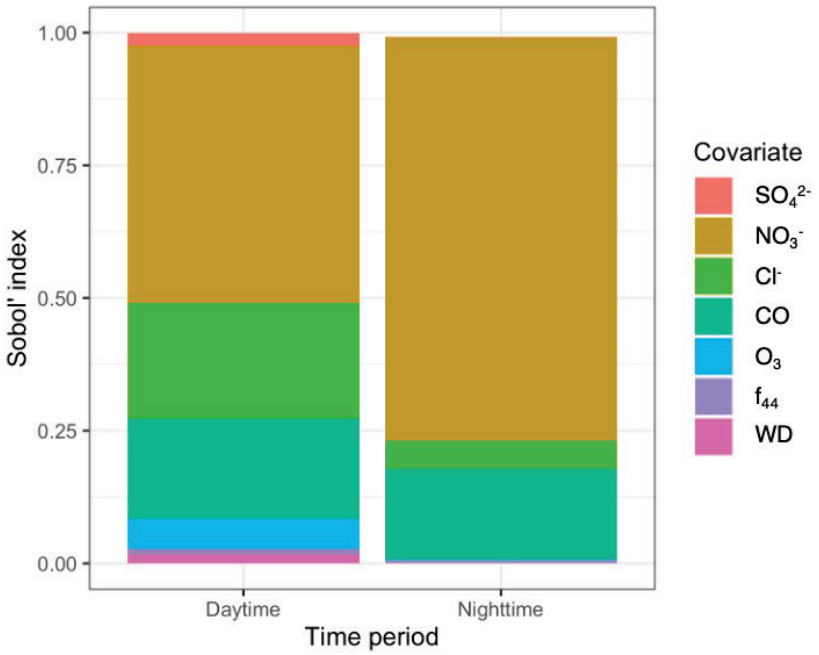

Figure 3.11: Sensitivity analysis for the Wilmington (Mar) detrended model.

not admit physical explanations. This furthers our understanding of OA regimes by revealing where interactions between covariates confound OA behavior, highlighting the nontrivial nature of OA.

In Riverside in March, OA is most strongly associated with covariates related to mobile emissions, like CO and secondary NO_3^-. In October, Riverside sees a stronger influence from upwind sources relative to March, but mobile emissions still contribute substantial explanatory power for OA behavior. This suggests the importance of reducing mobile emissions to address POA contributions to overall OA in Riverside, while the influence of OOA implies that controlling SOA formation may also alleviate overall OA levels. Our results show that OA in Riverside is associated with SO_4^{2-}, suggesting the role of acid-catalyzed SOA formation. This implies reducing anthropogenic sulfur emissions upwind of Riverside may support OA reductions in Riverside, although not all upwind sulfur is anthropogenic. Conclusions are similar for Wilmington in March, with the additional point that O_3 shows an association with OA such that O_3 controls may have the co-benefit of alleviating OA levels there. For OA in Bakersfield, there remains a strong association with covariates related to mobile emissions, but there is also a stronger biogenic influence. Additionally, the unusual influence of Cl^- in conjunction with the influence of wind direction may suggest the importance of upwind landfill emissions. Together, these observations suggest that mobile emissions may alleviate OA in Bakersfield like in Riverside and Wilmington, but also that additional site-specific control strategies that regulate specific point sources may support OA reductions. These results

imply the importance of targeted measurement and monitoring to reveal and disentangle the unique local conditions that vary across California. OA responds with varying strength to the same covariates at different locations, with some locations sensitive to covariates that exert little influential elsewhere, like SO_4^{2-} in Riverside and Bakersfield versus Wilmington, or Cl^- in Bakersfield versus Riverside and Wilmington.

With hourly time resolution, generalized additive modeling resolves features that daily averages, as in Chapter 2, could not. Temporal disaggregation into separate night and day model fits sheds light on how OA can behave differently between nighttime and daytime. Denoising to separate trends and diurnal cycles from short-term fluctuations facilitates robust model fits that characterize general OA behavior well, while short-term fluctuations are in some cases reasonably well characterized and in others not well characterized. This reflects the importance of rich, high-resolution data sets for capturing the complex dynamics of OA.

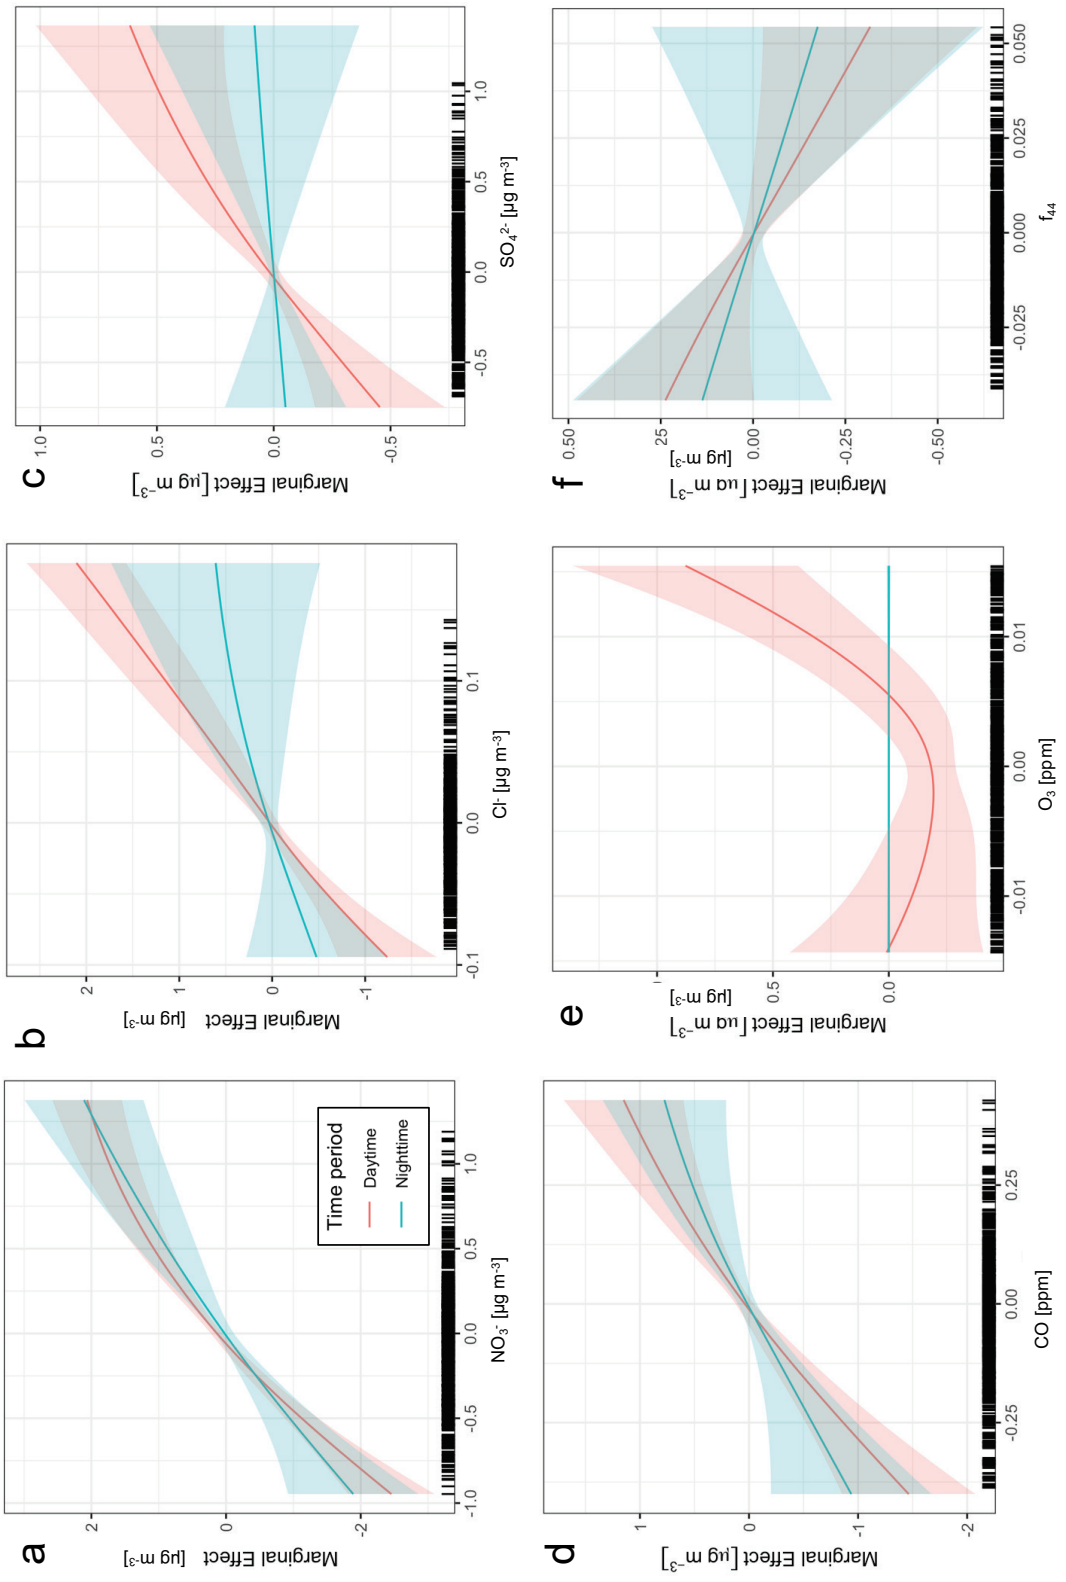

Figure 3.12: Marginal effects of (a) NO_3^-, (b) Cl^-, (c) SO_4^{2-}, (d) CO, (e) O_3, and (f) f_{44} for the Wilmington (Mar) detrended model with 95% confidence intervals shaded. Note: negative covariate values due to detrending.

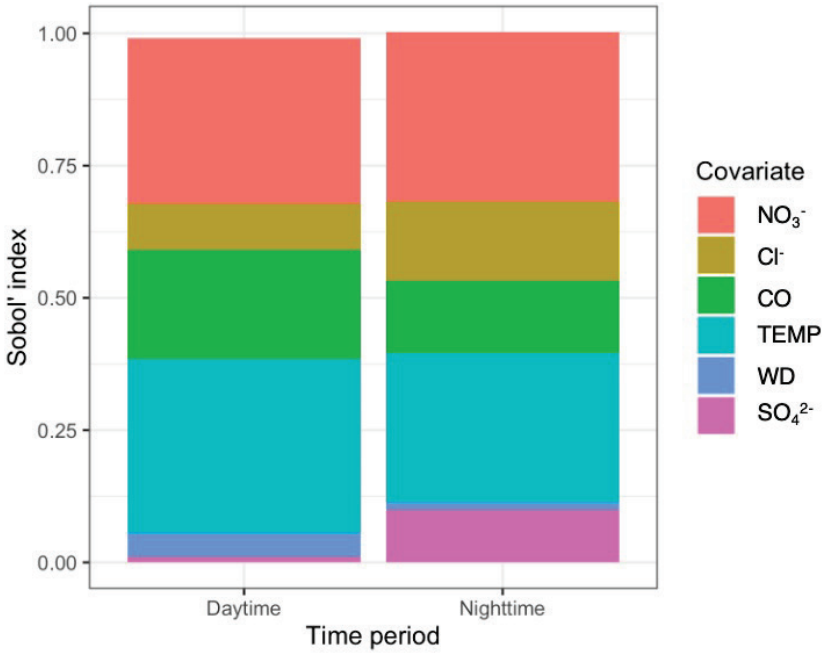

Figure 3.13: Sensitivity analysis for the Bakersfield (Apr) denoised model.

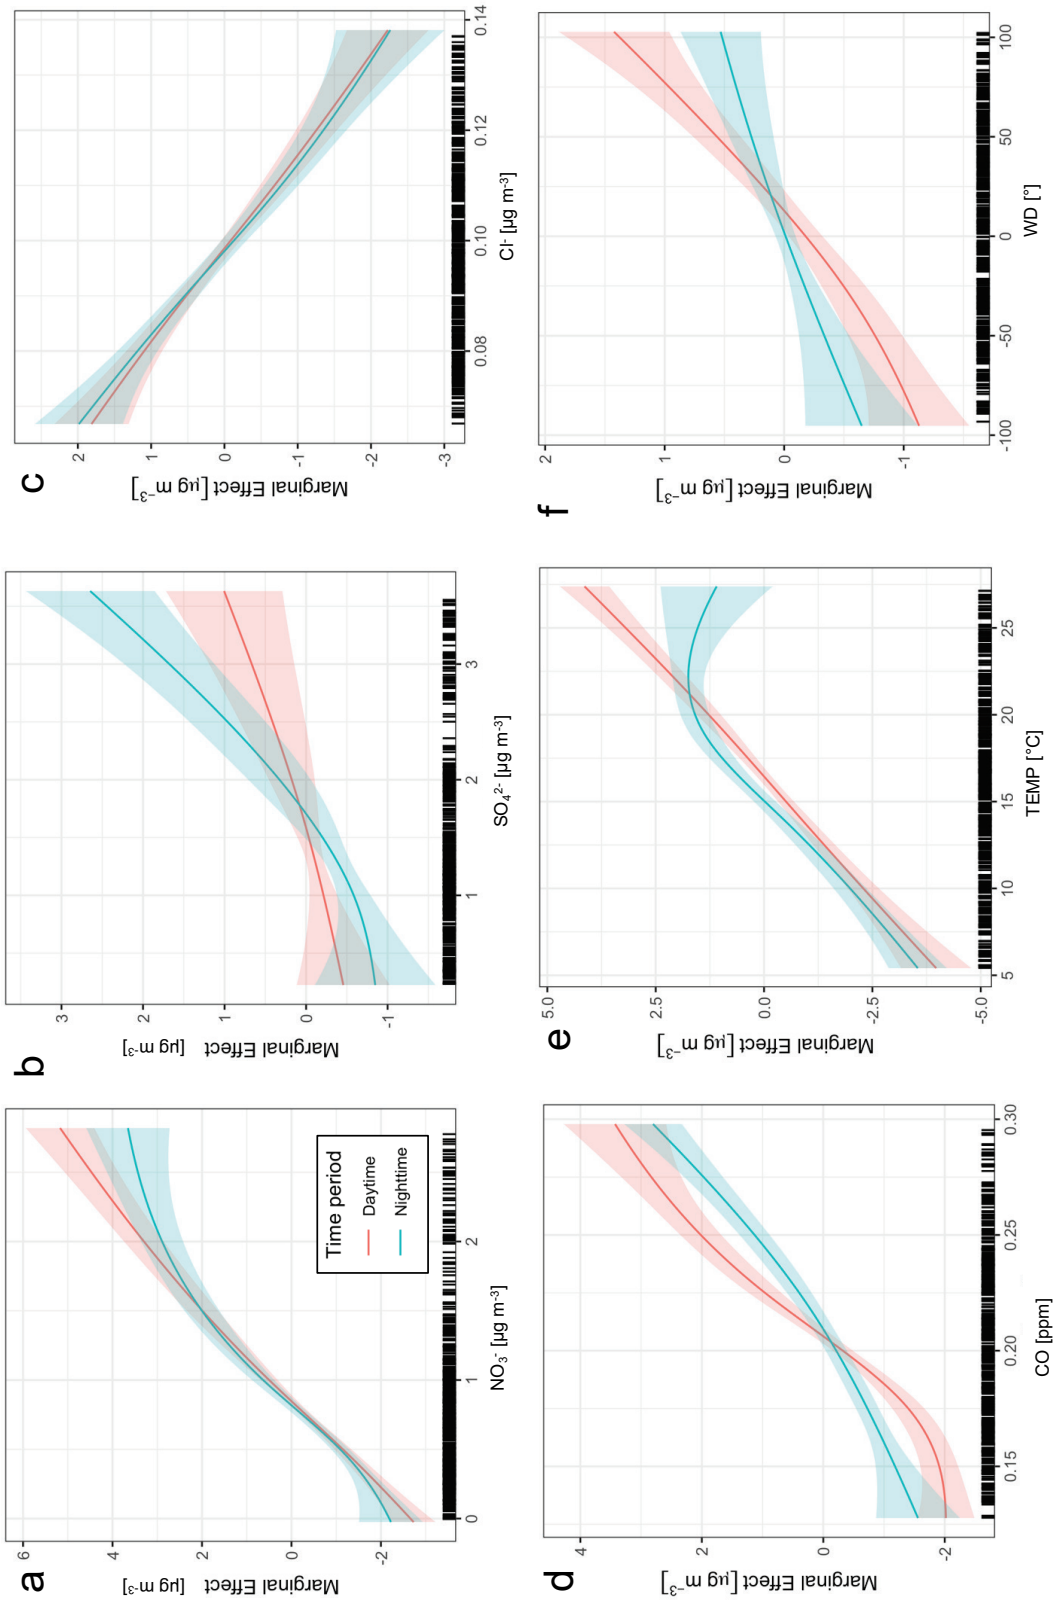

Figure 3.14: Marginal effects of (a) NO_3^-, (b) SO_4^{2-}, (c) Cl^-, (d) CO, (e) temperature, and (f) wind direction for the Bakersfield (Apr) denoised model with 95% confidence intervals shaded.

Chapter 4

Accelerating a CMAQ Chemistry Solver

This chapter is adapted from:

Duncan Quevedo, Khanh Do, et al. (Feb. 2025). "GPU Implementation of a Gas-Phase Chemistry Solver in the CMAQ Chemical Transport Model". In: *ACS ES&T Air* 2.2, pp. 226–235. DOI: 10.1021/acsestair.4c00181

The dissertation author was the primary investigator and author of this paper.

4.1 Introduction

Where Chapters 2 and 3 leveraged measurement data to analyze particulate matter in California and make direct location-based policy recommendations, this chapter turns to model development with the aim of laying a foundation for higher resolution regulatory modeling that informs local scale policymaking. Although policy impacts are more indirect than in previous chapters, this chapter's results and their implications push air quality modeling towards finer scale simulations that can support regulatory policy informed by local air pollution dynamics. The class of air quality model considered here is the chemical transport model.

Chemical transport models (CTMs) simulate numerous varied and complex atmospheric processes. These include advection, diffusion, cloud processes, aerosol processes, gas-phase chemistry, and deposition. These processes are modulated by input data including meteorology and emissions. The governing equation describing the simulation of these processes is the species continuity equation, a partial differential equation (PDE) of the form

$$\frac{\partial C_i}{\partial t} + \nabla \cdot (\mathbf{u} C_i) = (\nabla \cdot \mathbf{K} \nabla) C_i + R_i + E_i - S_i \tag{4.1}$$

where C_i is the concentration of species i, t is time, $\mathbf{u}$ is the wind velocity vector, $\mathbf{K}$ is the eddy diffusivity tensor, R_i and E_i are the net chemical rate of change and emissions of species i, respectively, and S_i is the loss of species i due to deposition (Seinfeld and Pandis 2016; Jacobson 2005).

CTMs tend to employ an operator-splitting strategy in their solution schemes (Seinfeld and Pandis 2016; Jacobson 2005). Operator splitting assumes that each operator in Equation 1 can be solved independently with little loss of solution accuracy on sufficiently small timescales. For instance, Equation 1 may be split into a family of simpler equations:

$$\frac{\partial C_i}{\partial t} = \begin{cases} -\nabla \cdot (\mathbf{u} C_i) \\ (\nabla \cdot \mathbf{K} \nabla) C_i \\ R_i \\ E_i \\ -S_i \end{cases} \tag{4.2}$$

Each case of Equation 4.2 is solved separately by its own code module within a CTM's software package (Byun and Schere 2006). Additional processes, including aerosol and cloud processes, are also treated in their own modules. Therefore, each module can be timed separately to determine how long it takes to execute its solution. The chemistry operator is determined by the chemical mechanism chosen to represent atmospheric gas-phase chemistry and takes the form of a net reaction rate. This reaction rate is determined by the solution of a system of coupled ordinary differential equations (ODEs) that needs to be solved for every grid cell in a domain.

Although species lumping is typical to reduce the number of species represented in a chemical mechanism, any mechanism needs to include many species to be representative of atmospheric gas-phase chemistry (W. P. Carter 2010). This can result in very large systems of coupled ODEs. The Statewide Air Pollution Research Center (SAPRC, now APRC) mechanism SAPRC07, for instance, contains 569 gas-phase reactions (W. Carter 2023). Moreover, these reactions are nonlinear and occur over a variety of timescales, producing stiff systems. These characteristics result in longer runtimes for the code modules that solve these systems. In fact, the gas-phase chemistry module often represents a major computational bottleneck in a CTM (Sandu et al. 1997; Kelp, Jacob, Kutz, et al. 2020). Indeed, gas-phase chemistry is often solved using schemes like Euler backwards iterative (EBI) solvers for their efficiency, at the cost of solution accuracy. More accurate methods, while slower, can improve the quantification of quantities of interest like sensitivity parameters. Additionally, there is growing interest in leveraging machine learning emulators to bypass traditional numerical integration of these systems (Kelp, Jacob, Kutz, et al. 2020; Kelp, Jacob, Lin, et al. 2022). However, there is room for improvement in these emulators' solution approximations in terms of numerical stability and accuracy. Furthermore, emulation has not been standardized for regulatory decision making and emulation does not necessarily capture the physical dynamics that govern first principles models like CTMs.

One may view emulators as a software-based approach to accelerating gas-phase chemistry modules. Another approach is hardware-based. CTM simulations are traditionally executed on high-performance computing (HPC) clusters with computational processing distributed in parallel across multiple central processing units (CPUs), which are the standard hardware units for computing. An alternative class of hardware capable of carrying out the floating-point arithmetic necessary for the numerical solution of gas-phase chemistry systems is the graphics processing unit (GPU).

With the advent of general purpose GPU (GPGPU) computing, we can leverage the massively parallel nature of GPUs to compute floating-point operations on large data arrays in parallel. Although GPU technology has been driven by advancements in real-time 3D computer graphics as required for interactivity in video games, the nature of graphics rendering—devoting individual GPU threads to rendering individual pixels—drives GPU thread counts higher and higher as pixel counts and graphics complexity increases (Nickolls and Dally 2010). This leads to the ability to compute floating-point operations on large data structures in parallel. Consequently, GPUs present a means of handling larger matrices to reduce the length of serialized loops required to cover study domains.

GPUs outperform CPUs in floating-point operation throughput (*CUDA C++ Programming Guide* 2024). Therefore, GPUs represent a potential hardware solution for accelerating gas-phase chemistry modules in CTMs. Moreover, since GPU acceleration is hardware-based, the original CTM algorithms can be retained, albeit with certain modifications to facilitate parallelism. This circumvents the numerical stability and accuracy issues that can arise in emulation. Already, stable GPU implementations of CTM algorithms have been achieved: the horizontal advection scheme in CAMx and the gas-phase chemistry schemes in CAM4-Chem and EMAC (Cao et al. 2023; Sun et al. 2018; Alvanos and Christoudias 2017).

The hardware acceleration of the horizontal advection module in CAMx achieves impressive performance, with the GPU implementation executing the module up to 28 times as fast as the standard CPU implementation (Cao et al. 2023). However, the authors report that this module only accounts for around 14% of total simulation time, with chemistry accounting for over 55% of walltime, and do not report the overall acceleration when accounting for all modules. Therefore, we can expect that accelerating chemistry would have a greater impact on total simulation walltime.

Accelerating the chemistry module in CAM4-Chem enabled simulations to execute about 4 times as fast as on standard hardware (Sun et al. 2018). However, the simulations in the study are limited to box models with only chemistry active, as opposed to a full CTM simulation.

Hardware accelerated chemistry in a full EMAC simulation was shown to speed application execution by a factor of 1.75 (Alvanos and Christoudias 2017). Accelerated model output for select species mostly stays within 5% of the value of the standard model where discrepancies result from architectural differences in how hardware handles floating point operations. This illustrates the potential performance gains achievable by accelerating gas-phase chemistry modules in CTMs.

In this work, we migrate a gas-phase chemistry solver in the United States Environmental

Protection Agency (US EPA)'s Community Multiscale Air Quality (CMAQ) CTM to GPU hardware to accelerate the solution of the reaction system and relieve the computational burden it imposes. We refer throughout this article to the migrated heterogeneous CPU-GPU code as "CMAQ-CUDA" in reference to its use of the Compute Unified Device Architecture (CUDA) general purpose GPU (GPGPU) computing solution. We say heterogeneous to refer to the fact that science processes other than the gas-phase chemistry module, CHEM, are executed on CPUs whereas CHEM is executed on GPUs.

4.2 Materials and Methods

CMAQ is widely used for regulatory compliance studies at the regional scale (Kim, Fu, and Miller 2010; Simon, Baker, and Phillips 2012). CMAQ ingests meteorology fields and emissions inventories as input and uses these data to inform science processes. CMAQ science processes include horizontal advection, horizontal diffusion, vertical advection, vertical diffusion, gas-phase chemistry, aerosol processes, cloud processes, and dry deposition (Byun and Schere 2006). The advection and diffusion modules solve advection or diffusion equations, respectively. Gas-phase chemistry solves the ODE systems that represent chemical mechanisms. Aerosol processes represent particulate physics and chemistry. Cloud processes handle scavenging, wet deposition, and aqueous chemistry. Dry deposition represents gravitational settling and adsorption onto surfaces. These modules are executed in series in accordance with operator splitting and with their own internal time-stepping schemes. When a certain time step, known as the synchronization time step, is reached, execution of one module concludes, and the next module is invoked to update its science process to the same synchronization time step. This time stepping scheme is diagrammed in Figure 4.1.

We develop CMAQ-CUDA for CMAQ version 5.4. We focus on the gas-phase chemistry module as it represents one of the major computational bottlenecks for the CMAQ CTM (CCTM) executable, with computation burdens illustrated in Figure 4.2.

CMAQv5.4 has several different gas-phase mechanisms available and three ODE solver options to numerically integrate the gas-phase ODE system. The mechanisms are versions of the Carbon Bond (CB) mechanism, Regional Atmospheric Chemistry Mechanism (RACM), SAPRC07 mechanism, Community Regional Atmospheric Chemistry Multiphase Mechanism (CRACMM), and the 2D-VBS mechanism (W. P. Carter 2010; Sarwar et al. 2008; Goliff, Stockwell, and Lawson 2013; Pye et al. 2023; Yarwood 2010). The solvers are the Euler backwards iterative (EBI), third-order Rosenbrock (ROS3), and sparse matrix vectorized Gear (SMVGEAR) solvers (Byun and Schere 2006; Sandu et al. 1997; Z. Jacobson and Turco 1994; *CMAQv5.2 Operational Guidance Document* 2017). These solvers are listed in descending order of speed and ascending order of accuracy (*CMAQv5.2 Operational Guidance Document* 2017). The EBI solver is actually a family of solvers because its code implementation must be tailored to the chosen mechanism. In contrast, the ROS3 and SMVGEAR solvers are mechanism agnostic: any given mechanism can be solved by these codes. Since ROS3 is the faster of the two mechanism-agnostic solvers, we focus our efforts on migrating ROS3 to

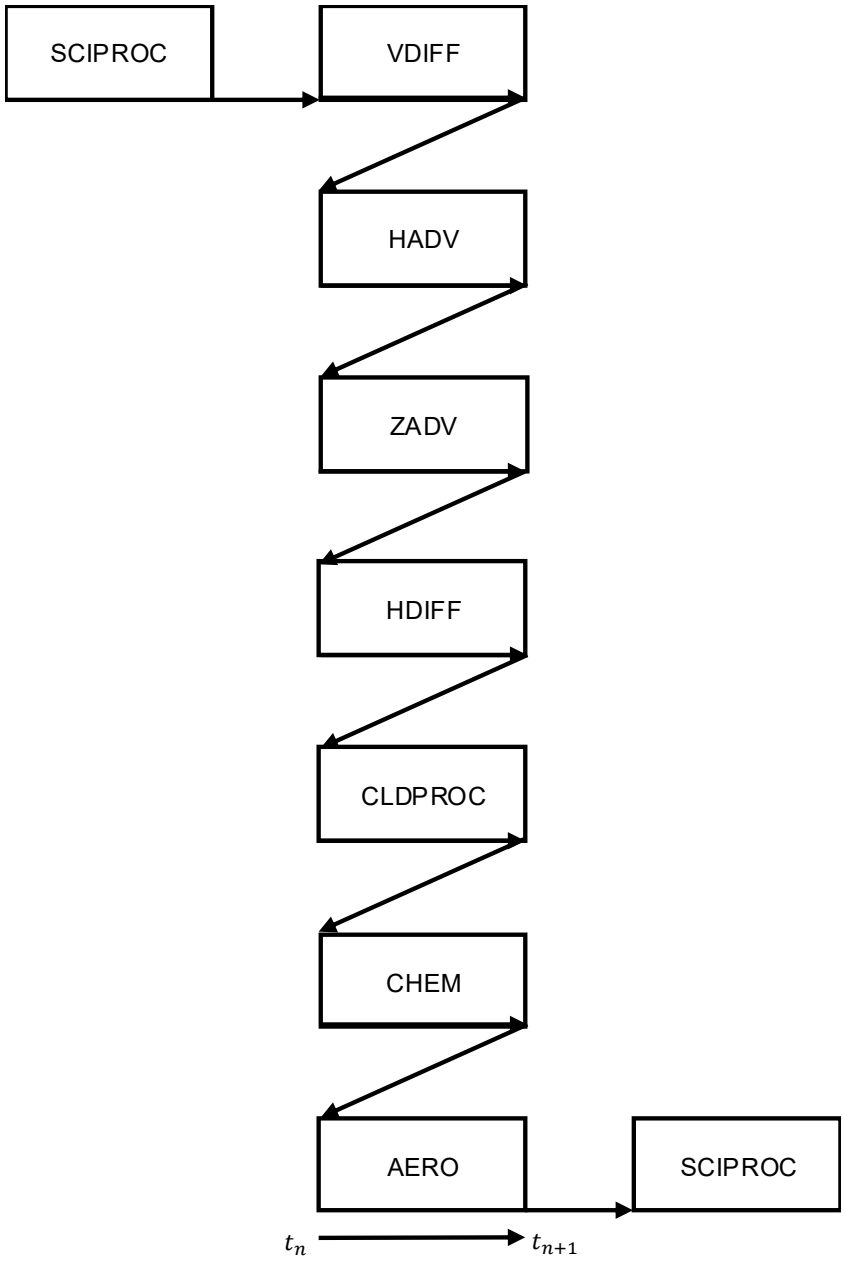

Figure 4.1: Schematic of CMAQ science process time stepping structure. SCIPROC is the Fortran routine that calls the science process routines. VDIFF is vertical diffusion. HADV is horizontal advection. ZADV is vertical advection. HDIFF is horizontal diffusion. CLDPROC denotes cloud processes. CHEM is gas-phase chemistry. AERO denotes aerosol processes. Times t_n and t_{n+1} are the nth and subsequent synchronization time steps, respectively.

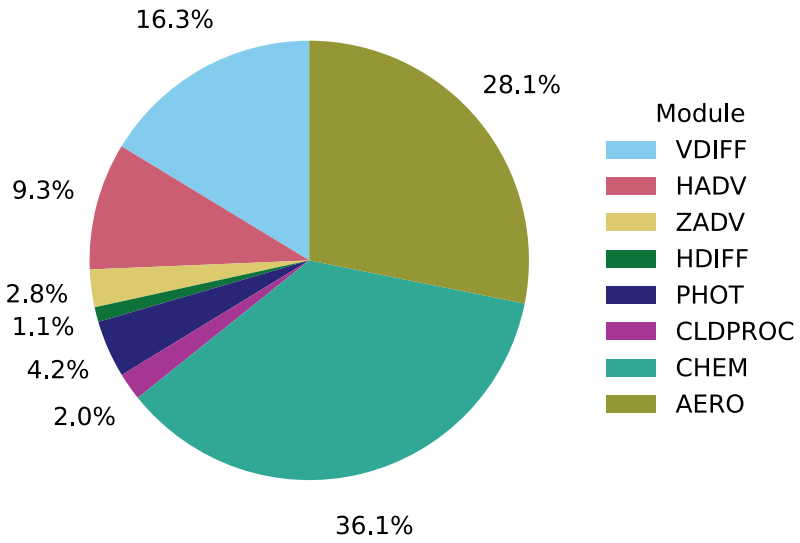

Figure 4.2: Pie chart of module timing burdens per time step for an 8-core simulation using the ROS3 solver and CB6R5 chemical mechanism with aero7 SOA treatment and standard cloud chemistry for the CMAQ 2018 12NE3 benchmark scenario.

GPUs.

Rosenbrock solvers like ROS3 are multistage semi-implicit schemes (Sandu et al. 1997). An s-stage Rosenbrock scheme for the solution of the autonomous equation $\dot{y} = f(y)$ has the form

$$k_i = hf\left(y_n + \sum_{j=1}^{i-1} \alpha_{ij} k_j\right) + hJ \sum_{j=1}^{i} \gamma k_j \tag{4.3}$$

where y_n is the numerical solution at time n, f prescribes the form of the ODE, h is the time step, J is the Jacobian, and the constants s, b_i, α_{ij}, and γ are chosen based on considerations of numerical stability and order of accuracy (Sandu et al. 1997). For ROS3, $s = 3$ (*CMAQv5.2 Operational Guidance Document* 2017; Damian et al. 2002).

The implementation of ROS3 within the CCTM executable is written to leverage parallelization across multiple CPU cores. CPU arallelization is accomplished with the Message-Passing Interface (MPI), which is a software standard that facilitates parallel computing by mapping processes that are generated to execute instruction threads to logical processors and enabling inter-process communication (Walker and Dongarra 1996). The simulation domain is divided across MPI processes and each process handles its own subdomain. To accelerate the evaluation of the CHEM module on each subdomain, the matrices describing the scheme are split into smaller block matrices, and each subdomain loops through these blocks. The block size is chosen at compile time. Figure 4.3 diagrams the domain-process-GPU mapping

for our simulation configuration.

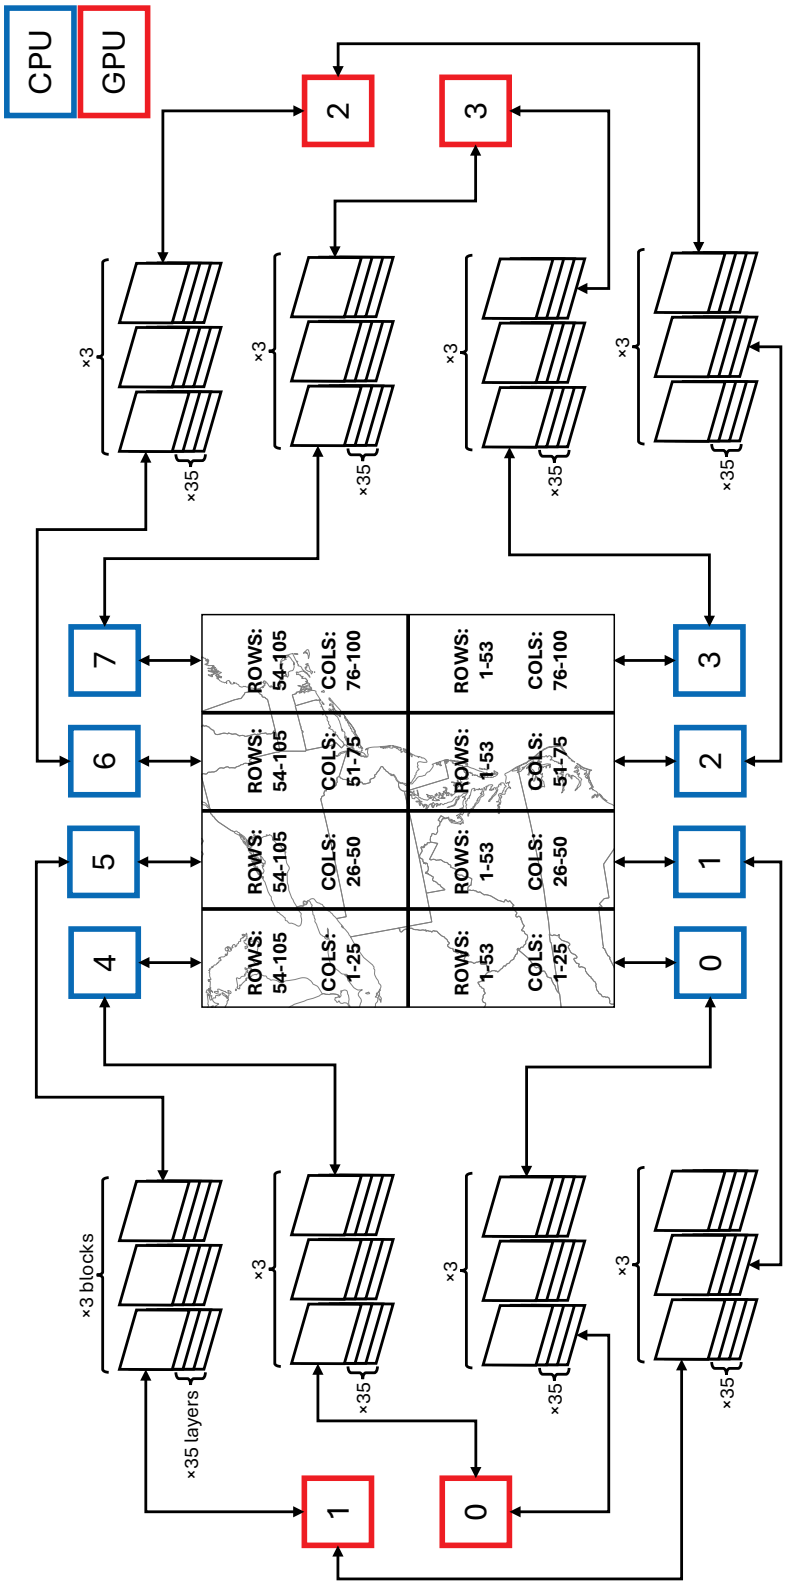

Figure 4.3: Schematic of domain subdivision and mapping to MPI processes. Each MPI processes is given its own CPU core and shares GPUs as instructed within our CMAQ-CUDA runtime configuration.

In the standard CMAQv5.4 case, each process splits its subdomain into 928 blocks of 50 grid cells, whereas the CMAQ-CUDA case splits each subdomain into just 3 blocks of 17,408 grid cells each, with the caveat that in both cases the final block may be smaller when the block size does not evenly divide the subdomain.

The results we present in this study come from simulations set up to mirror the EPA's 2018 12NE3 case study used to benchmark CMAQv5.4 (US EPA 2023). The study area encompasses the northeastern United States with a grid that starts at 84.41W, 35.16N with 100 12-km-wide columns eastward, 105 12-km-wide rows northward, and 35 vertical layers for a total of 367,500 grid cells. The benchmark simulates hourly concentrations for July 1st and 2nd, 2018 and uses the CB6R5 chemical mechanism with aero7 SOA treatment and standard cloud chemistry. We use the same time period, time resolution, and chemical mechanism as well as with the additional mechanisms SAPRC07 and RACM2 with aero6 SOA treatment and standard cloud chemistry to conduct our tests.

Our simulations' compute environment has Intel Xeon Skylake 6130 CPUs rated at 2.1 GHz with NVIDIA RTX 2080 Ti GPUs. Models are compiled with an OpenMPI 5.0.2 wrapper around the Intel OneAPI 2023.1.0 compiler for CPU code and the NVHPC SDK version 23.9 for GPU code. Our runtime configuration uses 8 CPU cores and 4 graphics cards.

Three popular GPGPU solutions are OpenCL, OpenACC, and NVIDIA CUDA. The ROS3 code, like the rest of CMAQ, is written in the Fortran programming language. We, therefore, leverage CUDA Fortran in our code migration from CPU to GPU as OpenCL is based on C/C++. This avoids the need for interfacing with C from Fortran and enables a native Fortran implementation of CMAQ-CUDA. Although OpenACC supports Fortran, CUDA Fortran offers greater control over optimizations.

CMAQ's ROS3 implementation is split across four subroutines: RBFEVAL, RBJACOB, RBDECOMP, and RBSOLVE. RBFEVAL computes the derivative with respect to time, RBJACOB computes the Jacobian matrix, RBDECOMP computes the LU decomposition of the system to facilitate its solution, and RBSOLVE computes the numerical solution to the linear system of equations based on the LU decomposition. Another subroutine, RBSOLVER, calls the preceding subroutines in the order listed, iterates until solution convergence, and outputs the concentration field. A further subroutine, RBDRIVER, calls RBSOLVER in a loop over blocks and fills the MPI process subdomain's concentration field with RBSOLVER's arrays until the whole subdomain's concentration field is updated. Each MPI process does this to update the entire simulation domain's concentration field. This structure is diagrammed in Figure 4.4.

A typical Fortran nested `DO` loop in these routines has the following structure for species index `J`, grid cell index `I`, and array `A`:

```
DO J = 1 TO NUMBER OF SPECIES
    DO I = 1 TO NUMBER OF CELLS
        UPDATE A( I,J )
```

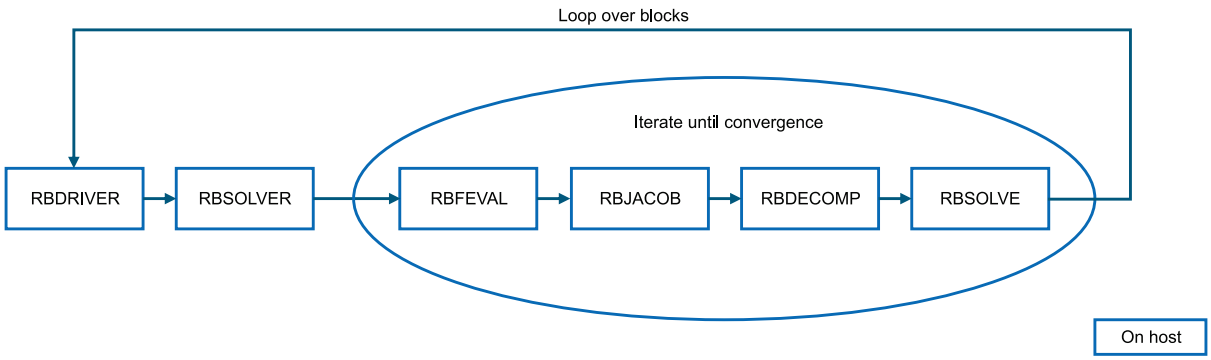

Figure 4.4: Schematic of the CMAQ ROS3 code structure. Arrows represent data transfer. All routines are on the CPU.

```
    END DO
END DO
```

These loops step through grid cells to update the spatial dimension of array `A` and do this for every species in the mechanism.

Because of operator splitting, transport processes (i.e., advection and diffusion) are inactive during execution of the CHEM module, so each grid cell in the simulation domain may be treated independently during CHEM execution, and every cell in a block can be updated in parallel. We, therefore, target the block size dimension of the ROS3 arrays for GPU parallelization: each grid cell is handled by individual GPU threads in parallel.

```
DO J = 1 TO NUMBER OF SPECIES
    DO I = 1 THREAD INDEX TO NUMBER OF CELLS, STEP BY NUMBER OF THREADS
        UPDATE A( I,J )
    END DO
END DO
```

The index `I` now represents a GPU thread index, with each thread having its own unique index and executing its own copy of the GPU device routine, or kernel. The inner loop becomes a grid-stride loop, where the loop stride is the size of the grid of threads on the GPU. This enables the code to loop through portions of the block if the number of threads initialized on the GPU is insufficient to cover the block at once, although this would slow the code down. Ideally, we initialize enough threads to handle the entire block in parallel. In this case, the grid-stride loop executes just once. This leads to a new ROS3 structure, which we diagram in Figure 4.5.

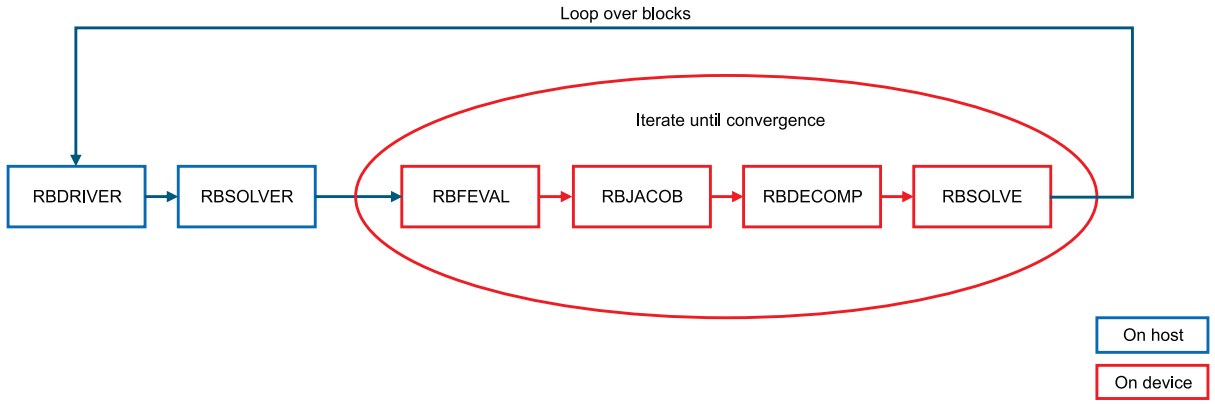

Figure 4.5: Schematic of the CMAQ-CUDA ROS3 code structure. Arrows represent data transfer.

With the convergence loop entirely on the device, data transfer only occurs once at the start and once at the end of each synchronization time step. Data transfer between GPU subroutines does not change the data's location in memory, as the convergence loop subroutines are contained in one Fortran module, so arrays are simply updated in place. This greatly relieves the bandwidth bottleneck exposed by data transfer and results in a performant CMAQ-CUDA implementation.

Timing is calculated as the time the model spends in the CHEM science process module, which accounts for the time that the primary MPI process spends on computation in that module and, for CMAQ-CUDA, data transfer overhead. Discrepancies between CMAQ and CMAQ-CUDA concentration fields are reported as mean bias or normalized mean error. With CMAQ output denoted y_i and CMAQ-CUDA output denoted $\hat{y}_i$, we define the following metrics:

$$NME = \frac{\sum_{i=1}^{N} |\hat{y}_i - y_i|}{\sum_{i=1}^{N} y_i} \tag{4.4}$$

where N is the number of grid cells and/or time steps and i is the index.

4.3 Results and Discussion

The objective of this work is to accelerate CMAQ's gas-phase chemistry module, CHEM, by leveraging GPU hardware. Figure 4.6 displays the timing of the CHEM module per synchronization time step of CMAQv5.4 versus CMAQ-CUDA for the chemical mechanisms CB6R5, SAPRC07TC, and RACM2 with aero6 SOA treatment (aero7 for CB6R5) and standard cloud chemistry. These mechanisms contain 349, 752, and 411 reactions, respectively.

We see that, in all cases, CMAQ-CUDA executes CHEM with the ROS3 solver substantially faster than CMAQv5.4. Comparing medians, CMAQ-CUDA executes CB6R5 in 50%

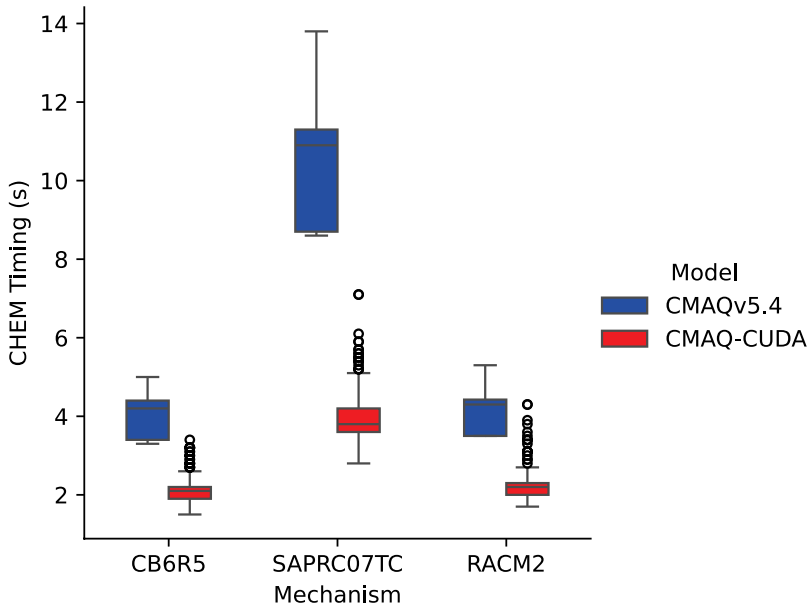

Figure 4.6: Boxplots of CHEM module timing per synchronization time step for CMAQv5.4 and CMAQ-CUDA across three mechanisms.

as much time as CMAQv5.4, SPARC07TC in 35% as much time, and RACM2 in 51% as much time. While there is more variability in CMAQ-CUDA timing, the slowest outliers remain near or below the median timing for CMAQv5.4. The variability present in each distribution is due to variation in the number of loops required for the solver to converge, as time steps with greater degrees of nonlinear behavior require more iterations to converge.

Figure 4.7 presents the total simulation walltime for our three mechanisms. Two features stand out. First, the differential is greatest for the most complex mechanism, indicating that studies simulating larger mechanisms stand to benefit the most from GPU acceleration. Second, the differences between CMAQ-CUDA walltimes across the different mechanisms are smaller than the differences between CMAQv5.4 walltimes across the same mechanisms. This indicates that CMAQ-CUDA scales well in the sense that the rate at which it slows with increasing mechanism complexity is less than the rate at which CMAQv5.4 slows with increasing mechanism complexity. Moreover, for each mechanism, we see notably smaller walltimes even though we have only accelerated one of CMAQ's eight science process modules. This suggests that accelerating other modules in the same way could greatly reduce the walltimes required for CMAQ simulations, saving valuable compute resources.

We explore the impact of varying the block size parameter in Figure 4.8. Blocking is an efficiency consideration for the CHEM module, so it is unsurprising that we see little to no variation in other modules' timing as we vary block size. CMAQv5.4 is fastest at small block sizes and slows exponentially as we increase the block size. CMAQ-CUDA is also sensitive to

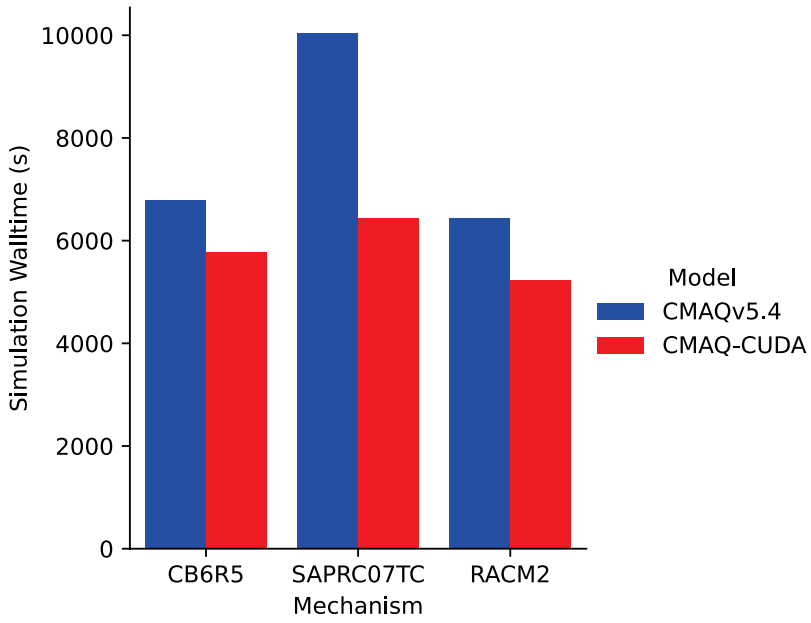

Figure 4.7: Bar plots of total simulation walltime for CMAQv5.4 and CMAQ-CUDA across three mechanisms.

block size, but in this case, we see speedup—up to a point—when we increase the block size. The speedup associated with increasing block size is indicative of increasing parallelization. In the case of GPUs, small block sizes greatly slow execution because this results in more blocks per subdomain. Each block requires kernel calls with their own scheduling overhead in addition to data transfer overhead. Hundreds of small blocks indicate that that overhead gets compounded hundreds of times over and in serial. On the other hand, we see slowdown at very large block sizes as this results in the grid-stride loops executing more than once. The ideal block size indicated by Figure 8 is specific to the NVIDIA RTX 2080 Ti graphics card. It is the number of threads that can reside simultaneously on the GPU when accounting for multiple MPI processes sharing a GPU. Refer to the SI for details.

GPUs are generally designed to excel at single-precision calculations. CMAQ's CHEM module performs double-precision computations. Each double-precision number requires two registers to reside on the GPU for computation. The result is that compiling our kernel that is intensive in double-precision operations reserves 152 registers per thread, regardless of block size. This translates into 37.5% device occupancy, which measures how well a kernel utilizes a GPU's capacity for computation. Higher occupancy generally equates to better performance. To attain greater occupancy for CMAQ-CUDA, we would need to reduce register pressure. One way to achieve this would be to perform single-precision computation, but doing so sacrifices solution accuracy. Indeed, we implemented and tested a single-precision version of

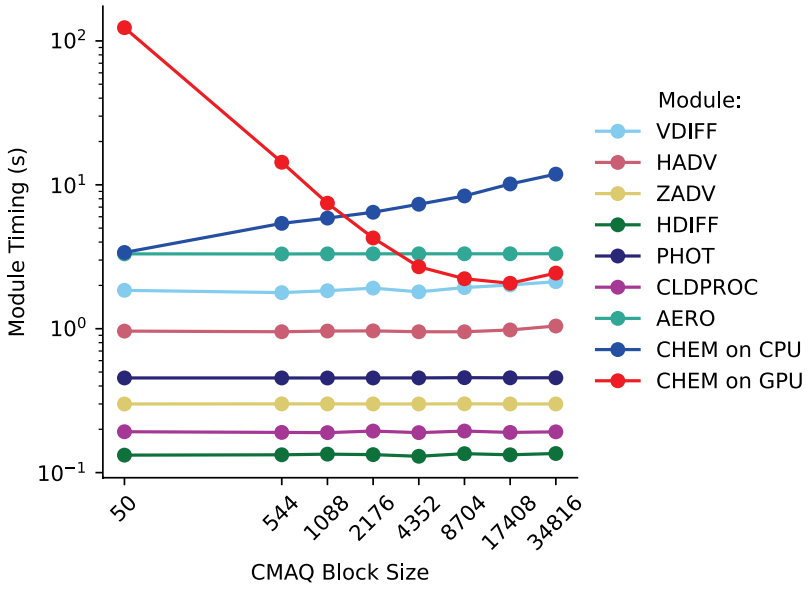

Figure 4.8: CMAQ science process module timing per synchronization time step as block size varies for CB6R5.

CMAQ-CUDA. The solver failed to converge.

While future work should seek to reduce register pressure and accelerate CMAQ-CUDA, even our current low-occupancy implementation attains the significant performance gains documented above, and these are even with the NVIDIA RTX 2080 Ti—a consumer graphics card. Tailoring the block size to the hardware specifications of newer, faster graphics cards may further accelerate CMAQ-CUDA even if occupancy changes little.

In addition to acceleration, we must also ensure that our GPU implementation of ROS3 maintains agreement with its CPU implementation. We want and expect CMAQ-CUDA's output to replicate the standard CMAQ output given the same model inputs, parameters, and configurations. Indeed, according to CMAQ documentation, we are typically satisfied with concentration field discrepancies within 1% when comparing compilations across different computer systems (*CMAQv5.2 Operational Guidance Document* 2017).

We compare MB in the surface layer for three key species between a full simulation and a simulation executed with CHEM as the only active module in Figure 4.9. This is to investigate the effect of other modules on the overall simulation agreement between CMAQv5.4 and CMAQ-CUDA.

When CHEM is the only active module, CMAQv5.4 and CMAQ-CUDA are in significantly better agreement. This suggests that transferring concentration fields between modules incurs additional error in CMAQ-CUDA, though we note that for simulations with CHEM (but no other modules) deactivated, errors are identically zero (Figure 6.4). This

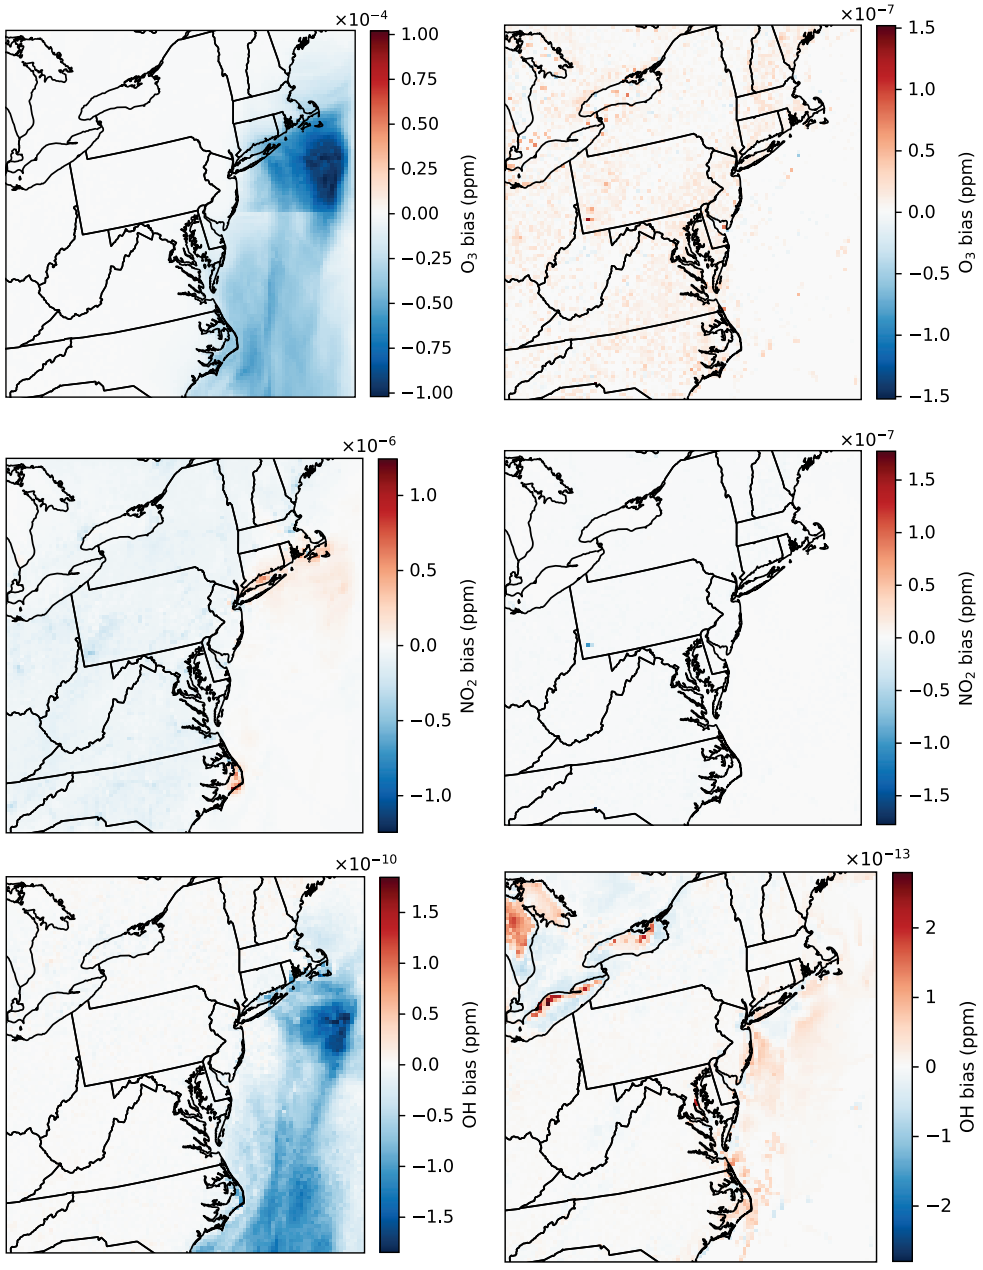

Figure 4.9: Maps of MB averaged in time for three key species in CB6R5. The left column is the MB between full simulations, the right column is the MB between simulations run with only the CHEM module active.

suggests that the small errors in the CHEM module are magnified by operations in other modules. We posit that this is due to roundoff error magnification in reactive trace species

whose concentrations are so small that such errors are disproportionately large. This could be due to architectural differences in how our CPUs and GPUs handle roundoff. We investigate this hypothesis in Figure 4.10.

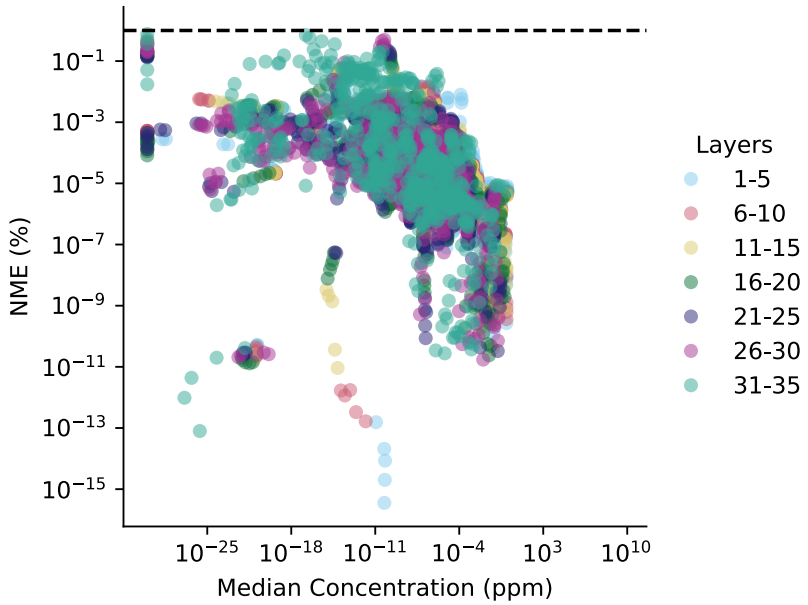

Figure 4.10: NME against median concentration for all species and all layers in the CB6R5 CHEM-only simulation. Each species appears 35 times, once for each vertical layer. NME is averaged in space (per layer) and time. Our 1% threshold is marked.

We see a tendency for NME to increase as median concentrations decrease, supporting our hypothesis. This may also help explain the patterns we see in our full simulations. Certain reactive trace species concentrations may be especially small and prone to roundoff error magnification over water where anthropogenic emissions are minimal.

Figure 4.11 compares NME in the surface layer for our three mechanisms as well as a benchmark comparison for which the NME is taken between our CMAQv5.4 simulation and the published EPA simulation results. We see that, at the surface, CMAQ-CUDA NME stays within our threshold for most species, and all CMAQ-CUDA NME distributions are concentrated below the benchmark comparison. This indicates that the discrepancy due to hardware differences between our cluster's CPUs and GPUs is smaller than between our architecture and the EPA cluster's architecture.

We expand our error analysis in Figure 4.12 to examine more vertical layers. There does not appear to be a strong dependence on layer until we reach the 26-30 layer bin, where the distributions begin to shift towards larger NME. Following our earlier hypothesis, it may be the case that errors grow in the upper layers because we exit the planetary boundary layer and certain reactive trace species' concentrations decrease.

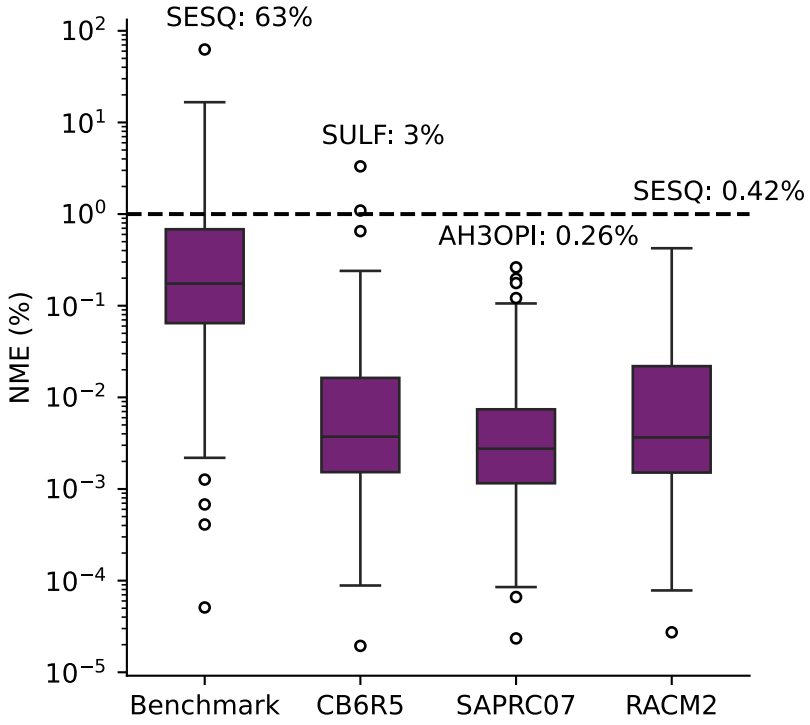

Figure 4.11: NME for the surface layer across the benchmark and three mechanisms. Each member of the distributions is a different species' NME averaged in space and time. Highest errors are annotated with the corresponding species and its NME value. Our 1% threshold is marked.

In all cases, most species remain within our threshold. In all cases, there are outliers beyond our threshold, and in the top 10 layers, some species in the distributions' tails exceed our threshold. In Figure 6.5 we see that, in our CHEM-only simulation, all species at all layers, without exception, remain below our threshold. This meets expectations considering our preceding discussion.

We acknowledge the following limitations in our work, in addition to the roundoff error magnification just discussed. Because we migrate CMAQ's ROS3 solver, the integrated source apportionment method (ISAM) tool is unavailable in CMAQ-CUDA. ISAM requires the EBI solver. Further, the direct decoupled method (DDM) sensitivity analysis tool is not preserved in our current implementation, nor is integrated reaction rate (IRR) analysis.

Despite these limitations, CMAQ-CUDA proves the viability of migrating CMAQ modules to GPU code and highlights the potential performance gains of doing so. Next steps include integrating DDM and IRR into CMAQ-CUDA to enable and accelerate these analysis tools. We also want to migrate additional science process models to CUDA Fortran, starting with the next largest computational burdens: AERO and VDIFF (Figure 4.2).

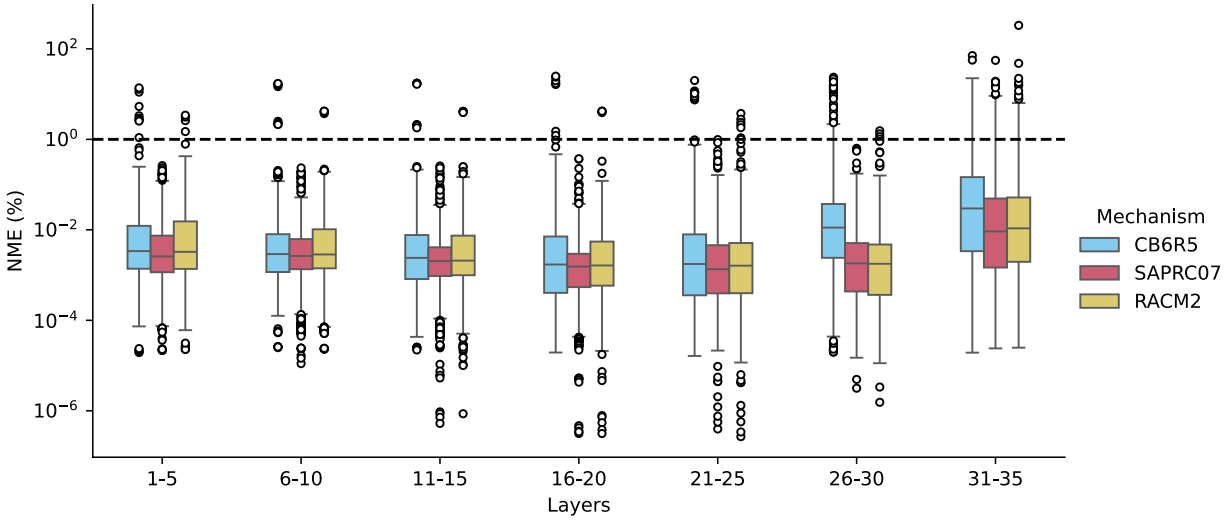

Figure 4.12: NME for binned vertical layers across three mechanisms. Each member of the distributions is a different species' NME averaged in space and time for the corresponding layer, such that each box plot has 5 members for each species: one for each layer in the corresponding bin. Our 1% threshold is marked.

4.4 Conclusions

CMAQ-CUDA successfully migrates CMAQ's ROS3 solver to GPU hardware. Our implementation accelerates CMAQ's CHEM module within the modular framework of CMAQ, enabling simulations with all science process modules active. Because we see agreement degrade when CHEM interacts with other modules, we note the importance of expanding our hardware migration to other modules in addition to CHEM.

CMAQ-CUDA highlights the performance gain inherent in GPU hardware acceleration. We have, therefore, reduced the computational burden imposed by CHEM on CMAQ, and in doing so, emphasized the potential for further model acceleration offered by migrating other science process modules to GPU hardware. We have also demonstrated desktop-scale computational improvements in earlier versions of this work, which supports the use of the code in situations where access to HPC resources is limited (Do et al. 2023).

With accelerated chemistry, CMAQ-CUDA can support the integration of complex chemistry into forecasting operations as well as the resolution of model domains into finer grid structures than is feasible with traditional hardware. These benefits can inform improved air quality regulatory policy because higher resolution simulations can elucidate air pollution dynamics that are obscured by coarse grids, like sharp concentration gradients in the vicinity of neighborhoods located near industrial activity. On a coarse grid, predicted concentrations may be diluted throughout the grid cell's volume, whereas in reality, concentrations may

be much higher, concentrated around the emissions source. This would artificially reduce the neighborhood's predicted exposure, and if there aren't co-located monitoring stations to alert regulators to the true exposure burden, then these communities may not experience beneficial interventions. Therefore, model developments like those presented here are important for reducing communities' exposure burdens.

Chapter 5

Developing the AMBRS Project Platform

5.1 Introduction

Chapter 4 explores model development with implications for improved air quality policy-making through enhanced fine scale simulations that can support local control strategies. Where Chapter 4 developed an accelerated implementation of the CMAQ model's chemistry module, this chapter turns to development of a platform for improving arbitrary aerosol microphysics modules by quantifying biases baked into the modules' designs. Like Chapter 4, this can indirectly improve air quality policy, in this case by improving model aerosol predictions. Because modeled aerosols contribute substantial climate projection uncertainty, these results can also inform climate policy.

Background

Aerosol particle behavior and particulate effects on air quality and climate are governed by a variety of physical and chemical processes. Particle size and composition affect particle microphysics processes like condensation and coagulation. Condensation is affected by particle size because gas adsorption efficiency is influenced by particle surface area (Seinfeld and Pandis 2016), while composition influences gas absorption through pH effects (S. Gao et al. 2004; Y. Zhang et al. 2019; Lei et al. 2022). Coagulation is influenced by particle sticking efficiency, which depends on the relative sizes of colliding particles and is affected by the presence of liquid coatings, whether water or organic (Riva et al. 2016; Schmedding et al. 2019). This small sample of factors influencing microphysics processes gives rise to complex dynamics that produce highly coupled physicochemical systems governing the evolution of particle populations. Accurately representing these dynamics is essential for simulating particulate effects on air quality and climate (Seigneur et al. 1986). For instance, the particle size distribution is a key determinant of air quality metrics like lung-deposited surface area (LDSA) and climate metrics like cloud condensation nuclei (CCN). Since particulate lung deposition

fraction is a function of particle size, LDSA is controlled by the particle size distribution. CCN also depends on the size distribution, because cloud nucleation through water vapor condensation onto CCN aerosol depends on particle size. These same metrics are directly or indirectly influenced by particle composition as well, which controls particle hygroscopicity and therefore CCN activity, as well as toxicity, which influences the exposure risk posed by LDSA. In this chapter, we will explore the influence of the particle size distribution representation by comparing a modal representation with a particle-resolved representation. We first compare and contrast these representations and another major alternative, the sectional representation.

Because the number and variety of microphysics processes and influences thereupon is sizable and complicated, accurate simulation of particle microphysics presents computational challenges. Model representations of particle size distributions and treatments of microphysics are necessarily simplified for computational tractability in large-scale 3-D Earth system models like the Energy Exascale Earth System Model (E3SM) (Golaz et al. 2022). E3SM's atmosphere component uses the Modal Aerosol Module (MAM) to treat aerosol chemistry and microphysics (Liu, Easter, et al. 2012; Liu, Ma, et al. 2016). MAM represents the particle size distribution with a modal parameterization that describes the full size distribution as a sum of individual lognormal modes that are described by moment equations, with the kth moment defined as

$$M_k = \int_{-\infty}^{\infty} d_p^k n(\ln d_p) \mathrm{d} \ln d_p \tag{5.1}$$

where d_p is the particle diameter and n is the particle number density.

The lognormal moments can be alternatively defined in terms of the modes' geometric mean diameters d_{gn} and geometric standard deviations σ_g:

$$M_k = N d_{gn}^k \exp\left(\frac{k^2}{2} \ln^2 \sigma_g\right) \tag{5.2}$$

where N is the number of particles in the mode.

Alternatives to the modal representation of particle size distributions include sectional and particle-resolved representations. Sectional representations, in essence, represent the size distribution as a histogram over particle diameter. Across bins, or sections, of a sectional representation, particle number concentration and composition vary, but within sections, populations are internally mixed. Internal mixing refers to the mixing state where all particles have the same composition. In modal representations, particle composition varies across modes but internal mixing is assumed within modes. On the other side of the mixing state spectrum, externally mixed populations are those where every particle is composed of a single constituent but distinct particles are allowed their own distinct constituent. Particle-resolved representations can be defined in part by their relaxation of mixing state assumptions, allowing different particles to contain multiple constituents but not enforcing that their mass fractions are the same across particles, falling between internal and external mixing on the mixing state spectrum.

Layered on top of this variety in particle size representations are particulate microphysics and chemistry. Treatments of physicochemical processes necessarily differs between modal, sectional, and particle-resolved representations, but treatments can differ even between modal models. MAM4 in E3SM, for instance, treats inorganic particulate phase partitioning using the Model for Simulating Aerosol Interactions and Chemistry (MOSAIC) (Zaveri et al. 2008), whereas the Community Multiscale Air Quality (CMAQ) chemical transport model's aerosol module, which can be thought of as akin to "MAM3," employs ISORROPIA (Nenes, Pandis, and Pilinis 1998) for inorganic particulate phase partitioning.

This diverse set of size representations, model assumptions and simplifications, and physicochemical process treatments converge to yield substantial variability in simulated aerosol predictions. Indeed, this variability contributes significant uncertainty to climate projections and modeled aerosol metrics like aerosol optical depth (AOD). The greatest uncertainty in radiative forcing projections reported by the Intergovernmental Panel on Climate Change comes from aerosols (Calvin et al. 2023), while simulated AOD variability in model intercomparison projects, measured as one standard deviation, can reach 45% (Vogel et al. 2022).

Fierce et al. (2024) find that MAM4 simulations can require up to an order of magnitude higher environmental water vapor supersaturation to activate an equivalent fraction of CCN as the PartMC benchmark. PartMC does not parameterize the size distribution with moment equations or otherwise, instead directly building the distribution from a sample of computational particles, freeing constraints on the size distribution and thereby allowing the particle population to grow more naturally. Similarly, by virtue of its particle-resolved representation, PartMC does not impose internal mixing like MAM4.

Like Fierce et al. (2024), Zheng et al. (2021) choose PartMC as their benchmark. Because PartMC is particle-resolved, however, Zheng et al. (2021) train an emulator on PartMC to predict mixing state at a global scale. They find mixing state to differ between PartMC and MAM4 significantly in some regions of the domain, with MAM4 predicting more extreme mixing states relative to the more intermediate mixing states predicted by the PartMC emulator.

While these two studies make strides towards quantifying the magnitude of structural error in MAM4 as a representative reduced aerosol model, more work is needed to elucidate and isolate the sources of that error at a process level except by qualitative inference. Other investigations into aerosol model uncertainty generally focus on radiative forcing, an emergent property yet another step removed from shedding quantitative light on sources of structural error. This represents a major knowledge gap and an opportunity for the aerosol modeling community to identify problematic model design choices to prioritize for model development.

The Aerosol Model Benchmark Repository and Standards Project

To fill that knowledge gap, the Aerosol Model Benchmark Repository and Standards (AMBRS) Project provides a platform to facilitate perturbed parameter ensemble (PPE) box

model simulations. This enables fast and robust model intercomparison, supporting error quantification. Like the works of Zheng et al. (2021) and Fierce et al. (2024), AMBRS proposes PartMC as a benchmark model and supports MAM4 integration. By driving PPEs for box models, AMBRS enables benchmarking against the full complexity of PartMC rather than an emulator. This provides insight into how design choices made in reduced models like MAM4 compare to those made in a particle-resolved model like PartMC, elucidating the impacts of those simplifications and suggesting how inter-model variability might arise.

AMBRS is a software platform written primarily in Python and designed to present a common interface to disparate aerosol models through unified model inputs. AMBRS began as a collaboration between Dr. Laura Fierce, Dr. Benjamin Murphy, Professor Nicole Riemer, Dr. Jeffrey Johnson, and Dr. Matthew Dawson. Engaging the aerosol modeling community is front of mind in AMBRS development, so AMBRS includes a CMake-based tool, AMBuilder, to automate the box model compilation process for a streamlined setup process. This way, modelers across varying degrees of familiarity with PartMC and MAM4 can benchmark their scenarios of interest with minimal hassle. For instance, Seigneur et al. (1986) provide a set of scenarios to test aerosol model performance under different representative conditions. AMBRS would provide simulation results from PartMC configured for given conditions so users can compare their models' simulation results to the PartMC benchmark. More concretely, Seigneur et al. (1986) provide conditions representative of an urban setting, which is a stringent test for coagulation performance because of the significant change over time in the size distribution. PartMC is known to converge to analytical coagulation solutions (DeVille, Nicole Riemer, and West 2019), so it provides a strong benchmark for this scenario. Users can measure their models' coagulation performance by comparing to the PartMC urban conditions benchmark.

AMBRS currently supports PartMC, as the proposed benchmark model, and MAM4, a representative reduced model widely used in large-scale Earth system models (ESMs). PartMC can be built coupled with the Chemistry Across Multiple Phases (CAMP) module and we contributed code to couple MAM4 to CAMP and automate CAMP configuration through the AMBRS interface.

5.2 Materials and Methods

AMBRS

AMBuilder

The top level of the AMBRS platform is AMBuilder, which uses CMake to automate the compilation process for the PartMC and MAM4 box models. AMBRS automatically builds and installs PartMC, MAM4, optionally CAMP, and all of their dependencies, requiring only a handful of command line invocations. To build PartMC and MAM4 with CAMP coupling:

```
> cmake -S . -B build -DCMAKE_INSTALL_PREFIX=<prefix> -DENABLE_CAMP=ON
> cd build
> make
> make install
```

These four commands will download and install PartMC, MAM4, CAMP, and their dependencies into the `<prefix>` directory without any further user input required.

AMBRS Driver

After AMBuilder has compiled PartMC and MAM4 (and, optionally, CAMP), the box models can be used to run PPEs through the AMBRS driver Python interface. The driver is used to specify ensemble parameters and distributions to sample. To specify a PPE, one must provide the aerosols and gases along with their molar masses and, for aerosols, their densities and hygroscopicities to the `ambrs GasSpecies` and `AerosolSpecies` classes, respectively. Once gases and aerosols are defined, a PPE may be specified:

```
spec = ambrs.EnsembleSpecification(
    name = "ppe_spec",
    aerosols = tuple(aerosols),
    gases = tuple(gases),
    size = ambrs.AerosolModelSizeDistribution(
        modes = [
            ambrs.AerosolModeDistribution(
                name = name,
                species = species,
                number = number,
                geom_mean_diam = geom_mean_diam,
                log10_geom_std_dev = log10_geom_std_dev,
                mass_fractions = mass_fractions,
            )
        ]
    ),
    gas_concs = gas_concs,
    flux = flux,
    relative_humidity = relative_humidity,
    temperature = temperature,
    pressure = pressure,
    height = height,
)
ensemble = ambrs.lhs(spec, n=n)
```

This code block illustrates the specification of an `ambrs` PPE. This PPE defines an aerosol particle population using the `AerosolModelSizeDistribution` class containing one mode specified by the `AerosolModeDistribution` class. The mode is defined by its constituent species, its initial number concentration in $\#/m^3$, its geometric standard deviation (GSD) in m, the log base 10 of its geometric standard deviation (unitless), and the mass fractions of its constituent species. The PPE is further specified by the initial concentrations in mol/mol of the gases given by `gas`, a flux parameter, the relative humidity given as a fraction, pressure in Pa, and the mixing layer height in m. The PPE specification is then realized by sampling an ensemble of size `n` with the `lhs` method of `ambrs`, which invokes Latin hypercube sampling of any parameters given as a sampling distribution with a `ppf` method.

To run the PPE, one must first choose what processes to simulate by initializing an `AerosolProcesses` object:

```
processes = ambrs.AerosolProcesses(
    gas_phase_chemistry = False,
    coagulation = True,
    condensation = False,
    nucleation = False
)
```

In this example, a simulation will be run where coagulation is the only active aerosol microphysics process. One should note that condensation refers to gas-particle phase partitioning for MAM4 but cloud water condensation for PartMC. To enable gas-particle phase partitioning in PartMC, one must configure CAMP with Henry's Law or SIMPOL.1 phase transfer reactions.

PartMC

PartMC is a stochastic particle-resolved aerosol model that implements particulate coagulation for a configurable number of computational particles which are sampled from a particle population through Monte Carlo methods (N. Riemer et al. 2009). PartMC is coupled to MOSAIC (Zaveri et al. 2008), with PartMC offloading gas- and particle-phase chemistry, particulate thermodynamics, and gas-particle phase partitioning to MOSAIC. PartMC itself also handles gas- and particle-phase emissions as well as dilution. The model's explicit simulation of individual particles obviates numerical diffusion concerns associated with approximate models, making PartMC a suitable benchmark model as established by N. Riemer et al. (2009) and McGraw et al. (2008) and demonstrated by Zheng et al. (2021) and Fierce et al. (2024), forming the basis for AMBRS' choice of PartMC as its benchmark model. PartMC has also been coupled with CAMP as an alternative to MOSAIC (Dawson et al. 2022).

MAM4

The Modal Aerosol Module (MAM) is a modal aerosol model originally developed for the Community Atmosphere Model version 5 (CAM5) in the Community Earth System Model version 1 (CESM1) (Liu, Easter, et al. 2012). MAM4 is one of several configurations for the Modal Aerosol Module (MAM), which was originally implemented for 3-mode and 7-mode configurations. MAM7 was developed as the benchmark configuration for CAM5, with the simpler MAM3 intended for long-term simulations spanning decades to centuries. The 4-mode version, MAM4, was developed later, also for CAM5, to improve on MAM3 (Liu, Ma, et al. 2016). The four modes in MAM4 are the canonical Aitken, accumulation, and coarse modes, along with a primary carbon mode. The species and inter-modal interactions treated by MAM4 are diagrammed in Figure 5.1. MAM includes gas-phase sulfur chemistry (DMS oxidation to SO_2 and SO_2 oxidation to H_2SO_4). The microphysics processes included are H_2SO_4 homogeneous nucleation; irreversible H_2SO_4 condensation to particulate SO_4^{2-} and reversible SOA condensation to and from organic vapors lumped into one species called SOAG; coagulation from the primary carbon and Aitken modes into the accumulation mode; and water uptake.

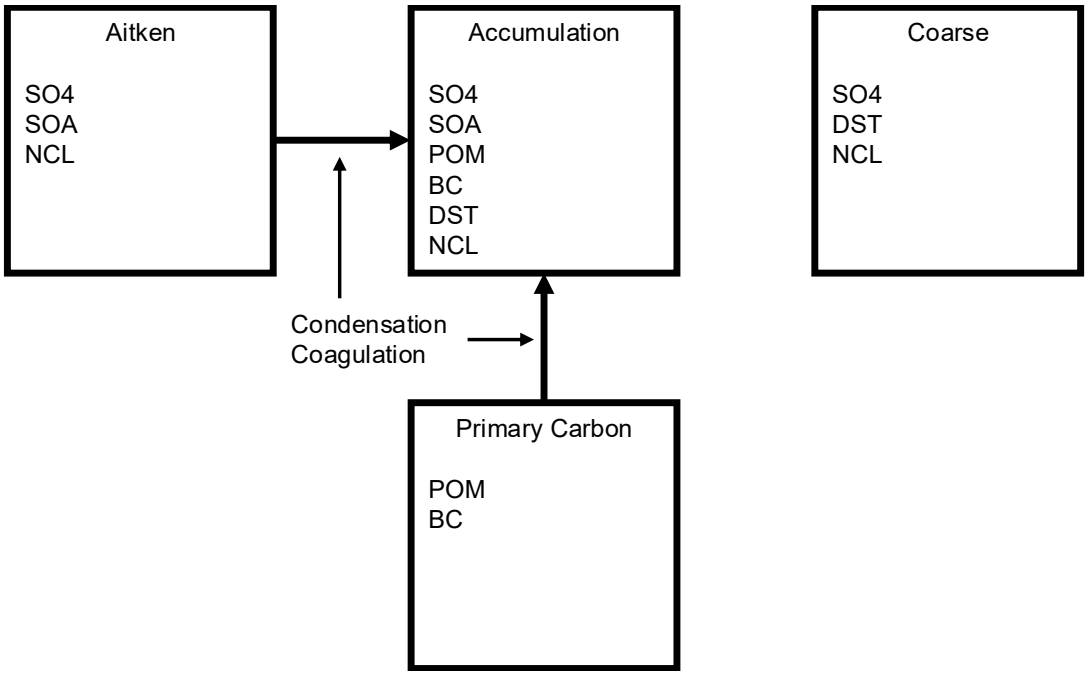

Figure 5.1: MAM4 modes schematic illustrating the species present in each mode and how mass can move between modes. Adapted from Liu, Ma, et al. (2016). SO4 is sulfate. SOA is secondary organic aerosol. NCL is sea salt. POM is primary organic matter. BC is black carbon. DST is soil dust.

CAMP

CAMP is a flexible multiphase chemistry module developed to support a wide variety of atmosphere models (Dawson et al. 2022). CAMP abstracts the aerosol representation to admit modal, sectional, and particle-resolved schemes and features a runtime JSON-based chemistry configuration interface. CAMP supports a variety of reaction types, including gas-phase emissions, gas- and condensed-phase Arrhenius reactions, Henry's Law and SIMPOL.1 phase transfer reactions, and first-order loss reactions, among others.

AMBRS Development

Dawson et al. (2022) coupled CAMP with PartMC, but AMBRS initially did not feature automated CAMP configuration through its Python interface. We expanded the AMBRS interface to automate CAMP configuration and we coupled MAM4 to CAMP to enable extended chemistry beyond MAM4's native schemes. To automate CAMP configuration through AMBRS, we introduced the `CAMP` class to `ambrs`, which is initialized with the relevant information CAMP needs based on the passed aerosol representation (i.e., modal or particle-resolved). The `CAMP` object will take its inputs and combine them with information from an `ambrs` PPE specification to generate a `camp.json` file for each ensemble scenario so that each scenario's CAMP configuration is aware of the sampled environmental conditions as well as gas and aerosol concentrations and parameters. To couple MAM4 to CAMP, we implemented the `mam4_camp` module, which contains subroutines to pass information between MAM4 and CAMP. Currently, this interface supports only gas-phase CAMP reactions.

5.3 Results and Discussion

We test our CAMP interfaces with a 12-hour SOA condensation simulation. Aerosol parameters are defined in Table 5.1. The mass fractions for each mode are 100% SOA, except for primary carbon which is 100% POM, although only the accumulation mode is initialized with a nonzero particle concentration. All gases are initialized to 0 ppb and SOAG is emitted at a rate of 1×10^{-8} mol/m^2/s. Pressure, temperature, and relative humidity are constant at 101325 Pa, 283.15 K, and 60%, respectively. These are based on the default hard-coded MAM4 box model values to facilitate direct comparison.

Table 5.1: SOA condensation simulation aerosol parameters by mode.

Parameter	Aitken	Accumulation	Coarse	Primary Carbon
Number concentration (#/m^3)	0	5×10^7	0	0
GMD (μm)	0.026	0.11	2.00	0.05
GSD	1.6	1.8	1.8	1.6

MAM4 computes reversible SOA condensation internally while SOAG emissions are offloaded to CAMP. PartMC handles SOAG emissions internally and offloads SOA conden-

sation to CAMP using a SIMPOL.1 vapor pressure parameterization configured to match MAM4's parameterization (Equation 5.3) (Liu, Easter, et al. 2012):

$$P^\circ(T) = P^\circ(298) \times \exp\left[\frac{-\Delta H_v}{R}\left(\frac{1}{T} - \frac{1}{298}\right)\right] \tag{5.3}$$

where $P^\circ(T)$ is the saturation vapor pressure in atm at temperature T in K, $P^\circ(298)$ is 1×10^{-10} atm, ΔH_v is the enthalpy of vaporization and equal to 156 kJ/mol, and R is the universal gas constant. If we take the $\log_{10}$ of Equation 5.3, we find

$$\log_{10} P^\circ(T) = \left(-\Delta H_v \frac{\log_{10} e}{R}\right)\frac{1}{T} + \left(-10 + \Delta H_v \frac{\log_{10} e}{298R}\right) \tag{5.4}$$

which is written in this form to identify the right-hand terms with the B_i in a SIMPOL.1 parameterization (Pankow and Asher 2008). Thus, Equation 5.4 can be represented in CAMP so that PartMC's treatment of SOA condensation is parameterized in the same way as MAM4's treatment.

We see that the evolution of SOA and SOAG is similar between PartMC and MAM4 in this synthetic scenario, but that SOAG reaches a higher maximum in PartMC (Figure 5.2). When we examine the difference between PartMC and MAM4, we see that the error, i.e., the difference, peaks in magnitude at 190 min (Figure 5.3). However, the error in the GMD and GSD grow in magnitude monotonically throughout the simulation (Figure 5.4). If we examine the SOAG and SOA error over the range of GMD and GSD errors, we see the SOA and SOAG errors peak in magnitude at intermediate GMD and GSD errors (Figure 5.5).

Comparing the size distributions at the timestep of the maximum error and the final timestep, we see that the difference in MAM4 and PartMC volume concentrations at the timestep corresponding to the maximum error is smaller than that at the final timestep (Figure 5.6). When we examine the evolution of the volume distributions, we see that the total system volume predicted by both models is very similar up to time $t = 190$ min, when the volume concentrations begin to diverge (Figure 5.7). Together, these observations suggest compensating error is responsible for the SOA and SOAG errors decreasing after $t = 190$ min. Although the total system volume up to $t = 190$ min is similar, the MAM4 volume distribution is shifted towards larger particles. This enhances SOAG condensation, leading to the smaller maximum in SOAG since it is condensing more quickly. After $t = 190$ min, MAM4 predicts less overall volume concentration, but the volume distribution remains larger than that of PartMC, so MAM4 condenses SOAG more quickly. Thus, the enhanced condensation rate compensates for the smaller system volume concentration, and the differences in SOA and SOAG concentrations decrease.

5.4 Conclusions

We have analyzed of the disagreement in SOA condensation stemming from different particle size distribution representations, namely the modal MAM4 model vs. the particle-resolved

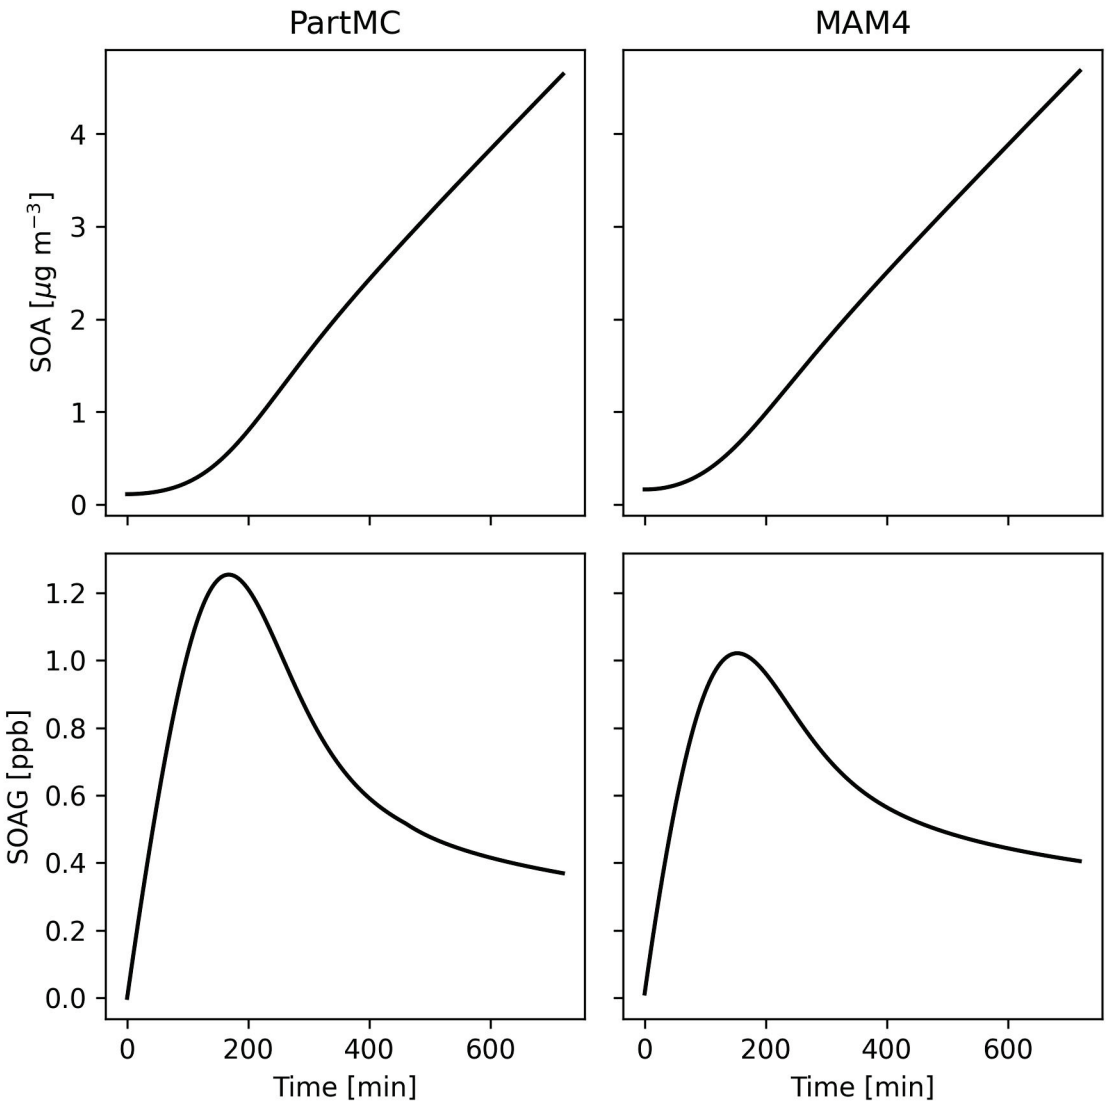

Figure 5.2: Time evolution of PartMC and MAM4 SOA mass concentration and SOAG mixing ratio for simulated SOA condensation.

PartMC model. We found that MAM4 ultimately produces less particle volume over the duration of the simulation, but the modal population grows larger in size to induce compensating error through enhanced condensation. Ultimately, this compensating error leads the system mass concentration to converge towards greater agreement as the simulation progresses. This analysis is enabled by our contributions to the AMBRS project. Our Python module that automates CAMP configuration ensures that PartMC-CAMP behaves similarly

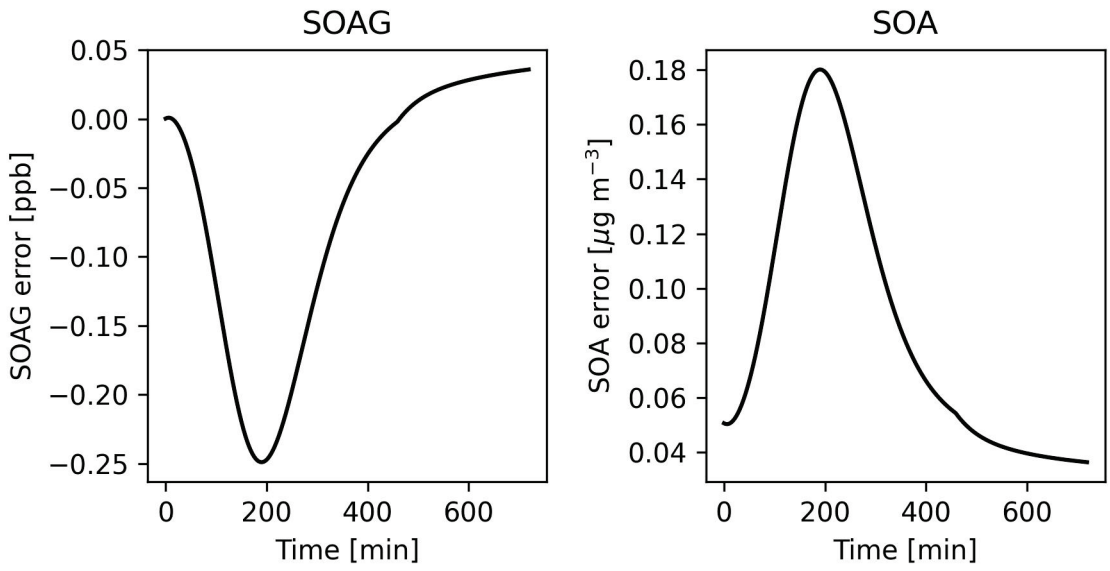

Figure 5.3: Time evolution of difference between PartMC and MAM4 (MAM4 - PartMC) SOA mass concentration and SOAG mixing ratio for simulated SOA condensation.

to MAM4 in this scenario, while MAM4-CAMP enables matching emissions in MAM4 to PartMC-CAMP. This analysis highlights the impact of assumptions in reduced models like MAM4. MAM4 assumes that the GSD of modes is fixed in time while GSD is free to evolve in PartMC. Indeed, PartMC GSD is computed *a posteriori* since the model imposes no assumptions or parameterizations on the size distribution by virtue of its explicit particle-resolved representation. We see that fixing the GSD in MAM4 has a small effect on bulk metrics like the system mass concentration of SOA but also that the agreement is the result of compensating errors that arise due to MAM4 approximations.

The impact of other assumptions, like the aggressive lumping of condensible organic gases into a single species, can now be interrogated by application of MAM4-CAMP. The interface can expand MAM4's gas phase chemistry treatment to break out the condensible organic gases and explore the differences in SOA condensation that may arise between a single condensing gas and a range of gases with varying volatilities. Future research directions include this exploration, combined with an expanded analysis of the modal representation, to characterize the effects of chemical representations and size parameterizations on model performance in terms of fidelity and efficiency.

With analyses like these, AMBRS, and our contributions thereto, can motivate and direct model development that improves aerosol predictions. Moreover, as AMBRS matures and adds more aerosol models to its repository, its benchmark cases can be used to identify the strengths and weaknesses different aerosol models have for different use cases, such as

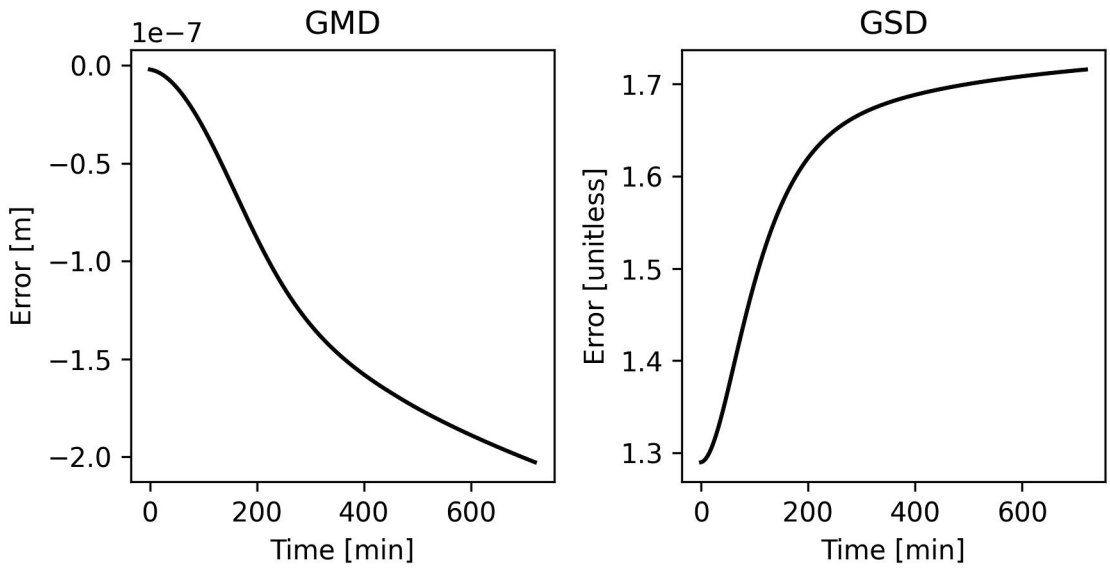

Figure 5.4: Time evolution of GMD and GSD differences (MAM4 - PartMC) for simulated SOA condensation.

predicting air quality metrics like $PM_{2.5}$ or predicting climate metrics like CCN. This can inform, for instance, the model configuration chosen for a regulatory demonstration simulation. With improved predictions from model development and informed model configuration, our contributions to the AMBRS Project can support improved regulatory policy design in both air quality and climate, since aerosol plays important roles in both.

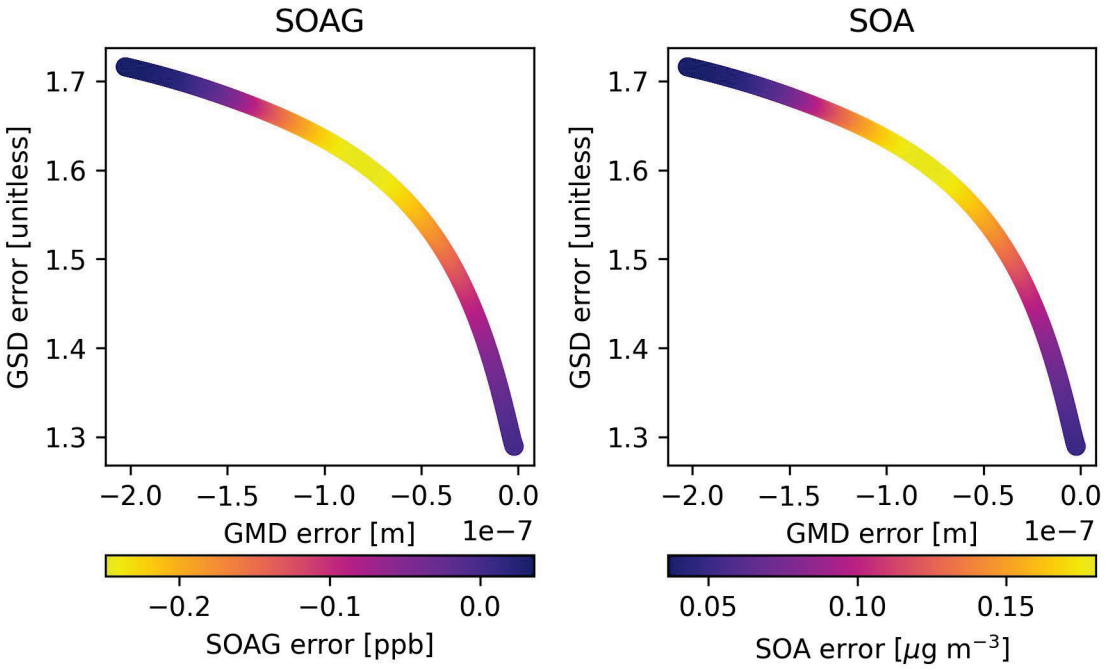

Figure 5.5: SOA and SOAG differences overlaid on GMD and GSD differences (MAM4 - PartMC) for simulated SOA condensation.

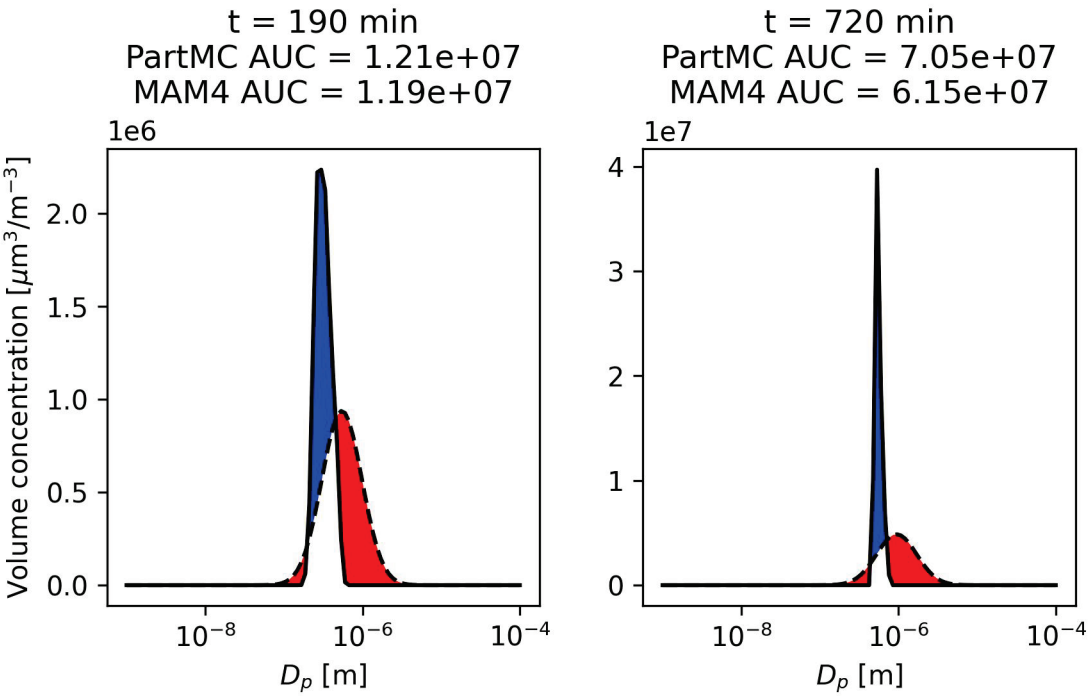

Figure 5.6: Volume distributions for MAM4 and PartMC at two timesteps. Blue shading indicates MAM4 is less than PartMC while red shading indicates the inverse. AUC is area under the curve.

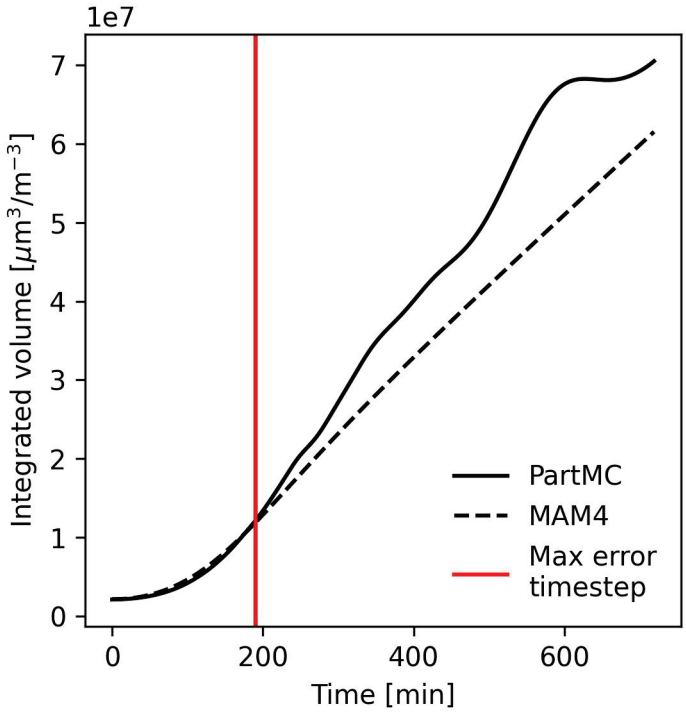

Figure 5.7: Evolution of integrated volume for MAM4 and PartMC.

Chapter 6

Conclusion

6.1 Summary

The work we collect and present throughout Chapters 2, 3, 4, and 5 explores a variety of modeling approaches and their potential to inform improved regulatory policy that could reduce exposure burdens across disparately affected communities, safeguarding the health of people everywhere. From measurement data-driven statistical modeling techniques to model acceleration and benchmarking, this body of work spans a broad set of methodologies, emphasizing the importance of varied but complementary approaches to addressing the persistent and pressing challenges posed by air pollution.

Measurement data elucidates the composition of air that communities breathe. Monitoring stations, like those used in Chapter 2, track air composition at their locations to provide a steady stream of data to quantify air quality and evaluate the effects of regulatory policies over long time horizons. Field campaigns, like those discussed in Chapter 3, while shorter in duration than permanent monitoring stations, complement monitors by producing rich data that can resolve more detail than many monitoring instruments. Chemical transport models like CMAQ offer a rigorous means of projecting air quality under different emissions scenarios, facilitating the *a priori* exploration of policy impacts on air quality, but are limited by computational resources. Powerful new hardware technology like GPUs, as in Chapter 4, can alleviate those computational limitations to simulate with higher fidelity, achieved by increasing chemical mechanism complexity, and higher grid resolution, achieved by decreasing the horizontal extent of grid cells. Simulation fidelity in terms of aerosol prediction accuracy can also be enhanced by benchmarking aerosol models, as explored in Chapter 5, to inform model development and simulation configurations that maximize prediction accuracy for the specific use case of interest.

Chapter 2, titled **"Analyzing Particulate Matter in California"**, characterizes the behavior of $PM_{2.5}$ and its major speciated components across a number of regulatory monitoring sites in California. We quantified the influence on $PM_{2.5}$ of a range of predictor covariates including meteorology and gas-phase precursor variables. GAMs predicted 24-hour average

$PM_{2.5}$ as a function of temperature, relative humidity, wind speed, wind direction, and solar radiation as meteorology predictors, along with NO_X HCHO, and O_3 as gas-phase predictors of $PM_{2.5}$. These covariates also predicted the NO_3^-, SO_4^{2-}, NH_4^+, OC, and EC components of $PM_{2.5}$. Our models predicted ambient $PM_{2.5}$ or one of its speciated components in Bakersfield, Fresno, Los Angeles, Riverside, Sacramento, and San Jose, with sensitivity analysis included alongside qualitative interpretations of marginal effects for Bakersfield, Riverside, and San Jose. The sensitivity analysis facilitated variable importance rankings for each model site to identify the most important predictor covariates. This ranking informed policy recommendations to reduce high $PM_{2.5}$ levels at these sites, which are to prioritize HCHO controls over NO_X controls in Bakersfield for more effective $PM_{2.5}$ reductions; to control both HCHO and NO_X in Riverside for effective $PM_{2.5}$ reductions; and the same for San Jose, although stronger controls may be necessary there for similar $PM_{2.5}$ reductions.

Chapter 3, titled **"Characterizing Organic Aerosol in California"**, extended the techniques of Chapter 2 to predict hourly average OA for field campaigns in the California cities of Riverside, Wilmington, and Bakersfield. GAMs predicted ambient OA as a function of speciated NR-PM_1 components including NO_3^-, SO_4^{2-}, NH_4^+, and Cl^-, alongside gas-phase covariates including NO_X, CO, and O_3 as well as meteorology covariates including temperature, relative humidity, wind speed and direction, and also the mass spectrometer m/z fractions f_{43} and f_{44}. Analogous to Chapter 2, we quantified the OA response to these predictor covariates and qualitatively interpret them at the study sites, as well as conducted sensitivity analysis to rank variables by importance to make the following policy recommendations. To most effectively control OA in Riverside, mobile emissions controls are necessary alongside controls on upwind sulfur emissions. In Wilmington, mobile emissions controls may have the co-benefits of alleviating high OA and high O_3. In Bakersfield, mobile emissions continue to be important, but may need to be even stronger to offset the impacts of biogenic emissions.

Chapter 4, titled **"Accelerating a CMAQ Chemistry Solver"**, implemented a GPU port of the Rosenbrock-3 solver in the Community Multiscale Air Quality model. We demonstrated that this implementation accelerates the chemistry solution without compromising the simulation. We showed up to 3× faster per-timestep chemistry solutions, which translates to a 1.6× faster full simulation walltime, saving significant compute resources with only one of CMAQ's science process modules accelerated. We reported degraded agreement between our accelerated CMAQ-CUDA implementation and the standard CMAQv5.4 implementation at high vertical layers and attributed this disagreement to roundoff error magnification, since there is an anticorrelation between the error and the concentration. Ultimately, our work in this chapter demonstrated the viability of GPU computation in a heterogeneous CTM implementation and laid the groundwork for hardware-accelerated first-principles atmospheric chemistry simulation in chemical transport models like CMAQ. We described the policy implications of this work, including the benefits of higher fidelity and higher resolution simulations.

Chapter 5, titled **"Developing the AMBRS Project Platform"**, presented our contributions to the Aerosol Model Benchmark Repository and Standards Project. We described the AMBRS platform and detailed our development of a Python module to automate con-

figuration of the CAMP aerosol chemistry module as well as our development of an interface between CAMP and MAM4. The Python configuration module facilitated configuring the PartMC benchmark model for like-to-like comparisons between PartMC and test models like MAM4, which enabled direct benchmarking and inter-comparison. Our MAM4-CAMP coupling interface facilitated complicating simplifications in MAM4's representation of gas-phase species, like its aggressively lumped condensible organic gas species, SOAG. This approach advanced towards structural error quantification, which can inform model development to constrain inter-model variability in Earth system models from their aerosol modules. We discussed the benefits of this type of analysis for improving air quality predictions and climate projections, which can improve policy for both fields.

6.2 Future Research Directions

Building on the work presented in this dissertation, future research directions include updating CMAQ-CUDA with improved scalability and a systematic interrogation of the simplifications made in MAM4.

CMAQ-CUDA, in the form presented in Chapter 4, is a first implementation of hardware-accelerated CMAQ chemistry. Chapter 4 presents performance demonstrations on a small test domain and with a modest runtime configuration that is not representative of production-scale runtime configurations. To maintain performance on production-scale runtime configurations, CMAQ-CUDA needs better GPU resource sharing. GPU sharing can be improved by executing simulations with the NVIDIA Multi Process Service (MPS) daemon active at runtime. Additional improvements may be obtained by restructuring the CMAQ-CUDA implementation of ROS3 to split the monolithic Fortran module into distinct files that import variables from other parts of CMAQ. This would also improve resource sharing by reducing variable redundancy, saving GPU memory. Moving to a unified compiler suite that relies solely on the NVIDIA HPC SDK (NVHPC) rather than a mixture of NVHPC and Intel or GNU compilers is required here, which will have the additional benefit of streamlining the compilation process.

To explore and quantify the impacts of simplifications in MAM4, like fixing the geometric standard deviation of modes or lumping all condensible organic gases into one model species with one volatility, we can develop and apply the MAM4-CAMP coupling interface presented in Chapter 5. Currently the interface only supports gas-phase chemistry. To investigate the impacts of breaking SOAG into a range of condensible organics with varying volatilities, support for gas-particle partitioning reactions like SIMPOL.1 must be implemented. Then MAM4-CAMP can be configured with any number of condensible organic gases, with CAMP storing and operating on the gas-phase species that condense into particulate SOA. PartMC can be configured to match and MAM4-CAMP can be benchmarked. With configurations for MAM4 and PartMC including standard MAM4 for both models, standard MAM4 with split-SOAG PartMC, and split-SOAG in both models, we can build up an error analysis that quantifies structural error, with PartMC configurations that mimic MAM4 controlling for

parametric error.

These directions will improve CMAQ-CUDA, readying it for deployment at production scales like large-domain, long-duration simulations and operational forecasting, and demonstrate the power of rigorous benchmarking to the aerosol modeling community.

6.3 Outlook

Outside of our contributions to the field, related efforts to improve and accelerate models are underway. The Multi-Scale Infrastructure for Chemistry and Aerosols (MUSICA), under development at the time of writing, seeks to unify the computational infrastructure under the hood of a chemical transport model and improve model resolutions with variable resolution grids (Tang et al. 2023). MUSICA development is also exploring GPU-enabled chemistry with the Model Independent Chemistry Module (MICM), an evolution of CAMP, that is also seeing integration into the National Oceanic and Atmospheric Administration (NOAA)'s Unified Forecasting System (UFS). Another GPU-enabled chemistry solution is in development by Díaz-Ibarra et al. (2026): TChem-atm. This chemistry solution stands out with its support for multiphase chemistry, its hardware agnosticism, and its host model flexibility, features shared by MICM.

It is an exciting time for atmospheric chemistry model development. In short order, we may have unified, hardware-accelerated atmosphere models with integrated multiphase chemistry support, improving model predictions, projections, and forecasts. Ultimately, developments like these will inform more efficacious air quality policy so that the challenges that persist can finally be overcome.

Appendix: Supplementary Material for Chapter 4

CMAQ-CUDA Code Progression

The implementation of CMAQ-CUDA has been an iterative process. The implementation we present in our main manuscript is version 3.0. Compared to earlier versions, CMAQ-CUDAv3.0 attains significant speedup by optimizing memory transfers and accommodating MPI-enabled and multi-GPU system configurations. Our first implementation of CMAQ-CUDA, version 1.0, migrated the subroutines called by RBSOLVER to CUDA Fortran (Do et al. 2023). Figure 6.1 diagrams this structure.

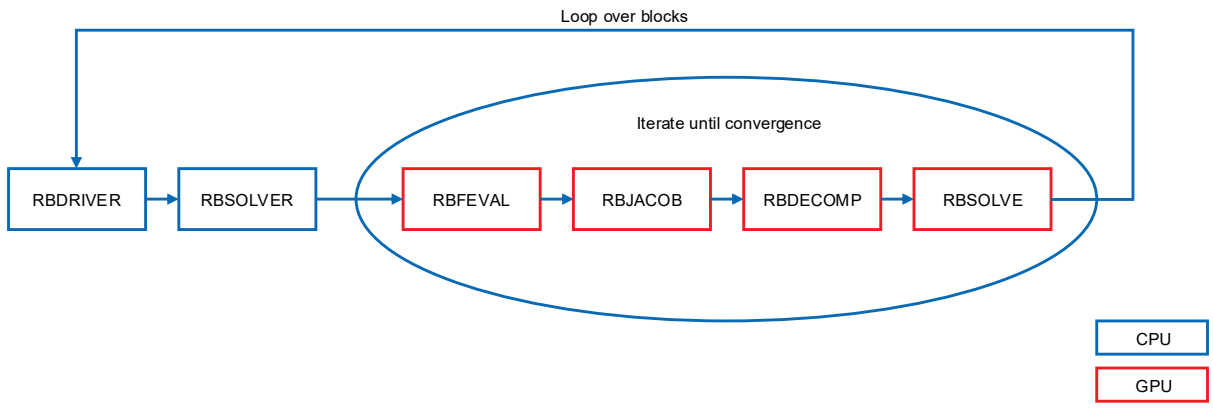

Figure 6.1: Schematic of CMAQ-CUDAv1.0 ROS3 code structure. Arrows represent data transfer. Each routine called by RBSOLVER is on the device in separate code, inducing host-device data transfer between calls.

This migration involved allocating copies of the subroutines' arrays on GPU device memory and rewriting the spatial loops as follows:

```
DO J = 1 TO NUMBER OF SPECIES
    IF I <= NUMBER OF CELLS
```

```
        UPDATE A( I,J )
    END IF
END DO
```

Notice that the inner loop is now an `IF` statement rather than a `DO` statement. This reflects the fact that this statement is executed in parallel, rather than in the traditional serial loop. The index `I` now represents a GPU thread index, with each thread uniquely indexed and executing its own copy of a kernel in parallel.

However, because the GPU needs to work with arrays stored in device memory, and because each kernel is its own separate code, CMAQ-CUDAv1.0 induces significant data transfer overhead. Every iteration until convergence, arrays are moved from CPU to GPU and back. This level of data transfer overhead results in slow CHEM module execution. To reduce data transfers, we implemented version 2.0 of CMAQ-CUDA, diagrammed in Figure 6.2. This implementation collects the four device routines into one kernel to eliminate the data transfer between kernels encountered in CMAQ-CUDAv1.0. However, the convergence loop remains on the CPU, so, while CMAQ-CUDAv2.0 reduces data transfer per iteration, it does not eliminate it. While an appreciable speedup over CMAQ-CUDAv1.0 is attained, the total CHEM module execution time is only on par with the standard CPU implementation.

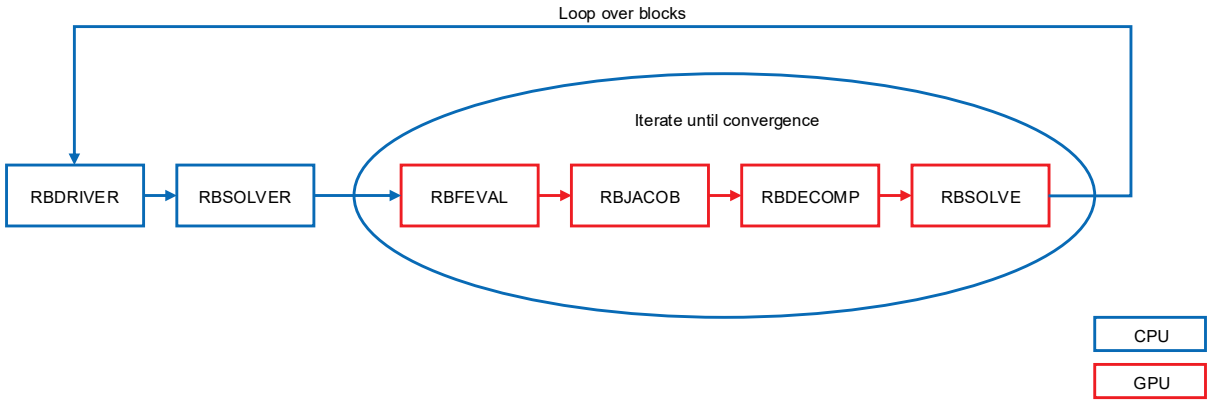

Figure 6.2: Schematic of CMAQ-CUDAv2.0 ROS3 code structure. Arrows represent data transfer. Data transfer is now limited to the start and end of a convergence loop.

To attain the desired speedup, the next version of CMAQ-CUDA, version 3.0, would need to further optimize data transfer overhead by reducing it as much as possible. CMAQ-CUDAv3.0 accomplishes this by moving the convergence loop onto the device. CMAQ-CUDAv3.0 is the version we present and diagram (Figure 4.5) in the main manuscript.

Determination of CMAQ's Block Size Compile-Time Parameter

In an effort to maximize device occupancy during kernel execution, we tailor CMAQ's block size parameter to GPU hardware and runtime specifications which are defined as

$$GRID = MAXBLOCKS \times MPCOUNT \times NDEV/NPROCS \tag{6.1}$$

$$TBLOCK = MAXTHREADS/MAXBLOCKS \tag{6.2}$$

where $GRID$ is the number of GPU thread blocks, $MAXBLOCKS$ is the maximum number of resident blocks per streaming multiprocessor, $MPCOUNT$ is the number of streaming multiprocessors on the GPU, $NDEV$ is the number of graphics cards, and $NPROCS$ is the number of MPI processes. In Equation 6.2, $TBLOCK$ is the number of GPU threads in a thread block and $MAXTHREADS$ is the maximum number of resident threads per streaming multiprocessor.

For single-precision operations, a kernel launches with $GRID \times TBLOCK$ total threads initialized. Therefore, $MAXTHREADS \times MPCOUNT \times NDEV/NPROCS$ total threads should be initialized for our kernel, where $MAXTHREADS \times MPCOUNT$ is the maximum number of resident threads on a GPU with $MPCOUNT$ streaming multiprocessors. Since we launch with $NDEV$ GPUs and $NPROCS$ MPI processes, we scale by their ratio to account for the total number of streaming multiprocessors on our multi-GPU node and how many MPI processes are using any one GPU.

Our simulations are executed on NVIDIA RTX 2080 Ti graphics cards, which each have 68 streaming multiprocessors that can each support up to 1024 threads across 16 thread blocks (*CUDA C++ Programming Guide* 2024). Our runtime configuration uses 4 cards and 8 MPI processes, so our theoretical occupancy-maximized block size parameter (not GPU thread block) is $GRID \times TBLOCK = (16 \times 68 \times 4/8) \times (1024/16) = 34,816$.

In Figure 4.8 of the main manuscript, however, we see that CMAQ-CUDA's optimal performance occurs at block size 17,408, or half of the theoretical occupancy-maximized block size. This is because the kernel computes with double-precision floating point numbers, which require twice as many registers per thread as single-precision numbers. This halves the thread count the compiler allocates for launching the kernel, resulting in the value we observe. Alvanos and Christoudias (2017) note this same effect in their work.

Varying MPI Process Count

We note that Alvanos and Christoudias (2017) highlight in their work that the factor by which GPUs accelerate EMAC diminishes as the application is launched with more MPI processes. Sun et al. (2018) report the same phenomenon in their work with CAM4-Chem. This effect is because simulation walltime on CPUs scales inversely with the number of MPI processes used to execute the application due to increasing parallelization, closing the

gap between CPU and GPU runtime. To investigate this effect in our case, we run a set of simulations while varying the number of MPI processes we use to launch CMAQ. The simulations are for the same 2018 12NE3 benchmark scenario as our main results and use the CB6R5 mechanism with aeo7 SOA treatment and standard cloud chemistry. We report our results in Figure 6.3.

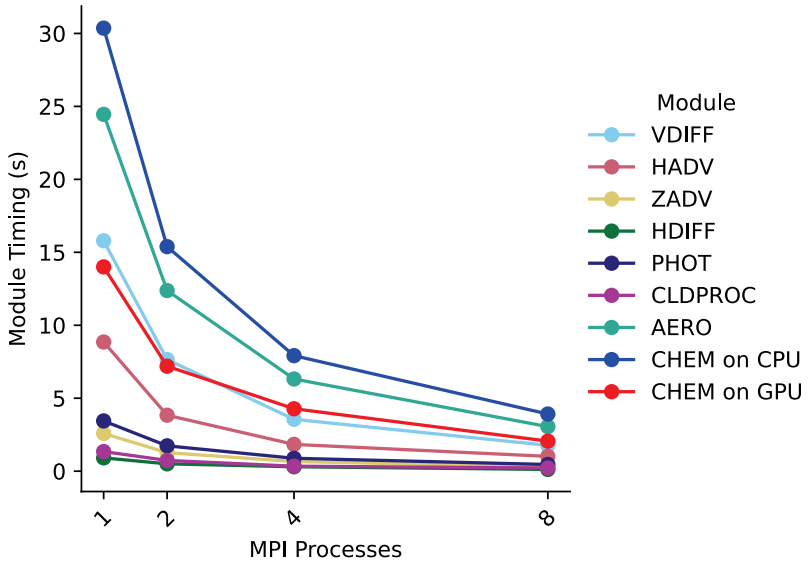

Figure 6.3: Timing of CMAQ science process modules as MPI process count varies. Each CMAQ-CUDA runtime configuration uses half as many graphics cards as MPI processes, rounded up.

We note the same effect in our case, where the difference in timing between CHEM on CPUs and CHEM on GPUs diminishes as we increase MPI process count. This may suggest that good candidate modules for GPU acceleration are those which do not scale well with process count. We note that the slowdown in CMAQ-CUDA module timing at low MPI process counts is due to the domain decomposition producing larger subdomains per process, requiring more block loops to cover the subdomain.

Investigating the Influence of Modules Other than CHEM

To isolate the disagreement between CMAQv5.4 and CMAQ-CUDA due only to our GPU-accelerated ROS3 implementation and not to interactions with other modules, we run simulations for both CMAQv5.4 and CMAQ-CUDA each with two cases: one where CHEM is

the only active science process module and one where CHEM is the only inactive science process module. Our simulation scenario is the same as before, again using CB6R5 with aero7 SOA treatment and standard cloud chemistry. We first investigate agreement in the surface layer in Figure 6.4.

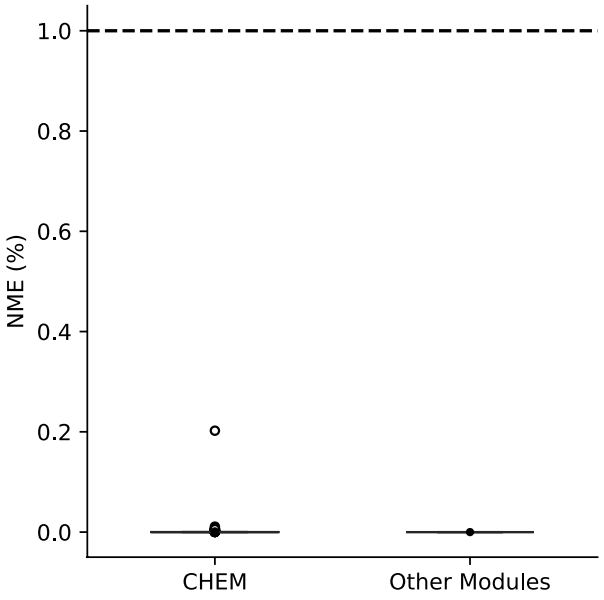

Figure 6.4: NME for the surface layer between simulations with only CHEM and simulations without CHEM for CB6R5. Each member of the distributions is a different species' NME averaged in space and time. Our 1% threshold is marked.

We see that our agreement in the CHEM-only case is significantly better than in the full simulation presented in Figure 4.11. All species are well below our threshold for acceptable difference, with the vast majority orders of magnitude lower. Because CMAQ-CUDA with CHEM inactive executes no GPU code, NME in that case is identically zero for all species.

When we examine more than the surface layer, in Figure 6.5 we again begin to see differences grow aloft. This is most probably due to roundoff error disproportionately affecting trace species at high altitudes, as in the case of the full simulation. However, unlike in the case of the full simulation results presented in Figure 4.12, in this CHEM-only case all species in all layers remain below our threshold. This supports our hypothesis of error magnification in other science process modules and Figure 4.10 suggests that the error being magnified may be roundoff error. Because roundoff errors occur differently in different architectures, such as CPUs versus GPUs, these differences are expected in CMAQ-CUDA. The fact that NME remains low throughout the whole three-dimensional domain when CHEM is the only active module supports confidence in the accuracy of CMAQ-CUDA.

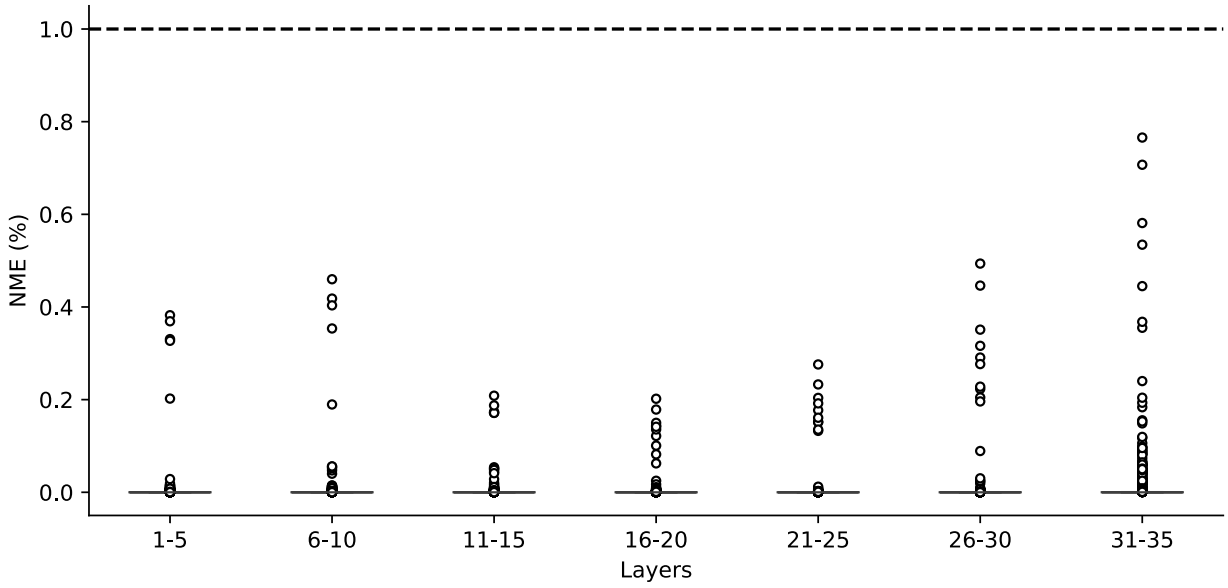

Figure 6.5: NME for binned vertical layers in the CHEM-only CB6R5 simulation. Each member of the distributions is a different species' NME averaged in space and time for the corresponding layer, such that each box plot has 5 members for each species: one for each layer in the corresponding bin. Our 1% threshold is marked.

References

40 CFR 58.12 – Operating schedules. (2022). URL: https://www.ecfr.gov/current/title-40/chapter-I/subchapter-C/part-58/subpart-B/section-58.12 (visited on 08/16/2022).

Ackermann-Liebrich, U et al. (Jan. 1997). "Lung function and long term exposure to air pollutants in Switzerland. Study on Air Pollution and Lung Diseases in Adults (SAPALDIA) Team." In: *American Journal of Respiratory and Critical Care Medicine* 155.1, pp. 122–129. ISSN: 1073-449X, 1535-4970. DOI: 10.1164/ajrccm.155.1.9001300.

Air Quality Data (PST) Query Tool (2022). URL: https://www.arb.ca.gov/aqmis2/aqdselect.php (visited on 08/16/2022).

Aldrin, M and I Haff (Apr. 2005). "Generalised additive modelling of air pollution, traffic volume and meteorology". In: *Atmospheric Environment* 39.11, pp. 2145–2155. ISSN: 13522310. DOI: 10.1016/j.atmosenv.2004.12.020.

Alduchov, O.A. and R.E. Eskridge (Nov. 1, 1997). *Improved Magnus' form approximation of saturation vapor pressure.* DOE/ER/61011–T6, 548871, DOE/ER/61011–T6, 548871. DOI: 10.2172/548871.

Alvanos, Michail and Theodoros Christoudias (Oct. 10, 2017). "GPU-accelerated atmospheric chemical kinetics in the ECHAM/MESSy (EMAC) Earth system model (version 2.52)". In: *Geoscientific Model Development* 10.10, pp. 3679–3693. ISSN: 1991-9603. DOI: 10.5194/gmd-10-3679-2017.

ARB Almanac 2013 - Chapter 4: Regional Trends and Forecasts (2023).

Atkinson, Richard W. et al. (Nov. 15, 2001). "Acute Effects of Particulate Air Pollution on Respiratory Admissions: Results from APHEA 2 Project". In: *American Journal of Respiratory and Critical Care Medicine* 164.10, pp. 1860–1866. ISSN: 1073-449X, 1535-4970. DOI: 10.1164/ajrccm.164.10.2010138.

Bannan, Thomas J. et al. (2019). "A Large Source of Atomic Chlorine From ClNO2 Photolysis at a U.K. Landfill Site". In: *Geophysical Research Letters* 46.14, pp. 8508–8516. ISSN: 1944-8007. DOI: 10.1029/2019GL083764.

Bell, M L and D L Davis (June 2001). "Reassessment of the lethal London fog of 1952: novel indicators of acute and chronic consequences of acute exposure to air pollution." In: *Environmental Health Perspectives* 109.Suppl 3, pp. 389–394. ISSN: 0091-6765. DOI: 10.1289/ehp.01109s3389.

Blanchard, Charles L. et al. (July 2019). "Emission influences on air pollutant concentrations in New York State: I. ozone". In: *Atmospheric Environment: X* 3, p. 100033. ISSN: 25901621. DOI: 10.1016/j.aeaoa.2019.100033.

Bolden, Ashley L., Carol F. Kwiatkowski, and Theo Colborn (May 2015). "New Look at BTEX: Are Ambient Levels a Problem?" In: *Environmental Science & Technology* 49.9, pp. 5261–5276. ISSN: 0013-936X. DOI: 10.1021/es505316f.

Bullard, Robert D. (1996). "Environmental Justice: It's More Than Waste Facility Siting". In: *Social Science Quarterly* 77.3, pp. 493–499. ISSN: 0038-4941.

Byun, Daewon and Kenneth L. Schere (Mar. 1, 2006). "Review of the Governing Equations, Computational Algorithms, and Other Components of the Models-3 Community Multiscale Air Quality (CMAQ) Modeling System". In: *Applied Mechanics Reviews* 59.2, pp. 51–77. ISSN: 0003-6900, 2379-0407. DOI: 10.1115/1.2128636.

Cai, Xuyi and Robert J. Griffin (2003). "Modeling the formation of secondary organic aerosol in coastal areas: Role of the sea-salt aerosol organic layer". In: *Journal of Geophysical Research: Atmospheres* 108 (D15). ISSN: 2156-2202. DOI: 10.1029/2002JD003053.

Calvin, Katherine et al. (July 25, 2023). *IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. IPCC, Geneva, Switzerland.* Intergovernmental Panel on Climate Change (IPCC). DOI: 10.59327/IPCC/AR6-9789291691647.

Cao, Kai et al. (Aug. 1, 2023). "GPU-HADVPPM V1.0: a high-efficiency parallel GPU design of the piecewise parabolic method (PPM) for horizontal advection in an air quality model (CAMx V6.10)". In: *Geoscientific Model Development* 16.15, pp. 4367–4383. ISSN: 1991-9603. DOI: 10.5194/gmd-16-4367-2023.

Carter, William (Sept. 9, 2023). *Documentation of the SAPRC-22 Mechanism.*

Carter, William P.L. (Dec. 2010). "Development of the SAPRC-07 chemical mechanism". In: *Atmospheric Environment* 44.40, pp. 5324–5335. ISSN: 13522310. DOI: 10.1016/j.atmosenv.2010.01.026.

Chen, Jianjun et al. (Aug. 2014). "Seasonal modeling of PM2.5 in California's San Joaquin Valley". In: *Atmospheric Environment* 92, pp. 182–190. ISSN: 13522310. DOI: 10.1016/j.atmosenv.2014.04.030.

Cheung, Kalam et al. (2023). *South Coast Air Basin Attainment plan for 2006 24-Hour PM2.5 Standard.*

CMAQv5.2 Operational Guidance Document (June 30, 2017).

CUDA C++ Programming Guide (2024). URL: https://docs.nvidia.com/cuda/cuda-c-programming-guide/ (visited on 05/10/2024).

Damian, Valeriu et al. (Nov. 2002). "The kinetic preprocessor KPP-a software environment for solving chemical kinetics". In: *Computers & Chemical Engineering* 26.11, pp. 1567–1579. ISSN: 00981354. DOI: 10.1016/S0098-1354(02)00128-X.

Davis, Dl (Dec. 2002). "A Look Back at the London Smog of 1952 and the Half Century Since". en. In: *Environmental Health Perspectives* 110.12. ISSN: 0091-6765, 1552-9924. DOI: 10.1289/ehp.110-a734.

Dawson, Matthew L. et al. (May 9, 2022). "Chemistry Across Multiple Phases (CAMP) version 1.0: an integrated multiphase chemistry model". In: *Geoscientific Model Development* 15.9, pp. 3663–3689. ISSN: 1991-959X. DOI: 10.5194/gmd-15-3663-2022.

DeVille, Lee, Nicole Riemer, and Matthew West (June 1, 2019). "Convergence of a generalized Weighted Flow Algorithm for stochastic particle coagulation". In: *Journal of Computational Dynamics* 6.1, pp. 69–94. ISSN: 2158-2491. DOI: 10.3934/jcd.2019003.

Díaz-Ibarra, Oscar H. et al. (Feb. 2026). "TChem-atm (v2.0.0): scalable performance-portable multiphase atmospheric chemistry". en. In: *Geoscientific Model Development* 19.3, pp. 1281–1299. ISSN: 1991-9603. DOI: 10.5194/gmd-19-1281-2026.

Do, Khanh et al. (Nov. 30, 2023). *Graphics Processing Unit Assisted Computation for a Gas-Phase Chemical Solver in a Regulatory Chemical Transport Model*. DOI: 10.31223/X5P400.

Dockery, Douglas W. et al. (Dec. 9, 1993). "An Association between Air Pollution and Mortality in Six U.S. Cities". In: *New England Journal of Medicine* 329.24, pp. 1753–1759. ISSN: 0028-4793, 1533-4406. DOI: 10.1056/NEJM199312093292401.

Enayati Ahangar, Faraz et al. (June 2021). "Long-term trends of PM2.5 and its carbon content in the South Coast Air Basin: A focus on the impact of wildfires". In: *Atmospheric Environment* 255, p. 118431. ISSN: 13522310. DOI: 10.1016/j.atmosenv.2021.118431.

EPA, US (Jan. 6, 2023). *EPA Proposes to Strengthen Air Quality Standards to Protect the Public from Harmful Effects of Soot*. URL: https://www.epa.gov/newsreleases/epa-proposes-strengthen-air-quality-standards-protect-public-harmful-effects-soot (visited on 04/21/2023).

Farmer, D. K. et al. (Apr. 13, 2010). "Response of an aerosol mass spectrometer to organonitrates and organosulfates and implications for atmospheric chemistry". In: *Proceedings of the National Academy of Sciences* 107.15, pp. 6670–6675. DOI: 10.1073/pnas.0912340107.

Faust, Jennifer A. (Feb. 2023). "PFAS on atmospheric aerosol particles: a review". en. In: *Environmental Science: Processes & Impacts* 25.2, pp. 133–150. ISSN: 2050-7895. DOI: 10.1039/D2EM00002D.

Fent, Kenneth W. et al. (May 2018). "Airborne contaminants during controlled residential fires". In: *Journal of Occupational and Environmental Hygiene* 15.5, pp. 399–412. ISSN: 1545-9624. DOI: 10.1080/15459624.2018.1445260.

Fenton, Suzanne E. et al. (2026). "Per- and Polyfluoroalkyl Substance Toxicity and Human Health Review: Current State of Knowledge and Strategies for Informing Future Research". en. In: ().

Ferretti, Federico, Andrea Saltelli, and Stefano Tarantola (Oct. 2016). "Trends in sensitivity analysis practice in the last decade". In: *Science of The Total Environment* 568, pp. 666–670. ISSN: 00489697. DOI: 10.1016/j.scitotenv.2016.02.133.

Fierce, Laura et al. (Sept. 1, 2024). "Quantifying structural errors in cloud condensation nuclei activity from reduced representation of aerosol size distributions". In: *Journal of Aerosol Science* 181, p. 106388. ISSN: 0021-8502. DOI: 10.1016/j.jaerosci.2024.106388.

Finlayson-Pitts, B. J. (Dec. 1, 2003). "The Tropospheric Chemistry of Sea Salt: A Molecular-Level View of the Chemistry of NaCl and NaBr". In: *Chemical Reviews* 103.12, pp. 4801–4822. ISSN: 0009-2665. DOI: 10.1021/cr020653t.

Gao, Song et al. (Dec. 1, 2004). "Particle Phase Acidity and Oligomer Formation in Secondary Organic Aerosol". In: *Environmental Science & Technology* 38.24, pp. 6582–6589. ISSN: 0013-936X. DOI: 10.1021/es049125k.

Gao, Ziqi et al. (Aug. 2022). "Separating emissions and meteorological impacts on peak ozone concentrations in Southern California using generalized additive modeling". In: *Environmental Pollution* 307, p. 119503. ISSN: 02697491. DOI: 10.1016/j.envpol.2022.119503.

Garg, Bhagwan D. et al. (Nov. 2000). "Brake Wear Particulate Matter Emissions". In: *Environmental Science & Technology* 34.21, pp. 4463–4469. ISSN: 0013-936X. DOI: 10.1021/es001108h.

Gentner, Drew R. et al. (Feb. 7, 2017). "Review of Urban Secondary Organic Aerosol Formation from Gasoline and Diesel Motor Vehicle Emissions". In: *Environmental Science & Technology* 51.3, pp. 1074–1093. ISSN: 0013-936X. DOI: 10.1021/acs.est.6b04509.

Golaz, Jean-Christophe et al. (Dec. 2022). "The DOE E3SM Model Version 2: Overview of the Physical Model and Initial Model Evaluation". en. In: *Journal of Advances in Modeling Earth Systems* 14.12, e2022MS003156. ISSN: 1942-2466, 1942-2466. DOI: 10.1029/2022MS003156.

Goliff, Wendy S., William R. Stockwell, and Charlene V. Lawson (Apr. 2013). "The regional atmospheric chemistry mechanism, version 2". In: *Atmospheric Environment* 68, pp. 174–185. ISSN: 13522310. DOI: 10.1016/j.atmosenv.2012.11.038.

Hafner, Hillary R and Bryan M Penfold (Jan. 4, 2018). *PAMS Data Validation and Analysis Training Material*.

Halonen, J I et al. (July 1, 2008). "Urban air pollution, and asthma and COPD hospital emergency room visits". In: *Thorax* 63.7, pp. 635–641. ISSN: 0040-6376. DOI: 10.1136/thx.2007.091371.

Hastie, Trevor and Robert Tibshirani (Aug. 1, 1986). "Generalized Additive Models". In: *Statistical Science* 1.3, pp. 297–310. ISSN: 0883-4237. DOI: 10.1214/ss/1177013604.

Hu, Shaohua et al. (June 1, 2009). "Metals emitted from heavy-duty diesel vehicles equipped with advanced PM and NOX emission controls". In: *Atmospheric Environment* 43.18, pp. 2950–2959. ISSN: 1352-2310. DOI: 10.1016/j.atmosenv.2009.02.052.

Ivey, Cesunica E. et al. (Jan. 31, 2022). *Ozone Meteorology Study Final Report*. Diamond Bar, CA: For South Coast Air Quality Management District, p. 249.

Jacobs, Elizabeth T., Jefferey L. Burgess, and Mark B. Abbott (Apr. 2018). "The Donora Smog Revisited: 70 Years After the Event That Inspired the Clean Air Act". In: *American Journal of Public Health* 108.S2, S85–S88. ISSN: 0090-0036. DOI: 10.2105/AJPH.2017.304219.

Jacobson, Mark Z. (2005). *Fundamentals of atmospheric modeling*. 2. ed. Cambridge: Cambridge Univ. Press. 813 pp. ISBN: 978-0-521-54865-6 978-0-521-83970-9.

Jaffe, Dan et al. (Aug. 2008). "Influence of Fires on O3 Concentrations in the Western U.S." In: *Environmental Science & Technology* 42.16, pp. 5885–5891. ISSN: 0013-936X. DOI: 10.1021/es800084k.

Jaffe, Daniel A. et al. (Oct. 2013). "Impact of Wildfires on Ozone Exceptional Events in the Western U.S." In: *Environmental Science & Technology* 47.19, pp. 11065–11072. ISSN: 0013-936X. DOI: 10.1021/es402164f.

Jang, Myoseon et al. (Oct. 25, 2002). "Heterogeneous Atmospheric Aerosol Production by Acid-Catalyzed Particle-Phase Reactions". In: *Science* 298.5594, pp. 814–817. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1075798.

Janssen, Nicole A.H. et al. (Dec. 2011). "Black Carbon as an Additional Indicator of the Adverse Health Effects of Airborne Particles Compared with PM10 and PM2.5". In: *Environmental Health Perspectives* 119.12, pp. 1691–1699. ISSN: 0091-6765. DOI: 10.1289/ehp.1003369.

Ji, Yi et al. (July 2025). "Assessing the Impact of Wildfire Smoke Transport Through Chemical Transport Modeling, Satellite Retrievals, and Ground-Based Observations of Ozone in Rural Nevada". In: *ACS ES&T Air* 2.7, pp. 1237–1248. DOI: 10.1021/acsestair.5c00037.

Kelly, Thomas J. et al. (Aug. 1994). "Concentrations and transformations of hazardous air pollutants". en. In: *Environmental Science & Technology* 28.8, 378A–387A. ISSN: 0013-936X, 1520-5851. DOI: 10.1021/es00057a003.

Kelp, Makoto M., Daniel J. Jacob, J. Nathan Kutz, et al. (Dec. 16, 2020). "Toward Stable, General Machine-Learned Models of the Atmospheric Chemical System". In: *Journal of Geophysical Research: Atmospheres* 125.23, e2020JD032759. ISSN: 2169-897X, 2169-8996. DOI: 10.1029/2020JD032759.

Kelp, Makoto M., Daniel J. Jacob, Haipeng Lin, et al. (June 2022). "An Online-Learned Neural Network Chemical Solver for Stable Long-Term Global Simulations of Atmospheric Chemistry". In: *Journal of Advances in Modeling Earth Systems* 14.6, e2021MS002926. ISSN: 1942-2466, 1942-2466. DOI: 10.1029/2021MS002926.

Kim, Yunhee, Joshua S. Fu, and Terry L. Miller (June 2010). "Improving ozone modeling in complex terrain at a fine grid resolution – Part II: Influence of schemes in MM5 on daily maximum 8-h ozone concentrations and RRFs (Relative Reduction Factors) for SIPs in the non-attainment areas". In: *Atmospheric Environment* 44.17, pp. 2116–2124. ISSN: 13522310. DOI: 10.1016/j.atmosenv.2010.02.038.

Koolik, Libby H. et al. (Sept. 2024). "PM2.5 exposure disparities persist despite strict vehicle emissions controls in California". In: *Science Advances* 10.37, eadn8544. DOI: 10.1126/sciadv.adn8544.

Lane, Haley M. et al. (Apr. 2022). "Historical Redlining Is Associated with Present-Day Air Pollution Disparities in U.S. Cities". In: *Environmental Science & Technology Letters* 9.4, pp. 345–350. DOI: 10.1021/acs.estlett.1c01012.

Laskin, Alexander et al. (2012). "Tropospheric chemistry of internally mixed sea salt and organic particles: Surprising reactivity of NaCl with weak organic acids". In: *Jour-*

nal of Geophysical Research: Atmospheres 117 (D15). ISSN: 2156-2202. DOI: 10.1029/2012JD017743.

Lei, Ziying et al. (Aug. 2, 2022). "Initial pH Governs Secondary Organic Aerosol Phase State and Morphology after Uptake of Isoprene Epoxydiols (IEPOX)". In: *Environmental Science & Technology* 56.15, pp. 10596–10607. ISSN: 0013-936X. DOI: 10.1021/acs.est.2c01579.

Li, Huiming, Xin Qian, and Qin'geng Wang (Dec. 2013). "Heavy Metals in Atmospheric Particulate Matter: A Comprehensive Understanding Is Needed for Monitoring and Risk Mitigation". In: *Environmental Science & Technology* 47.23, pp. 13210–13211. ISSN: 0013-936X. DOI: 10.1021/es404751a.

Lippmann, Morton (Dec. 1991). "Health effects of tropospheric ozone". en. In: *Environmental Science & Technology* 25.12, pp. 1954–1962. ISSN: 0013-936X, 1520-5851. DOI: 10.1021/es00024a001.

Liu, X., R. C. Easter, et al. (May 21, 2012). "Toward a minimal representation of aerosols in climate models: description and evaluation in the Community Atmosphere Model CAM5". In: *Geoscientific Model Development* 5.3, pp. 709–739. ISSN: 1991-959X. DOI: 10.5194/gmd-5-709-2012.

Liu, X., P.-L. Ma, et al. (Feb. 8, 2016). "Description and evaluation of a new four-mode version of the Modal Aerosol Module (MAM4) within version 5.3 of the Community Atmosphere Model". In: *Geoscientific Model Development* 9.2, pp. 505–522. ISSN: 1991-959X. DOI: 10.5194/gmd-9-505-2016.

Maps of State and Federal Area Designations — California Air Resources Board (2022). California Air Resources Board. URL: https://ww2.arb.ca.gov/resources/documents/maps-state-and-federal-area-designations (visited on 08/09/2022).

Marjollet, Arnaud et al. (Feb. 19, 2015). *Guidance for Assessing and Mitigating Air Quality Impacts*.

McGraw, R. et al. (July 2008). "Aerosol dynamics using the quadrature method of moments: comparing several quadrature schemes with particle-resolved simulation". In: *Journal of Physics: Conference Series* 125.1, p. 012020. ISSN: 1742-6596. DOI: 10.1088/1742-6596/125/1/012020.

Mohai, Paul, David Pellow, and J. Timmons Roberts (Nov. 2009). "Environmental Justice". en. In: *Annual Review of Environment and Resources* 34.Volume 34, 2009, pp. 405–430. ISSN: 1543-5938, 1545-2050. DOI: 10.1146/annurev-environ-082508-094348.

Mysliwiec, Mitchell J. and Michael J. Kleeman (Dec. 1, 2002). "Source Apportionment of Secondary Airborne Particulate Matter in a Polluted Atmosphere". In: *Environmental Science & Technology* 36.24, pp. 5376–5384. ISSN: 0013-936X. DOI: 10.1021/es020832s.

Nenes, Athanasios, Spyros N. Pandis, and Christodoulos Pilinis (Mar. 1, 1998). "ISORROPIA: A New Thermodynamic Equilibrium Model for Multiphase Multicomponent Inorganic Aerosols". In: *Aquatic Geochemistry* 4.1, pp. 123–152. ISSN: 1573-1421. DOI: 10.1023/A:1009604003981.

Nickolls, John and William J Dally (Mar. 2010). "The GPU Computing Era". In: *IEEE Micro* 30.2, pp. 56–69. ISSN: 0272-1732. DOI: 10.1109/MM.2010.41.

NOAA/ESRL/GSL - RAOB (2022). URL: https://ruc.noaa.gov/raobs/ (visited on 08/16/2022).

Nussbaumer, Clara M. and Ronald C. Cohen (Mar. 16, 2021). "Impact of OA on the Temperature Dependence of PM 2.5 in the Los Angeles Basin". In: *Environmental Science & Technology* 55.6, pp. 3549–3558. ISSN: 0013-936X, 1520-5851. DOI: 10.1021/acs.est.0c07144.

Pankow, J. F. and W. E. Asher (May 29, 2008). "SIMPOL.1: a simple group contribution method for predicting vapor pressures and enthalpies of vaporization of multifunctional organic compounds". In: *Atmospheric Chemistry and Physics* 8.10, pp. 2773–2796. ISSN: 1680-7316. DOI: 10.5194/acp-8-2773-2008.

Pearce, John L., Jason Beringer, Neville Nicholls, Rob J. Hyndman, and Nigel J. Tapper (Feb. 2011). "Quantifying the influence of local meteorology on air quality using generalized additive models". In: *Atmospheric Environment* 45.6, pp. 1328–1336. ISSN: 13522310. DOI: 10.1016/j.atmosenv.2010.11.051.

Pearce, John L., Jason Beringer, Neville Nicholls, Rob J. Hyndman, Petteri Uotila, et al. (Jan. 1, 2011). "Investigating the influence of synoptic-scale meteorology on air quality using self-organizing maps and generalized additive modelling". In: *Atmospheric Environment* 45.1, pp. 128–136. ISSN: 1352-2310. DOI: 10.1016/j.atmosenv.2010.09.032.

Perperoglou, Aris et al. (Dec. 2019). "A review of spline function procedures in R". In: *BMC Medical Research Methodology* 19.1, p. 46. ISSN: 1471-2288. DOI: 10.1186/s12874-019-0666-3.

Pope, C. Arden et al. (Mar. 1995). "Particulate Air Pollution as a Predictor of Mortality in a Prospective Study of U.S. Adults". In: *American Journal of Respiratory and Critical Care Medicine* 151.3, pp. 669–674. ISSN: 1073-449X, 1535-4970. DOI: 10.1164/ajrccm/151.3_Pt_1.669.

Puy, Arnald et al. (Dec. 3, 2021). *sensobol: an R package to compute variance-based sensitivity indices*. arXiv: 2101.10103[stat].

Pye, Havala O. T. et al. (May 4, 2023). "Linking gas, particulate, and toxic endpoints to air emissions in the Community Regional Atmospheric Chemistry Multiphase Mechanism (CRACMM)". In: *Atmospheric Chemistry and Physics* 23.9, pp. 5043–5099. ISSN: 1680-7324. DOI: 10.5194/acp-23-5043-2023.

Quevedo, Duncan, Khanh Do, et al. (Feb. 2025). "GPU Implementation of a Gas-Phase Chemistry Solver in the CMAQ Chemical Transport Model". In: *ACS ES&T Air* 2.2, pp. 226–235. DOI: 10.1021/acsestair.4c00181.

Quevedo, Duncan, Ziqi Gao, et al. (Jan. 2024). "Multidecadal Analysis of Meteorological and Emissions Regimes for PM2.5 Across California". In: *ACS ES&T Air* 1.1, pp. 33–42. DOI: 10.1021/acsestair.3c00019.

Quinete, Natalia et al. (Oct. 2014). "Occurrence and distribution of PCB metabolites in blood and their potential health effects in humans: a review". en. In: *Environmental Science and Pollution Research* 21.20, pp. 11951–11972. ISSN: 1614-7499. DOI: 10.1007/s11356-014-3136-9.

R: Generalized additive models with integrated smoothness... (2022). URL: https://stat.ethz.ch/R-manual/R-devel/library/mgcv/html/gam.html (visited on 08/17/2022).

Raizenne, M et al. (May 1996). "Health effects of acid aerosols on North American children: pulmonary function." In: *Environmental Health Perspectives* 104.5, pp. 506–514. ISSN: 0091-6765, 1552-9924. DOI: 10.1289/ehp.96104506.

Riemer, N. et al. (2009). "Simulating the evolution of soot mixing state with a particle-resolved aerosol model". In: *Journal of Geophysical Research: Atmospheres* 114 (D9). ISSN: 2156-2202. DOI: 10.1029/2008JD011073.

Riva, Matthieu et al. (June 7, 2016). "Effect of Organic Coatings, Humidity and Aerosol Acidity on Multiphase Chemistry of Isoprene Epoxydiols". In: *Environmental Science & Technology* 50.11, pp. 5580–5588. ISSN: 0013-936X. DOI: 10.1021/acs.est.5b06050.

Rollins, A. W. et al. (2013). "Gas/particle partitioning of total alkyl nitrates observed with TD-LIF in Bakersfield". In: *Journal of Geophysical Research: Atmospheres* 118.12, pp. 6651–6662. ISSN: 2169-8996. DOI: 10.1002/jgrd.50522.

Sandu, A. et al. (Oct. 1997). "Benchmarking stiff ode solvers for atmospheric chemistry problems II: Rosenbrock solvers". In: *Atmospheric Environment* 31.20, pp. 3459–3472. ISSN: 13522310. DOI: 10.1016/S1352-2310(97)83212-8.

Sarwar, Golam et al. (Jan. 1, 2008). "Impact of an Updated Carbon Bond Mechanism on Predictions from the CMAQ Modeling System: Preliminary Assessment". In: *Journal of Applied Meteorology and Climatology* 47.1, pp. 3–14. ISSN: 1558-8432, 1558-8424. DOI: 10.1175/2007JAMC1393.1.

Schlosser, Joseph S. et al. (2017). "Analysis of aerosol composition data for western United States wildfires between 2005 and 2015: Dust emissions, chloride depletion, and most enhanced aerosol constituents". In: *Journal of Geophysical Research: Atmospheres* 122.16, pp. 8951–8966. ISSN: 2169-8996. DOI: 10.1002/2017JD026547.

Schmedding, Ryan et al. (Sept. 15, 2019). "-Pinene-Derived organic coatings on acidic sulfate aerosol impacts secondary organic aerosol formation from isoprene in a box model". In: *Atmospheric Environment* 213, pp. 456–462. ISSN: 1352-2310. DOI: 10.1016/j.atmosenv.2019.06.005.

Seigneur, Christian et al. (Jan. 1986). "Simulation of Aerosol Dynamics: A Comparative Review of Mathematical Models". In: *Aerosol Science and Technology* 5.2, pp. 205–222. ISSN: 0278-6826, 1521-7388. DOI: 10.1080/02786828608959088.

Seinfeld, John H. and Spyros N. Pandis (2016). *Atmospheric chemistry and physics: from air pollution to climate change*. Third edition. Hoboken, New Jersey: John Wiley & Sons. 1 p. ISBN: 978-1-119-22116-6 978-1-119-22117-3.

Simon, Heather, Kirk R. Baker, and Sharon Phillips (Dec. 2012). "Compilation and interpretation of photochemical model performance statistics published between 2006 and 2012". In: *Atmospheric Environment* 61, pp. 124–139. ISSN: 13522310. DOI: 10.1016/j.atmosenv.2012.07.012.

Sun, Jian et al. (Aug. 2018). "Computational Benefit of GPU Optimization for the Atmospheric Chemistry Modeling". In: *Journal of Advances in Modeling Earth Systems* 10.8, pp. 1952–1969. ISSN: 1942-2466, 1942-2466. DOI: 10.1029/2018MS001276.

Tang, Wenfu et al. (2023). "Capturing High-Resolution Air Pollution Features Using the Multi-Scale Infrastructure for Chemistry and Aerosols Version 0 (MUSICAv0) Global Modeling System". en. In: *Journal of Geophysical Research: Atmospheres* 128.7, e2022JD038345. ISSN: 2169-8996. DOI: 10.1029/2022JD038345.

Team, R Core (2021). *R: A Language and Environment for Statistical Computing*. Vienna, Austria.

Tessum, Christopher W. et al. (Apr. 2021). "PM $_{2.5}$ polluters disproportionately and systemically affect people of color in the United States". en. In: *Science Advances* 7.18, eabf4491. ISSN: 2375-2548. DOI: 10.1126/sciadv.abf4491.

Tingey, David T. et al. (1979). "The Influence of Light and Temperature on Isoprene Emission Rates from Live Oak". In: *Physiologia Plantarum* 47.2, pp. 112–118. ISSN: 1399-3054. DOI: 10.1111/j.1399-3054.1979.tb03200.x.

Understanding Particulate Matter: Protecting Public Health in the San Francisco Bay Area (2012).

US EPA (2023). *CMAQ Model Version 5.4 Benchmark Input Data*. In collab. with Elizabeth Adams. URL: https://dataverse.unc.edu/citation?persistentId=doi:10.15139/S3/BWMI8X (visited on 10/14/2024).

US EPA, OAR (Feb. 2026). *Air Quality - National Summary*. en. Data and Tools.

Vogel, Annika et al. (2022). "Uncertainty in Aerosol Optical Depth From Modern Aerosol-Climate Models, Reanalyses, and Satellite Products". In: *Journal of Geophysical Research: Atmospheres* 127.2, e2021JD035483. ISSN: 2169-8996. DOI: 10.1029/2021JD035483.

Vutukuru, Satish, Robert J. Griffin, and Donald Dabdub (May 27, 2006). "Simulation and analysis of secondary organic aerosol dynamics in the South Coast Air Basin of California: SOA DYNAMICS IN SOUTHERN CALIFORNIA". In: *Journal of Geophysical Research: Atmospheres* 111 (D10), S12. ISSN: 01480227. DOI: 10.1029/2005JD006139.

Walker, David and Jack Dongarra (1996). "MPI: a standard message passing interface". In: *Supercomputer* 12, pp. 56–68.

Wang, Yuzhou et al. (Nov. 2022). "Location-specific strategies for eliminating US national racial-ethnic PM2.5 exposure inequality". en. In: *Proceedings of the National Academy of Sciences* 119.44, e2205548119. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.2205548119.

Wood, S. N. (May 2000). "Modelling and smoothing parameter estimation with multiple quadratic penalties". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 62.2, pp. 413–428. ISSN: 1369-7412, 1467-9868. DOI: 10.1111/1467-9868.00240.

Wood, Simon N. (Jan. 2011). "Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models: Estimation of Semiparametric Generalized Linear Models". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73.1, pp. 3–36. ISSN: 13697412. DOI: 10.1111/j.1467-9868.2010.00749.x.

Wood, Simon N., Natalya Pya, and Benjamin Säfken (Oct. 1, 2016). "Smoothing Parameter and Model Selection for General Smooth Models". In: *Journal of the American Statis-*

tical Association 111.516, pp. 1548–1563. ISSN: 0162-1459, 1537-274X. DOI: 10.1080/01621459.2016.1180986.

Woody, Matthew C. et al. (Mar. 29, 2016). "Understanding sources of organic aerosol during CalNex-2010 using the CMAQ-VBS". In: *Atmospheric Chemistry and Physics* 16.6, pp. 4081–4100. ISSN: 1680-7324. DOI: 10.5194/acp-16-4081-2016.

Yarwood, Greg (2010). "Yarwood, Greg, et al. "Updates to the Carbon Bond mechanism for version 6 (CB6)". Chapel Hill, NC.

Z. Jacobson, Mark and Richard P. Turco (Jan. 1994). "SMVGEAR: A sparse-matrix, vectorized gear code for atmospheric models". In: *Atmospheric Environment* 28.2, pp. 273–284. ISSN: 13522310. DOI: 10.1016/1352-2310(94)90102-3.

Zanobetti, A, J Schwartz, and D W Dockery (Nov. 2000). "Airborne particles are a risk factor for hospital admissions for heart and lung disease." In: *Environmental Health Perspectives* 108.11, pp. 1071–1077. ISSN: 0091-6765, 1552-9924. DOI: 10.1289/ehp.001081071.

Zaveri, Rahul A. et al. (2008). "Model for Simulating Aerosol Interactions and Chemistry (MOSAIC)". In: *Journal of Geophysical Research: Atmospheres* 113 (D13). ISSN: 2156-2202. DOI: 10.1029/2007JD008782.

Zhang, X. et al. (Feb. 14, 2014). "Role of ozone in SOA formation from alkane photooxidation". In: *Atmospheric Chemistry and Physics* 14.3, pp. 1733–1753. ISSN: 1680-7316. DOI: 10.5194/acp-14-1733-2014.

Zhang, Yue et al. (Dec. 19, 2019). "Joint Impacts of Acidity and Viscosity on the Formation of Secondary Organic Aerosol from Isoprene Epoxydiols (IEPOX) in Phase Separated Particles". In: *ACS Earth and Space Chemistry* 3.12, pp. 2646–2658. DOI: 10.1021/acsearthspacechem.9b00209.

Zhao, Yunliang et al. (Oct. 16, 2013). "Sources of organic aerosol investigated using organic compounds as tracers measured during CalNex in Bakersfield: SOURCES OF ORGANIC AEROSOL". In: *Journal of Geophysical Research: Atmospheres* 118.19, pp. 11, 388–11, 398. ISSN: 2169897X. DOI: 10.1002/jgrd.50825.

Zheng, Zhonghua et al. (Dec. 3, 2021). "Quantifying the structural uncertainty of the aerosol mixing state representation in a modal model". In: *Atmospheric Chemistry and Physics* 21.23, pp. 17727–17741. ISSN: 1680-7316. DOI: 10.5194/acp-21-17727-2021.

Zhu, Shupeng et al. (Apr. 2019). "Comprehensively assessing the drivers of future air quality in California". In: *Environment International* 125, pp. 386–398. ISSN: 01604120. DOI: 10.1016/j.envint.2019.02.007.