

Calculation of Sensitivity Coefficients for Individual Airport Emissions in the Continental U.S. using CMAQ-DDM/PM

Scott Boone
University of North Carolina at Chapel Hill
Institute for the Environment
stboone@live.unc.edu

Saravanan Arunachalam^{*}
University of North Carolina at Chapel Hill
Institute for the Environment
sarav@email.unc.edu

ABSTRACT

Fine particulate matter (PM_{2.5}) is a federally-regulated air pollutant with well-known impacts on human health. The FAA's *Destination 2025* program seeks to decrease aviation-related health impacts across the U.S. by 50% by the year 2018. Atmospheric models, such as the Community Multiscale Air Quality model (CMAQ), are used to estimate the atmospheric concentration of pollutants such as PM_{2.5}. Sensitivity analysis of these models has long been limited to finite difference and regression-based methods, both of which require many computationally intensive model simulations to link changes in output with perturbations in input. Further, they are unable to offer detailed or *ad hoc* analysis for changes within a domain, such as changes in emissions on an airport-by-airport basis. In order to calculate the sensitivity of PM_{2.5} concentrations to emissions from individual airports, we utilize the Decoupled Direct Method in three dimensions (DDM-3D), an advanced sensitivity analysis tool recently implemented in CMAQ. DDM-3D allows calculation of sensitivity coefficients within a single simulation, eliminating the need for multiple model runs. However, while the output provides results for a variety of input perturbations in a single simulation, the processing time for each run is dramatically increased compared to simulations conducted without the DDM-3D module.

Use of the XSEDE Stampede computing cluster allows us to calculate sensitivity coefficients for a large number of input parameters. This allows for a much wider variety of *ad hoc* aviation policy scenarios to be generated and evaluated than would be possible using other sensitivity analysis methods or smaller-scaled computing systems. We present a design of experiments to compute individual sensitivity coefficients for 139 major airports in the US, due to six different precursor emissions that form PM_{2.5} in the atmosphere. Simulations based on this design are currently in progress, with full results to be published at a later date.

^{*}Corresponding author

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

XSEDE '14 Jul 13–18 2014, Atlanta, GA, USA
ACM 978-1-4503-2893-7/14/07.
<http://dx.doi.org/10.1145/2616498.2616504>

1. BACKGROUND

1.1 Aviation

Aviation is a critical component of transportation infrastructure in the United States. Since 2001, the proportion of domestic miles traveled by air compared to other modes has increased steadily from 9.5% in 2001 to 11.8% in 2011[8]. In 2013, air carriers flew nearly 578 billion domestic passenger-miles, an increase of 17% over the past decade[9]. The Federal Aviation Administration (FAA) expects this rate of growth to continue, projecting year-over-year increases of 2.5% over the next 25 years[13].

The FAA's Aviation Emission Design Tool (AEDT)[27] contains over 2,000 active airports in the United States; of these, the busiest 200 airports represent 85% of total flights and 95% of total fuel consumption. From this group, major commercial hubs represent an even more disproportionate level of activity, with the top ten U.S. airports alone contributing nearly one-quarter of all trips taken and one-third of all fuel consumed. Major hub airports are generally located in or near U.S. cities; their high levels of activity, steady growth rate and proximity to population centers provide strong motivation for investigation into the health effects associated with airport operations. The FAA's *Destination 2025* vision framework sets out a series of goals to improve the aviation industry's level of safety and future sustainability. One of these goals is the reduction of aviation emission-related health impacts by 50% by the year 2018 compared to 2005 baseline levels[31].

1.2 Particulate Matter

Atmospheric particulate matter¹ is broadly defined as any nonwater substance that exists in liquid or solid form, is of microscopic size, and is suspended in the atmosphere. We may additionally discriminate between classes of particulates by their size and chemical composition. Fine particulate matter, or PM_{2.5}, is generally defined as the set of particulates with aerodynamic diameter 2.5 microns or smaller.

PM_{2.5} may be broken down by chemical composition into two categories: primary and secondary. Primary PM_{2.5} is the set of particulates emitted directly into the atmosphere, e.g. smoke from incomplete combustion, sea salt from ocean spray, dust from brake pads and suspended dust; secondary PM_{2.5} is formed in the atmosphere from the oxidation of precursor species such as SO₂ and NO_x.

¹Also known generally as aerosols, though this term properly refers collectively to both the particles and the gas in which they are suspended.

Atmospheric concentrations of particulate matter are federally regulated by the National Ambient Air Quality Standards (NAAQS)[30]. Annual concentrations (averaged over a three-year period) may not exceed $12 \mu\text{g}/\text{m}^3$, while daily concentrations (measured as the 98th percentile, averaged over three years) may not exceed $35 \mu\text{g}/\text{m}^3$.

Long-term exposure to fine particulate matter has been shown to be strongly associated with increased health impacts due to its ability to penetrate deep into the respiratory system. Pope and Dockery reported that for a $10\mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ concentration, mortality from all causes increased by ten percent, including a six percent increase in cardiopulmonary mortality and an eight percent increase in lung cancer mortality[25]. Other estimates are higher; an update of the Harvard Six Cities Study (one of the original studies linking particulates with adverse health impacts) estimates an increase in premature mortality between 13% and 16% for an equivalent increase in $\text{PM}_{2.5}$ concentrations[17].

Previous estimates of total aviation contributions of $\text{PM}_{2.5}$ have shown that aviation LTO activity increases annual average concentrations by about 0.05%. However, this number is expected to increase to 0.20% by the year 2025[33]. An analysis of health impacts due to aviation emissions from a set of 99 U.S. airports estimated an increase from 75 premature deaths per year in 2005 to approximately 460 in 2025. This growth factor of 6.1 includes a factor of 2.1 attributable to aircraft emissions growth, a factor of 2.3 attributable to changing nonaviation concentrations, and a factor of 1.3 attributable to population growth[20].

1.3 Air Quality Modeling

Particulate matter and other pollutants are modeled using one of several types of chemical transport models (CTMs). An Eulerian CTM divides the domain into three-dimensional grid cells. Each cell is assumed to be well-mixed; that is, each chemical species in the cell is of uniform concentration². A commonly-used Eulerian CTM is the Community Multiscale Air Quality modeling system, or CMAQ[10].

Most Eulerian CTMs function in the same manner. For each grid cell and at each time step, chemical concentrations for each species are governed by the basic advection-reaction-dispersion equation:

$$\frac{\partial}{\partial t} Y_i = -\nabla(uY_i) + \frac{1}{\rho}\nabla(\rho K \nabla Y_i) + R_i + E_i$$

where Y_i represents concentration of a chemical species, u represents velocity of the medium, ρ the density of the medium, K the turbulence diffusivity, R_i the reaction rates, and E_i the emissions rate.

Because the chemical mechanisms and reaction rates are periodically updated, CTMs are modular in nature, allowing state-of-the science inputs to be used. The most recent gas-phase chemical mechanism for CMAQ is the Carbon Bond 5 chemical mechanism (CB05) [34]. CB05 contains about 80 chemical species³, including ozone, nitrogen dioxide,

²In contrast, Lagrangian CTMs follow emitted “puffs” or “plumes” of chemicals as they move through the atmosphere. Hybrid models combine these two approaches, tracking puffs of pollutants until they disperse to the degree that they may be included within a well-mixed grid cell.

³Some species with similar behaviors are grouped to simplify the model; for example, X02 denotes any hydrocarbon chain ending in two oxygen atoms. Throughout this work, species

and various hydrocarbons. The aerosol module, AERO6, adds another 60 aerosol species, calculating concentrations of primary and secondary particulates across a range of size distributions[11, 28].

1.4 Sensitivity Analysis

While building a working model of atmospheric conditions is interesting in its own right, the real value of the model comes from the ability to evaluate its sensitivity to inputs. This can take the form of zero-out analysis, or source apportionment, in order to quantify the total contributions of a source to ambient chemical concentrations; alternatively, comparison of model outputs can be made based on perturbations of input variables, such as atmospheric conditions, reaction rates, boundary conditions, or in our case, rates of emissions. These output comparisons can be used to evaluate various policy or environmental scenarios.

The simplest form of sensitivity analysis is known as *finite difference analysis*, or the subtractive or brute-force method. In this type of analysis, two or more model simulations are conducted: a single “base case”, and a series of sensitivity cases, each with a single or more input variables changed. The output concentration from each sensitivity case is then subtracted from the base case in order to discern the effect of the change in input. The advantage of this method is its accuracy, and its ability to fully capture the nonlinearities of the model; however, in order to evaluate n modeling scenarios, $n+1$ simulations must be conducted—a substantial computational cost.

A more sophisticated method is the *response surface model* (RSM). To construct a response surface, a series of “training runs” are conducted, each with small variations in the selected inputs. An ordinary regression is run on the outputs, and the resulting coefficients can be used to predict the results of hypothetical variations in the input factors used to train the model. While this method is effective in that it allows *ad hoc* scenarios to be considered without re-running atmospheric simulations, it requires a large number of training simulations—typically at least three per variable of interest[3]. Further, depending on the sampling methods used in the regression, expansion of the model to accommodate additional variation in inputs may require that some or all of the training simulations be re-conducted[21].

A third method, and the one central to our experiments, is the *decoupled direct method in three dimensions* (DDM-3D)[14, 24, 22]. While finite difference and RSM methods are conducted using output from a series of simulations, DDM-3D is an “inside the model” method that calculates sensitivity coefficients for a given set of sensitivity parameters at each modeled timestep. Numerically, the equation

$$C_{i,j} = \frac{\partial Y_i}{\partial X_j}$$

represents the sensitivity $C_{i,j}$ of the concentration of species Y_i to a change in sensitivity parameter X_j . Returning to our advection-reaction-dispersion equation, this takes the following form:

$$\frac{\partial}{\partial t} C_{i,j} = -\nabla(uC_{i,j}) + \frac{1}{\rho}\nabla(\rho K \nabla C_{i,j}) + J_i C_j + E_i$$

in **monospace** font refer to in-model species or pseudo-species, while those typeset conventionally refer to their real-world counterparts.

where J is the corresponding row from the Jacobian matrix representing interspecies chemical interaction kinetics[23]. It is important to note that the Jacobian must be calculated (unless its rate of change is quite slow) for each timestep[14].

A sensitivity parameter can be nearly any input variable in the model; in our case, we use emissions of six PM_{2.5} precursor species from a set of airports (i.e., the Cartesian product of our list of species and our list of airports). Once the coefficients C for each species i and input parameter j are calculated, we can evaluate hypothetical scenarios using the equation

$$Y'_i = Y_i + \Delta X_j C_{i,j}$$

Here, Y_i represents the base concentration of species i —that is, the concentration determined by CMAQ without the DDM calculations—while Y'_i represents the adjusted concentration. ΔX_j represents the multiplicative scaling factor of input parameter j . No change in the variable (i.e., the trivial case) would be shown by using $\Delta X_j = 0$. A “zero-out” scenario would use a value of $\Delta X_j = -1$; a thirty percent increase would use a value of $\Delta X_j = .3$. Note that while $C_{i,j}$ is output in concentration units, it should be interpreted in terms of sensitivity to scaled input parameters rather than concentration.

First-order⁴ DDM has been shown to compare favorably in output with brute force calculations for reductions of up to 20% for both primary and secondary emissions[16]. Because aviation LTO emissions represent far less than 20% of all PM_{2.5} precursor emissions for a typical regional modeling domain, our calculated sensitivity coefficients should be sufficiently accurate for perturbations of at least $\pm 100\%$. Only one simulation is required for a given set of sensitivity parameters, although each additional parameter adds to the required processing time. To expand the domain, additional sensitivity parameters can be calculated without re-running the entire suite of experiments.

2. METHODOLOGY

Previous studies quantifying the contribution of aircraft LTO emissions to ambient PM_{2.5} concentrations have used either finite difference or response surface methods, and accordingly have focused on either a few individual airports or the entire sector. Use of the DDM algorithm gives more flexibility in implementation by allowing any combination of designated airport and emission species to be designated as a sensitivity parameter. As a result, a more flexible set of scenarios can be modeled than with single-airport or domain-wide simulations. For example, routing more hub traffic through an airport in a more sparsely-populated region may decrease overall health impacts from aviation emissions without dramatically reducing overall air activity. However, due to the computationally intensive nature of DDM, it is important to design an experiment that balances resolution of output with a reasonable model runtime. Because of the extensive parallel computing requirements of our model, it is necessary to design a set of experiments to take advantage

⁴Higher-order implementations of the DDM algorithm—those that model nonlinear or cross-sensitivities to input parameters—have been implemented in CMAQ[35]. However, in light of our unusually large number of sensitivity parameters, the additional computational requirements are prohibitive in our design of experiments.

of the resources offered by a large-scale, high performance scientific computing cluster.

2.1 Domain

We restrict our domain to the continental United States (CONUS), using a 148×112 grid of $36 \text{ km} \times 36 \text{ km}$ cells typical for a model application of this scale⁵. Vertically, we limit our emissions to those from aircraft landing and takeoff operations (LTO; i.e. emissions within the lowest 3,000 ft). Cruise-altitude emissions do have additional impacts on both PM_{2.5} concentrations and health outcomes on a global scale[6]; however, it is more difficult to evaluate their impacts on a regional or continental scale. Additionally, while the atmospheric simulation is carried out with 34 vertical layers, health impacts are only evaluated at the surface layer, and thus we restrict our model output to sensitivities at ground-level.

Our flight segment and meteorological data come from 2005. To ascertain annual averages, the bare-minimum modeled period is two months, typically January and July; this period of time captures seasonal (winter-summer) variation while smoothing out short-term weather effects. Ideally, a simulation is conducted for a full year to reach an accurate annual average; however, in order to reduce runtime we decided on a compromise of four representative months (e.g. January, April, July, and October). For each contiguous modeled period, a “spin-up” period of two weeks is conducted in order to allow the model to reach a steady state from its initial conditions.

2.2 Sensitivity parameters

2.2.1 Airport selection

While an ideal experiment would allow us to truly quantify the sensitivity of PM_{2.5} concentrations to emissions from every airport in the United States, the marginal computational cost of each sensitivity parameter restricts their number. Because a relatively small number of the more than 2,000 airports represented in the AEDT are responsible for the majority of emissions, it is feasible to pick a subset of these airports that provide sufficient coverage in terms of geography and activity (in terms of fuel burn and passenger-trips) and still capture the majority of air quality and health impacts.

Our airport selection process consists of three criteria: (1) the top ninety-nine airports in the United States, selected by total fuel burned during takeoff or approach from U.S. airports, (2) smaller airports in NAAQS nonattainment areas conducting at least daily flight operations, and (3) the most active airport in any U.S. state lacking an airport meeting the above criteria. Using these guidelines, we selected 139 airports in the CONUS region (figure 1). In order to retain the ability to evaluate changes in emissions across all U.S. airports *en masse*, the remaining airports were grouped into a single group, giving us a total of 140 separate airport entities.

2.2.2 Emissions species

We divide emitted PM_{2.5} precursor species into six species groups (table 1). Different types of aircraft and various oper-

⁵While CMAQ is capable of much higher horizontal grid resolutions—regional and local experiments often employ grid cells of 12 km, 4 km, or even 1 km per side—total health impacts have not been found to vary significantly with the size of the grid cell[2].



Figure 1: Locations of the 139 airports included in the domain.

Table 1: Sensitivity parameters for $PM_{2.5}$ precursor subspecies[34].

Group	Species	Name
PSO ₄	PSO4	Primary sulfate
POC	POC	Primary organic carbon
PEC	PEC	Primary elemental carbon
VOC	ALD2	Acetaldehyde
	ALDX	Other aldehydes
	ETH	Ethene
	ETHA	Ethane
	ETOH	Ethanol
	FORM	Formaldehyde
	IOLE	Internal olefin bond
	MEOH	Methanol
	OLE	Terminal olefin bond
TOL	TOL	Toluene-like
	XYL	Xylene-like
SO ₂	SO2	Sulfur dioxide
NO _x	NO	Nitric oxide
	NO2	Nitrogen dioxide
	HONO	Nitrous acid

ation modes cause the precursors to be emitted in different quantities; tracking them separately allows for the maximum variety of scenarios to be evaluated.

Our initial approach used a single domain-wide emissions inventory for aircraft operations. We then “tagged” emissions to a designated airport on a grid-cell-by-grid-cell basis—that is, all LTO emissions from a grid cell were attributed to the major airport inside that grid cell. This method had two major shortcomings: first, aircraft ascent and descent trajectories often take them through more than one grid cell before reaching cruise altitude, thereby spreading their emissions over a larger area than our initial methodology could account for; secondly, some grid cells contain more than one major airport (e.g. San Francisco and Oakland International Airports). Using chorded flight segment emissions data from the AEDT, we created separate emissions files for each airport in the domain, eliminating crossover between airports and allowing full landing and takeoff segment paths to be assigned to each airport regardless of trajectory (figure 2).

One “flight” is a gate-to-gate aircraft operation—in other

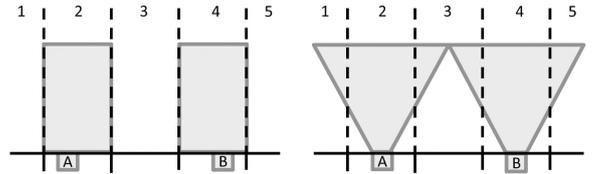


Figure 2: Left: initial, column-based emissions assignment strategy. Aviation emissions from airports A and B are assigned to cells 2 and 4, respectively. Right: when flight segment data is parsed, the column-spanning emissions patterns can be taken into effect; Airport A’s emissions are categorized to cells 1 through 3, while B’s are assigned to cells 3 through 5.

words, a trip. Each flight is made up of many flight segments, which are classified by location, distance, and aircraft operation mode. Using flight segment information obtained from AEDT, we select segments from takeoff through initial climb, and descent through landing. Flight segments were further limited to those below 3,000 ft. Departure flight segments were assigned to the airport of departure, and arrival flight segments were assigned to the airport of arrival. Segments were processed into gridded emission rate files using AEDTProc[4]. Background emission rates from EPA’s National Emissions Inventories (NEI-2005) were processed into grid-based emissions using the Sparse Matrix Operator Kernel Emissions (SMOKE)[29]. Meteorology for 2005 was processed from the Weather Research and Forecasting model (WRF), with outputs downscaled from NASA’s Modern-Era Retrospective Analysis for Research and Applications data (MERRA)[26]. We use CMAQ v5.0.2 with DDM-3D to process our inputs to generate sensitivity and concentration output files. We group particulate species into primary and secondary groups according the methodology in Arunachalam et al. 2008[1].

2.2.3 Airport grouping mechanism

Total model runtime is linearly related to the number of sensitivity parameters used; thus, any reduction in this number has a directly proportional reduction in total CPU-hours needed (table 4). One method to reduce the total number of sensitivity parameters used is to group airports with relatively non-overlapping sensitivity fields into a single parameter. A similar approach was taken with an RSM-based study of particulate emissions from power plants in the Eastern United States[15].

In order to group airports, we must first determine what constitutes a non-overlapping pair. Because CMAQ carries calculations to a very high level of numerical precision⁶,

⁶It is nontrivial to define accuracy and precision in the context of comparing CMAQ model output with the “real world”. For one, model output is comprehensive spatially and temporally but is averaged across both dimensions, while monitored data is the opposite—spatially and temporally precise, but lacking in coverage across both dimensions. Secondly, the monitors used to measure ambient particulate matter are “uncalibrated”, with consistency between monitors achieved by following a federally-regulated methodology rather than comparison to an established benchmark[32]. Finally, at high enough levels of numerical precision, the

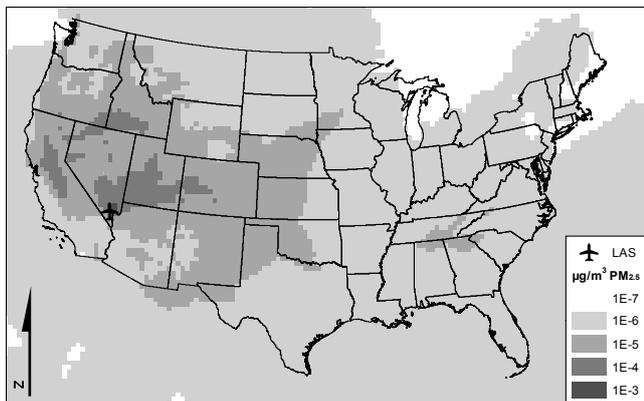


Figure 3: Log-scale regions of influence for Las Vegas-McCarren International Airport.

virtually every grid cell in the domain will have a nonzero sensitivity to emissions from even the smallest airport (figure 3). Thus, in order to group more than one airport into a single sensitivity parameter, we must establish a threshold below which emissions sensitivities are treated as zero. Using preliminary data from two one-week simulations—one in January, one in July—we determined an approximate estimate of the extent of influence for each airport in the domain. We then determined the distance r of the furthest cell above the threshold chosen above, and drew a circle of radius r around the home grid cell of the airport. A buffer of an additional two grid cells was applied to the outside of each circular region; if an airport was too small to generate sensitivities above the threshold, this two-cell buffer was applied to its home grid cell. Airports were then randomly sorted into non-overlapping groups. Using a threshold of $1 \times 10^{-4} \mu\text{g}/\text{m}^3$, we were able to fit all 139 airports into about thirty groups, with an order-of-magnitude plus two-grid-cell buffer around each airport’s domain of influence added for conservative purposes.

2.3 Computing Environment

Because their grid-cell based domains are inherently parallel, Eulerian models such as CMAQ are well-suited to implementation on parallel computing systems. Grid cells are divided into contiguous sub-domains, with each subdomain assigned to a different processor for parallel calculations. While a CMAQ simulation could in theory be parallelized on a system of processors up to the number of horizontal grid cells used—in our case, 16,576—our benchmarking efforts showed a considerable decrease in per-CPU efficiency past a certain number of processors due to the demand on input/output file access (figure 4). Doubling the per-simulation number of processors from eight to sixteen resulted in a halving of wall-clock runtime with virtually no change in total CPU-hours. However, increasing the number of processors from sixteen to thirty-two reduced wall-clock time by less than a factor of two, while increasing CPU hours by about 37%.

CMAQ outputs themselves are sensitive to factors such as levels of optimization and processor instruction sets. Regulatory limits on exposure are given in terms of $\mu\text{g}/\text{m}^3$, while commercially-available monitors have a detection limit of $0.04\text{--}1 \mu\text{g}/\text{m}^3$ [12].

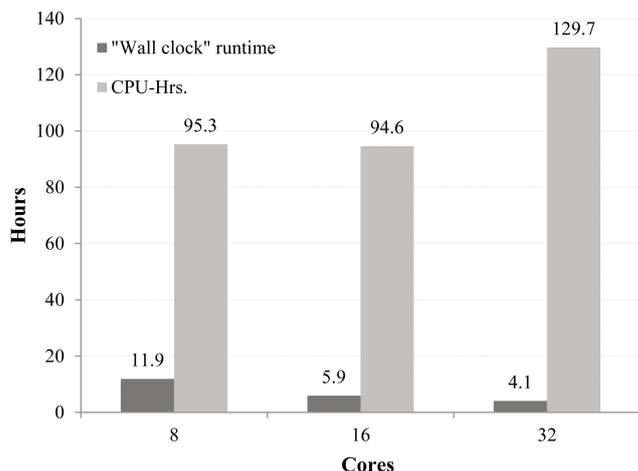


Figure 4: Model runtime for one simulated day, showing decreasing processor efficiency between 16 and 32 cores. 18 sensitivity parameters were used in these benchmarks.

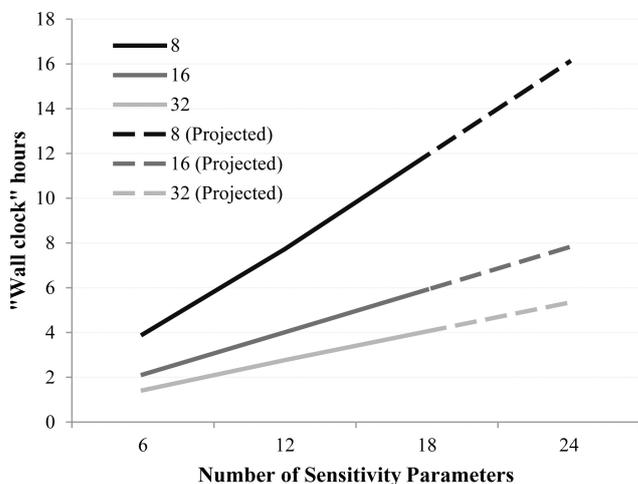


Figure 5: Model runtime, showing linear increase in processing time as parameters are added.

We noted a near-linear increase in model runtime with each additional sensitivity parameter (figure 5). Our benchmarking efforts also showed us that total memory available was a severe limiting factor in the number of sensitivity parameters per simulation (slightly more than one parameter per gigabyte of RAM available). While Stampede does have a number of large-memory (1TB) nodes available for use, designing our experiment around many parallel simulations instead of fewer, high-memory simulations proved to keep the overall duration of the experiment low. For an eighteen-parameter simulation conducted on sixteen processors, the total wall-clock duration was about six hours—a good balance between minimizing both the total number of CPU-hours and the wall-clock experiment duration.

Using eighteen parameters per simulation allows us to fit three groups (three groups, each with six species for a total of eighteen parameters) per simulation. Thus, to fit all thirty groups, we needed ten parallel model simulations, each

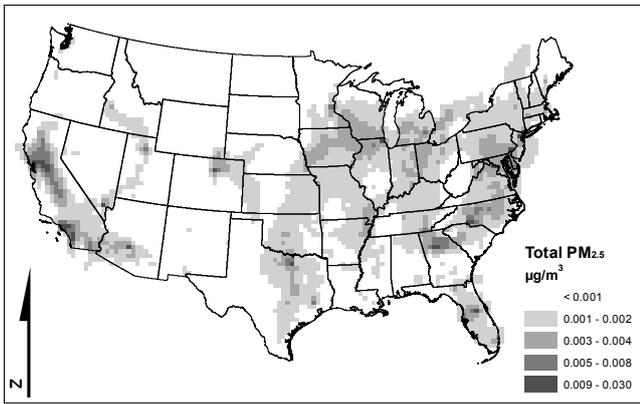


Figure 6: Total $PM_{2.5}$ sensitivity (primary and secondary) to all airports in the domain.

operating for forty-five days per modeled month. In order to run a quarterly experiment, we estimate a total wall-clock time of 450 days and about 173,000 CPU-hours. However, this is parallelizable by a factor of forty (ten groups, with four periods), for an optimal, perfectly-parallel wall clock runtime of 11.25 days (table 2).

Table 2: Model runtimes (16 processor, 18 parameter configuration).

Period	Wall Clock		CPU-Hrs
	Hours	Days	
Day	6	0.25	96
Run (45 days)	270	11.25	4,320
Period (10 runs)	2,700	112.50	43,200
Experiment (4 periods)	10,800	450.00	172,800

Another factor to consider with the use of DDM is the extremely large output files generated. Each day generates around fifteen gigabytes of data, including comprehensive, state-of-the-simulation restart file at the end of each day. However, for the purposes of evaluating health impacts, only about 4.6 GB need be retained. An experiment with four months of duration will generate between six and eight terabytes of uncompressed output.

3. RESULTS

The aggregation of modeled sensitivities from all sources results in an intuitive sensitivity grid, with high values near major airports (figure 6). However, the true power of our dataset is shown in disaggregate form, where sensitivity of each individual output species can be linked to any combination of airport and emissions group. For example, we can show that primary $PM_{2.5}$ sensitivities to emissions from Atlanta’s Hartsfield-Jackson International Airport (ATL) are grouped around the airport itself, while secondary sensitivities are located further downwind, reaching as far as central North Carolina (figure 7). We can also show secondary $PM_{2.5}$ sensitivities to emissions from Charlotte-Douglas International Airport (CLT), despite the fact that they overlap geographically with sensitivities to emissions from ATL.

As an example, we compare preliminary $PM_{2.5}$ sensitivities in two grid cells: one directly north of CTL, and the other directly south (CTL is nominally located in the southern

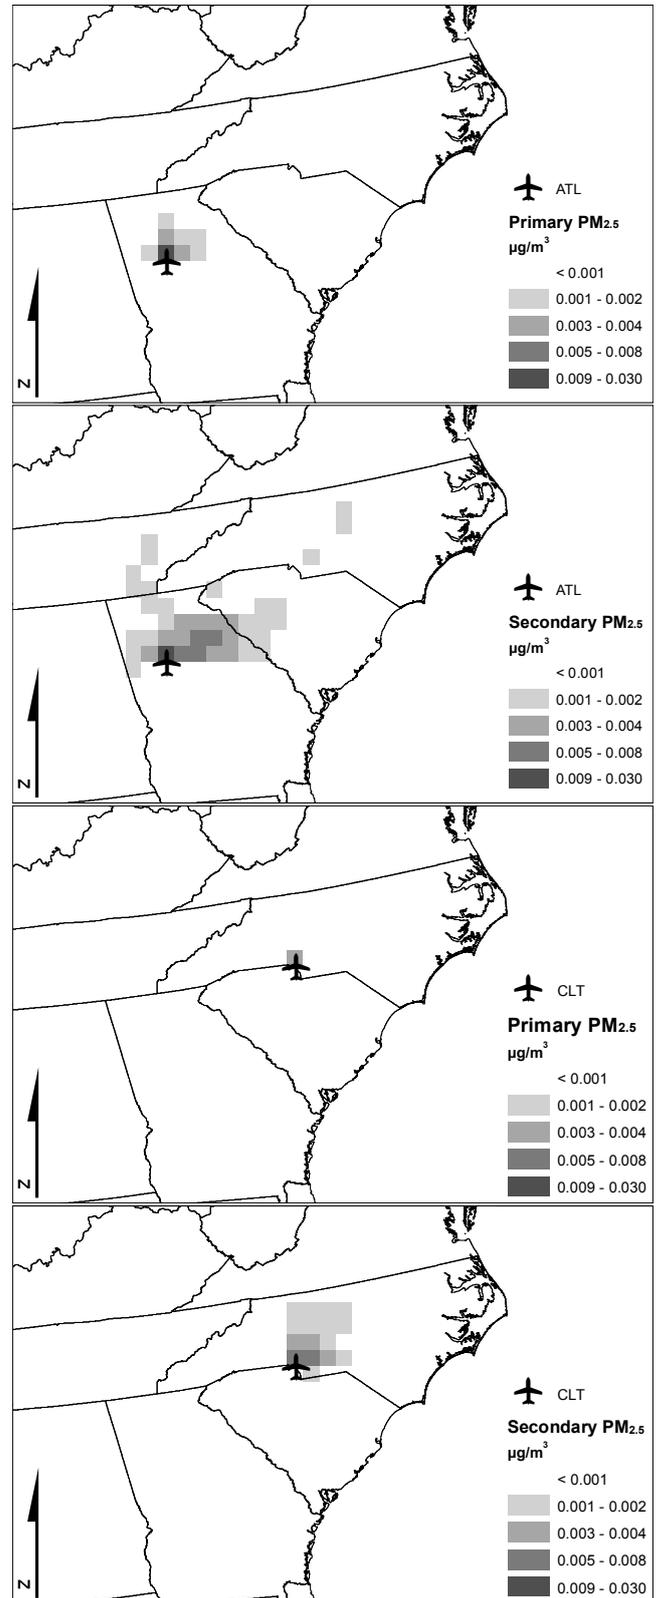


Figure 7: Top to bottom: Primary $PM_{2.5}$ sensitivities to emissions from Atlanta’s Hartsfield-Jackson International Airport (ATL); secondary $PM_{2.5}$ sensitivities to ATL; primary $PM_{2.5}$ sensitivities to emissions from Charlotte-Douglas International Airport (CLT); secondary $PM_{2.5}$ sensitivities to CLT.

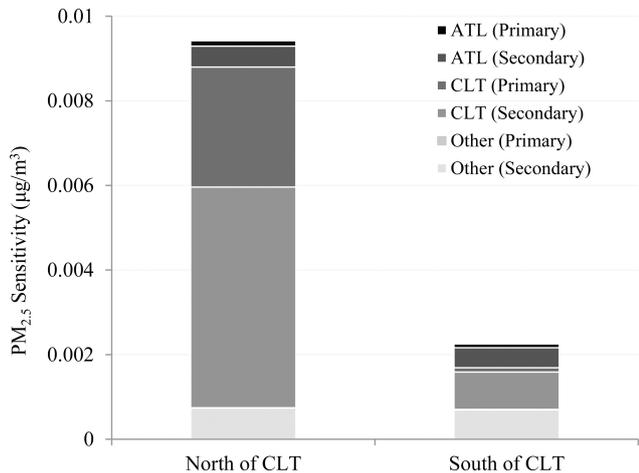


Figure 8: Relative contributions of ATL, CLT and all other airports to PM_{2.5} sensitivities in the grid cells directly north and south of CLT.

grid cell, though very close to the border). In the cell to the south, nearly one-third of PM_{2.5} is due to emissions from ATL, despite the fact that CTL is much closer (figure 8). However, in the grid cell to the north—in the direction of the prevailing winds—PM_{2.5} sensitivities are much higher. Further, we see that in the north cell, the proportion of primary PM_{2.5} sensitivity is much higher than in the south (54.3% vs 11.5%). Thus, for each grid cell in the domain, we can estimate the effects of policy introductions by scaling any one of our sensitivity parameters. For example, introducing low-sulfur fuel at major airports[5]; an increase in activity at airports near populated areas; or a re-balancing of aircraft operations from one hub to another.

In order to calculate health impacts, PM_{2.5} sensitivities must be combined with geographically-distributed population densities using health impact functions. Previous work using subtractive and RSM methodology have attributed between one and three hundred premature mortalities per year from aviation LTO emissions[3, 20, 7]. However, these estimates describe the response of the model to domain-wide or aggregated for multiple airports, rather than airport-level emissions. Correspondingly, they suggest policy changes at the domain-wide level. Changes in fuel composition, for example, have been estimated to have the potential to reduce annual premature mortalities by 900–4000 worldwide[5].

While there are several epidemiological studies relating health outcomes with total PM_{2.5}, there are fewer studies linking health impacts with its specific components. Studies have pointed out that the changing composition of fine particulate matter over time, especially in response to various control policies, could have an effect on concentration-response functions. Studies which produce data on PM_{2.5} speciation could aid in this task[18, 19].

4. CONCLUSIONS

Using parallel computing resources, we have designed and initiated a sensitivity modeling exercise spanning 139 airports and 6 species, representing 76.7% of flights and 91.2% of fuel burn. Preliminary model runs have shown that the project will require 172,800 CPU-hours of modeling time; however,

parallel implementation allows a optimal runtime of 11.25 days across 40 compute nodes (each with 16 processors). Use of XSEDE resources makes a simulation of this scale feasible. When completed, the simulation outputs will describe individual sensitivities of fine particulate concentrations to emissions from a group of airports that is diverse in both location and activity, allowing health impacts to be quantified for a variety of scenarios. Separation of sensitivities by PM_{2.5} precursor species will allow health impacts to be estimated with greater accuracy than allowed by modeling exercises using unspiciated particulate sensitivities. Finally, estimating the effects of growth in aviation activity at the airport level coupled with high-resolution future population estimates will allow the assessment of differential growth scenarios for airport activity related emissions and their human health impacts.

5. ACKNOWLEDGMENTS

This work would not have been possible without: S. Napelenok, U.S. EPA; M. Reed, UNC ITS; C. Coats, B.H. Baek, M. Woody, P. Vennam and S.Y. Chang, UNC IE; and M. Serre, UNC ESE. This work was funded by PARTNER under grants to UNC. PARTNER is funded by FAA, NASA, Transport Canada, U.S. DOD, and EPA. The aviation emissions inventories used for this work were provided by U.S. DOT Volpe Center and are based on data provided by the U.S. FAA and EUROCONTROL in support of the objectives of the ICAO Committee on Aviation Environmental Projection CO₂ Task Group. Any opinions, finding, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. DOT, Volpe Center, the U.S. FAA, EUROCONTROL, ICAO or PARTNER.

6. REFERENCES

- [1] S. Arunachalam, B. H. Baek, A. Holland, Z. Adelman, F. Binkowski, A. Hanna, T. Thrasher, and P. Soucacos. An improved method to represent aviation emissions in air quality modeling systems and their impacts on air quality. In *Proceedings of the 13th Conference on Aviation, Range and Aerospace Meteorology, New Orleans, LA*, volume 135626, 2008.
- [2] S. Arunachalam, B. Wang, N. Davis, B. H. Baek, and J. I. Levy. Effect of chemistry-transport model scale and resolution on population exposure to PM_{2.5} from aircraft emissions during landing and takeoff. *Atmos Environ*, 45(19):3294–3300, 2011.
- [3] A. Ashok, I. H. Lee, S. Arunachalam, I. A. Waitz, S. H. Yim, and S. R. Barrett. Development of a response surface model of aviation’s air quality impacts in the United States. *Atmos Environ*, 77:445–452, 2013.
- [4] B. Baek, S. Arunachalam, M. Woody, P. Vennam, M. Omary, F. Binkowski, and G. Fleming. A new interface to model global commercial aircraft emissions from the FAA Aviation Environmental Design Tool (AEDT) in air quality models. In *Proceedings of the 11th annual CMAS Conference.*, October 2012.
- [5] S. Barrett, S. Yim, C. Gilmore, L. Murray, S. Kuhn, A. Tai, R. Yantosca, D. Byun, F. Ngan, X. Li, J. Levy, A. Ashok, J. Koo, H. Wong, O. Dessens, S. Balasubramanian, G. Fleming, M. Pearlson, C. Wollersheim, R. Malina, S. Arunachalam,

- F. Binkowski, E. Leibensperger, D. Jacob, J. Hileman, and I. Waitz. Public health, climate, and economic impacts of desulfurizing jet fuel. *Envir Sci Tech*, 46:4275–4282, 2012.
- [6] S. R. Barrett, R. E. Britter, and I. A. Waitz. Global mortality attributable to aircraft cruise emissions. *Envir Sci Tech*, 44(19):7736–7742, 2010.
- [7] E. Brunelle-Yeung, T. Masek, J. J. Rojo, J. I. Levy, S. Arunachalam, S. M. Miller, S. R. Barrett, S. R. Kuhn, and I. A. Waitz. Assessing the impact of aviation environmental policies on public health. *Transp Policy*, 2014.
- [8] Bureau of Transportation Statistics. National Transportation Statistics Table 1-40. 2014.
- [9] Bureau of Transportation Statistics. T-100 segment data. 2014.
- [10] D. W. Byun and J. Ching. *Science algorithms of the EPA Models-3 community multiscale air quality (CMAQ) modeling system*. US Environmental Protection Agency, Office of Research and Development Washington, DC, USA, 1999.
- [11] A. G. Carlton, P. V. Bhave, S. L. Napelenok, E. O. Edney, G. Sarwar, R. W. Pinder, G. A. Pouliot, and M. Houyoux. Model representation of secondary organic aerosol in CMAQv4.7. *Envir Sci Tech*, 44(22):8553–8560, 2010.
- [12] J. C. Chow, P. Doraiswamy, J. G. Watson, L.-W. A. Chen, S. S. H. Ho, and D. A. Sodeman. Advances in integrated and continuous measurements for particle mass and chemical composition. *J Air Waste Ma*, 58(2):141–163, 2008.
- [13] CSSI, Inc. JPDO–NextGen–TDM23: Investigation of Aviation Emissions for Future Aviation Activity Scenario. Technical report, 2009.
- [14] A. M. Dunker. The decoupled direct method for calculating sensitivity coefficients in chemical kinetics. *J Chem Phys*, 81:2385–2393, 1984.
- [15] C. Jang. Applications and evaluations of response surface model (rsm) using cmaq in the united states. In *Proceedings of the 11th annual CMAS Conference.*, October 2012.
- [16] B. Koo, G. Wilson, R. Morris, A. Dunker, and G. Yarwood. Comparison of source apportionment and sensitivity analysis in a particulate matter air quality model. *Envir Sci Tech*, 43:6669–6675, 2009.
- [17] F. Laden, J. Schwartz, F. E. Speizer, and D. W. Dockery. Reduction in fine particulate air pollution and mortality: extended follow-up of the Harvard Six Cities study. *Am J Resp Crit Care*, 173(6):667, 2006.
- [18] J. Levy, M. Woody, B. Baek, U. Shankar, and S. Arunachalam. Current and future particulate-matter-related mortality risks in the United States from aviation emissions during landing and takeoff. *Risk Analysis*, 32:237–249, 2012.
- [19] J. I. Levy, D. Diez, Y. Dou, C. D. Barr, and F. Dominici. A meta-analysis and multisite time-series analysis of the differential toxicity of major fine particulate matter constituents. *Am J Epidemiol*, 175(11):1091–1099, 2012.
- [20] J. I. Levy, M. Woody, B. H. Baek, U. Shankar, and S. Arunachalam. Current and future particulate-matter-related mortality risks in the United States from aviation emissions during landing and takeoff. *Risk Anal*, 32(2):237–249, 2012.
- [21] T. Masek. A response surface model of the air quality impacts of aviation. Master’s thesis, Massachusetts Institute of Technology, 2008.
- [22] S. Napelenok, D. S. Cohan, M. Odman, and S. Tonse. Extension and evaluation of sensitivity analysis capabilities in a photochemical model. *Environ Modell Softw*, 23(8):994–999, 2008.
- [23] S. L. Napelenok. *Sensitivity Analysis in Air Quality Models for Particulate Matter*. PhD thesis, Georgia Institute of Technology, 2006.
- [24] S. L. Napelenok, D. S. Cohan, Y. Hu, and A. G. Russell. Decoupled direct 3D sensitivity analysis for particulate matter (DDM-3D/PM). *Atmos Environ*, 40(32):6112–6121, 2006.
- [25] C. A. Pope III, R. T. Burnett, M. J. Thun, E. E. Calle, D. Krewski, K. Ito, and G. D. Thurston. Lung cancer, cardiopulmonary mortality, and long-term exposure to fine particulate air pollution. *JAMA*, 287(9):1132–1141, 2002.
- [26] M. M. Rienecker, M. J. Suarez, R. Gelaro, R. Todling, J. Bacmeister, E. Liu, M. G. Bosilovich, S. D. Schubert, L. Takacs, G.-K. Kim, et al. Merra: Nasa’s modern-era retrospective analysis for research and applications. *J Climate*, 24(14), 2011.
- [27] C. Roof, A. Hansen, G. Fleming, T. Thrasher, A. Nguyen, C. Hall, F. Grandi, B. Kim, S. Usdrowski, and P. Hollingsworth. Aviation Environmental Design Tool (AEDT) system architecture. *Doc# AEDT-AD-01*, 2007.
- [28] H. Simon and P. V. Bhave. Simulating the degree of oxidation in atmospheric organic particles. *Envir Sci Tech*, 46(1):331–339, 2011.
- [29] U.S. Environmental Protection Agency. National Emissions Inventory (NEI). 2005.
- [30] U.S. Environmental Protection Agency. National Ambient Air Quality Standards (NAAQS). 2012.
- [31] U.S. Federal Aviation Administration. Destination 2025, 2012.
- [32] W. Wilson, J. C. Chow, C. Claiborn, W. Fusheng, J. Engelbrecht, and J. G. Watson. Monitoring of particulate matter outdoors. *Chemosphere*, 49(9):1009–1043, 2002.
- [33] M. Woody, B. H. Baek, Z. Adelman, M. Omary, Y. F. Lam, J. J. West, and S. Arunachalam. An assessment of aviation’s contribution to current and future fine particulate matter in the United States. *Atmos Environ*, 45:3424–3433, 2011.
- [34] G. Yarwood, S. Rao, M. Yocke, and G. Whitten. Updates to the carbon bond chemical mechanism: CB05. *Final report to the US EPA, RT-0400675*, 2005.
- [35] W. Zhang, S. Capps, Y. Hu, A. Nenes, S. Napelenok, and A. Russell. Development of the high-order decoupled direct method in three dimensions for particulate matter: enabling advanced sensitivity analysis in air quality models. *Geosci Model Dev*, 5:255–268, 2012.

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle		5. Report Date	
		6. Performing Organization Code	
7. Author(s)		8. Performing Organization Report No.	
9. Performing Organization Name and Address		10. Work Unit No. (TRAIS)	
		11. Contract or Grant No.	
12. Sponsoring Agency Name and Address		13. Type of Report and Period Covered	
		14. Sponsoring Agency Code	
15. Supplementary Notes			
16. Abstract			
17. Key Words		18. Distribution Statement	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages	22. Price