

SPSA/SIMMOD Optimization of Air Traffic Delay Cost

Nathan L. Kleinman

Department of Mathematics, Brigham Young University
Provo, UT 84602 USA
kleinman@brutus.mts.jhu.edu

Stacy D. Hill

The Johns Hopkins University Applied Physics Laboratory
Johns Hopkins Road, Laurel, MD 20723 USA
stacy.hill@jhuapl.edu

Victor A. Ilenda

The Johns Hopkins University Applied Physics Laboratory
Johns Hopkins Road, Laurel, MD 20723 USA
victor.ilenda@jhuapl.edu

Abstract

The cost of delay is a serious and increasing problem in the airline industry. Air travel is increasing, and already domestic airports incur thousands of hours of delay daily, costing the industry \$2 billion a year. One strategy for reducing total delay costs is to hold planes for a short time at the gate in order to reduce costly airborne congestion. In a network of airports involving hundreds of flights, it is difficult to determine the amount to hold each flight at the gate. This paper discusses how the optimization procedure simultaneous perturbation stochastic approximation (SPSA) can be used to process delay cost measurements from air traffic simulation packages and produce an optimal gate holding schedule. As a test case, the SIMMOD air traffic simulation package was used to model a simple four-airport network. Initial delay costs are reduced up to 10.3%.

1. Introduction

Air traffic delay is a rapidly increasing problem in the United States. The National Transportation Research Board reported that in 1990, over 20,000 hours of delay were incurred at each of 21 airports nationwide, with 12 more airports expected to reach that mark by 1997 (see [12]). Airlines report that delay costs their industry \$2 billion annually (roughly the same amount as the industries total losses in 1991), and the Federal Aviation Administration (FAA) expects the demand for air travel to increase 25% by the year 2000 [20]. Congestion is increased further by airlines' desire to use "hub" airport systems. Here, airlines schedule large numbers of flights from outlying airports so that they arrive at the hub airport at approximately the same time. Passengers are then exchanged, and a new group of flights leaves the hub airport, again at roughly the same time. As an example, Odoni [11] cites Atlanta, a major hub for Delta Airlines, where at least six times a day "banks" of 100 flights arrive and depart within approximately one-hour periods.

Clearly the need exists for reducing air traffic delay. Current options include constructing new airports or runways, encouraging or constraining airlines to spread out arrivals and

departures more evenly at congested airports in order to reduce "peak" period congestion, and using larger aircraft in order to transport more passengers per flight. But these methods are either very expensive or unlikely to be implemented soon [20]. Vranas, Bertsimas, and Odoni [20] state that "ground holding policies" offer a more promising way to reduce delay costs. Air traffic delay can be divided into three categories: induced holding at the gate, delay while taxiing, and airborne delay. Ground holding policies attempt to assign a small amount of delay to each flight prior to leaving the gate in order to reduce network-wide congestion and, particularly, to reduce the much more costly airborne delay.

Determining the amount of gate delay to assign to each flight on a particular day for an entire network of airports is a very large optimization problem. Software packages exist which simulate (some with a high degree of detail) many aspects of flights and airport operations in a network of airports. But these packages do not *optimize*; they only take a ground holding policy from the user and output the associated delay cost. To set up and solve a detailed stochastic programming problem which models a network of airports and finds an optimal ground holding policy would be intractable. Vranas, Bertsimas, and Odoni [20] describe an integer programming model for the multi-airport ground-holding problem; however, an integer programming approach would not model the activity of the air network in as much detail as would simulation packages. For instance, their model does not account for taxiway and airspace congestion. Furthermore, it is often necessary to optimize objective functions that are highly nonlinear and for which the form is unknown and only noisy measurements are available. In 1993 researchers at The MITRE Corporation [6] stated that their current methods for determining ground-holding policies were optimal *if* the following assumptions were made: (1) the National Air Space is deterministic, that is, there are no random capacities or travel times, (2) all airway capacities are infinite or nonrestrictive, and (3) an aircraft's flight plan includes only one destination. They noted that these assumptions were too restrictive, but they found no other optimization method at that time which did not include

at least two of them. Later, Helme [5] proposed a multicommodity minimum cost flow method which allowed somewhat random capacities and negated the second and third assumptions. But, this method used deterministic take-off times, a linear objective function, and discrete delay times (problem size increases dramatically with finer discretizations). Again, an optimization method is needed which incorporates the high level of modelled detail and stochastic nature of air traffic simulation packages.

Simultaneous perturbation stochastic approximation (SPSA) is an iterative technique for finding local optimizers of linear or nonlinear objective functions from many types of systems [16, 17, 18]. SPSA is like other Kiefer-Wolfowitz stochastic approximation algorithms (such as finite difference stochastic approximation) in that SPSA only requires measurements (possibly noisy) of an objective function to form gradient estimates and converge to a local optimum. SPSA differs from FDSA in that SPSA only requires *two* objective function evaluations per gradient estimate while FDSA requires $2p$ evaluations, where p is the number of system parameters being estimated. This gives SPSA a significant advantage in high dimensional problems, especially when evaluating the objective function is expensive or time consuming. Further improvements can be achieved by using the method of "common random numbers" to reduce the variance of the SPSA estimation error (see [7]). SPSA has been applied successfully in optimal sensor location problems [14], signal light timing determination in a simulated traffic network [8, 15], nonlinear adaptive control [19], and queueing system optimization [1].

SPSA appears to be an ideal tool for solving the ground-holding problem. This paper discusses how SPSA can be used in conjunction with an air traffic simulation package to improve ground-holding policies in a network of airports. Because SPSA only requires objective function measurements, a detailed simulation package can be used to estimate the delay cost associated with a particular ground-holding policy. Each iteration SPSA will generate a new ground-holding policy, based on previous delay cost measurements, until a desired level of improvement is reached.

The purpose of this paper is to introduce and describe an SPSA-based method for solving air traffic-related optimization problems. In Section 2. we outline the form and characteristics of the SPSA algorithm and explain how SPSA can be combined with air traffic simulation output to form an iterative optimization method. In Section 3. we describe the air traffic network simulation package SIMMOD and explain why this particular simulation tool is a good candidate for supplying simulated delay cost measurements to the SPSA optimization algorithm.

In Section 4. we discuss a 168-dimensional test case for this method which involves a simple hypothesized network of four airports and flights among those airports. We start with a flight schedule and an initial ground-holding policy which assigns no gate-delay to every flight, and then we use the SPSA optimization method to find a ground holding policy which improves the total delay cost in the network. This test illustrates how planners can use this method to improve existing ground-holding policies.

We also perform tests for another potential application of the optimization method. Suppose, before a ground-holding policy is implemented for a particular schedule of flights, that new weather information is received, telling planners that capacity at a certain airport will be reduced during a particular time period. The SPSA optimization method could then be used to adjust the network-wide ground-holding policy to account for the period of low capacity at the airport experiencing

poor weather.

For each test, initial and improved delay cost values are given. Comparisons are also made of the SPSA algorithm with and without the use of common random numbers. Note that although these results are for a fictitious network of airports, the optimization method is general and can be applied in an identical manner to a simulation model of a real network.

2. Simultaneous Perturbation Stochastic Approximation

The simultaneous perturbation stochastic approximation (SPSA) algorithm was introduced and developed by Spall [16, 17, 18]. The algorithm uses objective function measurements to iteratively update system control parameters until parameter values are reached which locally optimize the objective function. Specifically, let $\theta \in R^p$ be a vector whose components represent system parameters we wish to control (each flight's ground-holding delays, for example). Then let $L(\theta)$ represent the objective function we wish to optimize. The goal is to find a zero of the gradient of this objective function. That is, we want a θ such that

$$g(\theta) \equiv \frac{\partial L}{\partial \theta} = 0.$$

The SPSA algorithm attempts to find a local minimizer θ^* by starting at a fixed $\hat{\theta}_0$ and iterating according to the following scheme:

$$\hat{\theta}_{k+1} = \hat{\theta}_k - a_k \hat{g}_k(\hat{\theta}_k). \quad (1)$$

Here $\{a_k\}$ is a *gain* sequence of positive scalars satisfying certain conditions [18] (in particular, $a_k \rightarrow 0$ and $\sum_{k=1}^{\infty} a_k = \infty$), and \hat{g}_k is an estimate of the gradient g whose l -th component is defined as

$$\hat{g}_{kl} = \frac{y_k^+ - y_k^-}{2c_k \Delta_{kl}}. \quad (2)$$

Here, y_k^+ represents a (perhaps noisy) measurement of $L(\hat{\theta}_k + c_k \Delta_k)$, and y_k^- is a similar measurement of $L(\hat{\theta}_k - c_k \Delta_k)$. The sequence $\{c_k\}$ is a sequence of positive scalars such that $c_k \rightarrow 0$ (see [18]), and $\Delta_k \in R^p$ is a vector of p mutually independent random variables satisfying conditions in [18]. For example, the components of Δ_k could be independent Bernoulli(± 1) distributed random variables [13].

Observe the numerator in equation (2) is the same for each l . Thus, only two measurements of the objective function are required to obtain the SPSA gradient estimate at each iteration. To illustrate, if $\hat{\theta}_k$ represents our current estimate of the best ground-holding policy θ^* , then $\hat{\theta}_k + c_k \Delta_k$ and $\hat{\theta}_k - c_k \Delta_k$ are "perturbed" policies. Then y_k^+ and y_k^- could be total delay cost values obtained by performing two simulations of the air traffic network, one using the ground-holding policy $\hat{\theta}_k + c_k \Delta_k$ and one using $\hat{\theta}_k - c_k \Delta_k$. These values would then be used in (2) and (1) to obtain a new estimate $\hat{\theta}_{k+1}$ of the best ground-holding policy θ^* . This process would be repeated until a desirable level of improvement in delay cost is reached.

The method of common random numbers [2, 4, 7, 9, 10], which can reduce the variance of the SPSA estimation error, is implemented by using the same random number seeds to drive both the simulation which yields y_k^+ and the simulation which yields y_k^- each iteration. This way, many of the events in the simulation yielding y_k^+ will use the same random numbers as the corresponding events in the simulation yielding y_k^- . This can reduce the variance of \hat{g}_k and thus reduce the variance of the SPSA error estimates. See [7] for a complete description.

Note that the SPSA algorithm is very general and can be applied in many different situations to optimize many different kinds of objective functions. Even within the air traffic realm, there are many examples. For instance, θ could represent the arrival and subsequent departure times for a bank of flights, while $L(\theta)$ could be the percentage of passengers able to make connecting flights within that bank. Or, θ could be a ground-holding policy and $L(\theta)$ a function of fuel consumption. The flexibility of the algorithm stems from the fact that only (noisy) objective function *measurements* are required, not full objective function or gradient information.

3. SIMMOD

The objective function measurements required by the SPSA algorithm can come from a real system or from a computer simulation which models a real system, depending on the purpose of the optimization. Although SPSA has been used for real-time control, our purpose here, and especially with the example in Section 4., is to outline a method for making decisions on policies minutes, hours, or days prior to the time of implementation. Thus, detailed simulations are appropriate means for obtaining objective function measurements.

As mentioned in the introduction, a number of simulation packages exist which model activity in a network of airports. For our tests we chose to use the simulation package SIMMOD which was developed by CACI Products Company for the FAA. SIMMOD is used by the FAA and is well recognized as a credible tool for modelling inter- and intra-airport activity. SIMMOD models and tracks individual aircraft of different types as they move in a network of airports from gate to taxiway to runway to airspace and ultimately to the final destination gate.

In the air, SIMMOD implements speed control, vectoring, and several different holding procedures in order to maintain proper separation distances between different types of aircraft. Sequence control and spacing for merging aircraft are also modelled. SIMMOD also takes into account route and sector capacity constraints as well as wind speeds.

In the interface between airspace and airport SIMMOD models many detailed operations including arrival and departure procedures, separation constraints for Instrument Flight Rules (IFR) and Visual Flight Rules (VFR) approaches, runway selection, missed approaches, and takeoff and landing runway distances and times.

SIMMOD provides a detailed model of ground activity as well. SIMMOD accounts for separation distances along taxiways, runway crossing while taxiing, and different types of departure queues as planes approach the runway. SIMMOD also can determine optimal taxiway and gate choices for each flight given the availability of each gate and the specific set of gates used by each airline.

Before running simulations, the user inputs information about the airspace, airfields, and simulation events. The user is allowed to set up gates, taxiways, runways, and airspace routes as he wishes. The user also inputs a schedule of flights and other events such as runway closures, wind changes, etc. As mentioned previously, each of the components of the airspace and airfields and each flight can be modelled in detail, but SIMMOD provides default settings for many parameters. This gives the user the flexibility to provide more detail for the aspects of the air traffic system he is most concerned with and less detail for other aspects.

Another valuable feature of SIMMOD is its capability to produce reports of many kinds of delay and usage results. SIMMOD provides aggregate statistics as well as reports broken down by airline, gate, runway, airport, sector, or route. In the

tests described in Section 4., the airborne and taxiway delay times were important parts of our objective function measurements.

4. Numerical Example of the SPSA Optimization Method

The purpose of performing the numerical example described in this section is to illustrate an application of the SPSA optimization method. As shown above, this method is very flexible. It can accommodate many different objective functions and varying degrees of simulation complexity, depending on the needs and resources of the user. Hence, the numerical results obtained for the following example problem serve as qualitative evidence of the usefulness and applicability of the SPSA optimization method.

4.1. Problem Setup

In this problem we use SPSA to find a ground-holding policy that significantly reduces the total delay cost for aircraft movement within a network of airports. The network in our example is comprised of four simplified airports. Air routes, each comprised of a series of air nodes and links, connect the airports as shown in Figure 1. Each airport consists of one runway, one gate, and taxiways leading from the gate to each end of the runway. The structure of the airspace and airfields and the location of the nodes and links were input to SIMMOD by means of a graphical input program within SIMMOD. Node characteristics, such as air and taxiway holding strategies, capacity, and altitude; link characteristics, such as capacity, passing restrictions, and average speed for different aircraft types; and other information, such as takeoff, landing, and in-flight separation distances, were input to SIMMOD as well. The input information for our project was reviewed and validated by SIMMOD support staff at SABRE Decision Technologies.

Next, a flight schedule was created. The schedule consisted of 168 flights (enough flights to cause a realistic amount of congestion in the system). The departures were distributed among the four airports and spread out nearly evenly over a three-hour period. Over the three-hour period each runway averaged 14 departures per hour (roughly the same rate as in the sample project provided with the SIMMOD software). In addition, during the third hour, arrivals of earlier flights competed with the later departures for use of the same runways. The last flights finished landing and unloading approximately seven simulated hours after the earliest flights began.

4.2. Objective Function and Parameter Definition

During each simulation SIMMOD keeps track of the amount of time aircraft are required to hold (due to congestion ahead), both on the taxiways and in the air. This information is given as output so that an objective function value can be calculated. Specifically, we wish to minimize the total cost of delay in the network of airports described above, given the flight schedule and an initial ground-holding policy for the 168 flights. Geisinger of the FAA reported that for the average flight, taxiway delay cost 2.38 times more than air traffic controller-induced gate-holding delay, and airborne delay cost 3.86 times more than gate-holding delay per hour [3]¹. Thus

¹Geisinger reported that the 1986 hourly values for gate-holding delay, taxiway delay, and airborne delay were \$591, \$1,407, and \$2,283, respectively.

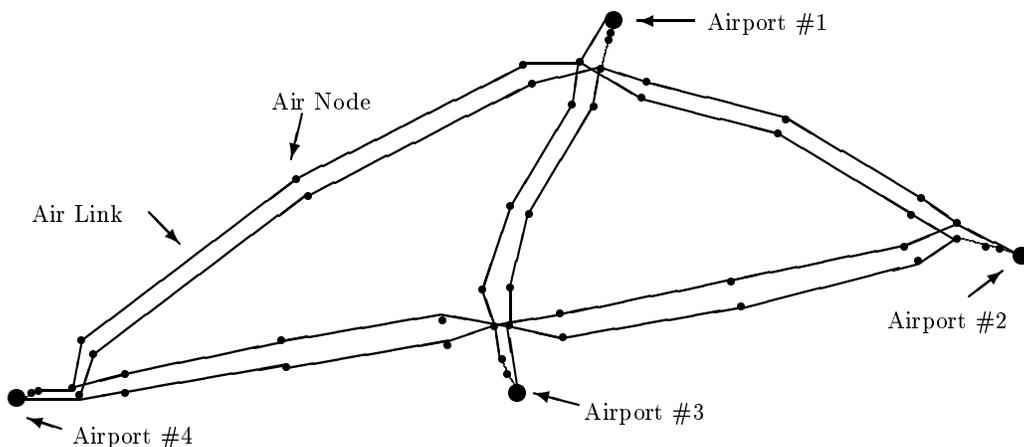


Figure 1: Airspace network

our total delay cost objective function was

$$L(\theta) = m_g(\theta) + 2.38m_t(\theta) + 3.86m_a(\theta) \quad (3)$$

where $m_g(\theta)$, $m_t(\theta)$, and $m_a(\theta)$ are the total number of minutes of gate, taxiway, and airborne delay throughout the 168 flights, $\theta \in R^p$, and θ_i is the number of minutes of air traffic controller-induced gate hold for flight i . Note that each iteration $m_g(\theta_k) = \sum_{i=1}^{168} \theta_{ki}$. Because the θ_{ki} are produced by the SPSA algorithm, m_g is a controlled quantity, whereas the systemic delays m_t and m_a (which depend on θ_k) are obtained as output from the SIMMOD simulations.

4.3. Results

Tests were performed for two different scenarios. In the first scenario normal air link capacities were used (approximately as restrictive as the aircraft separation constraints which required at least five miles between two airplanes on an air link). Here, we started with a ground-holding policy which assigned no ground-holding to each flight, that is, $\theta_{0,i} = 0$, $i = 1, \dots, 168$. Then, using the process described in Section 2., the SPSA optimization algorithm found a ground-holding policy which improved the total delay cost in the system. The first two rows of Table 1 summarized the results for this scenario. For the first row (normal capacity IRN), five 30-iteration runs were performed using independent random number (IRN) seeds to start each of the two simulations every iteration. The average initial objective function value ($L(\hat{\theta}_0)$) was 6,161, and the average final objective value ($L(\hat{\theta}_{30})$) was 5,592, giving a 9.2% reduction. For the second row (normal capacity CRN), five 30-iteration runs were again averaged, but this time the same random number seeds were used to start both simulations in a given iteration. This was done in an attempt to use the common random numbers (CRN) variance reduction technique [7]. Here, the average initial and final objective values were 6,161 and 5,526, for a 10.3% reduction in total delay cost.

In the second scenario, we illustrate how the SPSA optimization method can take a ground-holding policy that is optimal for a normal capacity situation and adjust it so that it is optimal for a situation where capacity is reduced for a time at a particular airport. In our experiment, we used one of the *final* ground-holding policies from the normal capacity CRN case as our *initial* ground-holding policy. We then drastically reduced the capacity of the air links going in and out of one of the airports for a half-hour period to simulate poor weather

Table 1: Objective Value and Reduction Averages from the SPSA Optimization Method

		Initial Objective Value	Final Objective Value	Decrease
Normal Capacity	IRN	6,161	5,592	9.2%
	CRN	6,161	5,526	10.3%
Reduced Capacity	IRN	7,233	6,562	9.3%
	CRN	7,233	7,078	2.1%

conditions. The SPSA algorithm then started with the "initial" ground-holding policy and produced a final policy which yielded lower delay costs for this reduced capacity scenario. The last two rows of Table 1 summarize the results for this scenario. Here, three 45-iteration runs were averaged in both the IRN and CRN cases. In the IRN case the average initial and final objective values were 7,233 and 6,562 for a 9.3% reduction. In the CRN case the average initial and final delay values were 7,233 and 7,078, yielding a 2.1% reduction.

As expected, the delay costs were higher in the reduced capacity scenario than in the normal capacity scenario. Furthermore, in the normal capacity scenario, attempting to use the CRN variance reduction method produced a larger decrease in delay cost, as expected. But in the reduced capacity scenario, the CRN case produced smaller reductions. One explanation may be that more iterations need to be performed to see in practise the theoretical benefits of attempting to use the CRN method. For each run the gain sequences $\{a_k\}$ and $\{c_k\}$ from equations 1 and 2 were defined as $a_k = ak^{-\alpha}$ and $c_k = ck^{-\gamma}$ with $\alpha = 1.0$, $c = 0.05$, and $\gamma = 0.167$. In the normal capacity scenario, a was 1×10^{-6} , and a was 2×10^{-6} in the reduced capacity scenario (see [18] and [7] for a more detailed discussion of the gain sequences).

5. Conclusions

This paper outlines the SPSA optimization algorithm and describes its usefulness as a tool for solving many kinds of air traffic optimization problems. Since the SPSA algorithm requires only noisy measurements of the objective function to be optimized, detailed air traffic simulation software can be used

to accurately model an air traffic network and provide system performance measurements. This combination of a powerful optimization technique and detailed simulation models provides planners with a tool that is potentially much better than current optimization techniques.

To illustrate this method, this paper describes how the SPSA optimization method can be used in conjunction with the SIMMOD air traffic simulation software package to find ground-holding policies which yield improved network-wide delay costs. The results of a test case involving a network of four simple airports are given. Here 9.2% and 10.3% reductions in total delay cost are achieved in the normal capacity scenario (for the IRN and CRN cases), and similarly 9.3% and 2.1% reductions are achieved in a scenario where the initial ground-holding policy (a final ground-holding policy from the normal capacity scenario) is iteratively adjusted by the SPSA algorithm so that it accounts for reduced capacity at a particular airport.

One direction for future study would be to use this optimization method to find optimal ground-holding policies for actual flights in a SIMMOD-simulated network of real airports. Another direction of study would be to consider optimizing other objective functions such as minimizing fuel consumption or to control parameters other than ground-holding times.

Acknowledgements

This work was partially supported by U.S. Navy Contract N00039-95-C-0002 and the JHU/APL IR&D Program. The authors wish to thank Belinda Hargrove of SABRE Decision Technologies for her help with the SIMMOD simulation package.

References

- [1] M. C. Fu and S. D. Hill. Optimization of discrete event systems via simultaneous perturbation stochastic approximation. *Transactions of the Institute of Industrial Engineers, Special Issue: Computer Simulation in Industrial Engineering Research and Practice*, 29(3), Mar. 1997.
- [2] S. Gal, R. Y. Rubinstein, and A. Ziv. On the optimality and efficiency of common random numbers. *Mathematics and Computers in Simulation*, 26:520–512, 1984.
- [3] K. Geisinger. Airline delay: 1976-1986. Technical Report FAA-APO-88-13, Federal Aviation Administration, Washington, DC, Mar. 1989.
- [4] P. Glasserman and D. D. Yao. Some guidelines and guarantees for common random numbers. *Management Science*, 38(6):884–908, June 1992.
- [5] M. P. Helme. Reducing air traffic delay in a space-time network. *IEEE International Conference on Systems, Man, and Cybernetics*, 1:236–242, 1992.
- [6] M. P. Helme, K. Lindsay, S. V. Massimini, and G. Booth. Optimization of traffic flow to minimize delay in the national airspace system. Technical Report MP93W14, The MITRE Corporation, 1993.
- [7] N. L. Kleinman, J. C. Spall, and D. Q. Naiman. Simulation-based optimization with stochastic approximation using common random numbers. *Management Science*, 1996. Submitted.
- [8] M. I. Koch, D. C. Chin, and R. H. Smith. A network-wide approach to optimal signal light timing for integrated transit vehicle and traffic operations. In *Conference Proceedings A of the Seventh National Conference on Light Rail Transit*, volume 2, 1996.
- [9] P. L'Ecuyer and G. Perron. On the convergence rates of IPA and FDC derivative estimators for finite-horizon stochastic simulations. *Operations Research*, 42(4):643–656, 1994.
- [10] P. L'Ecuyer and G. Yin. Budget-dependent convergence rate of stochastic approximation. *SIAM Journal on Optimization*, 1995. Submitted.
- [11] A. R. Odoni. The flow management problem in air traffic control. In A. R. Odoni, L. Bianco, and G. Szego, editors, *Flow Control of Congested Networks*, pages 269–288, Berlin, 1987. Springer-Verlag.
- [12] M. D. Peterson, D. J. Bertsimas, and A. R. Odoni. Models and algorithms for transient queueing congestion at airports. *Management Science*, 41(8):1279–1295, Aug. 1995.
- [13] P. Sadegh and J. C. Spall. Optimal random perturbations for stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control*, 1996. Submitted.
- [14] P. Sadegh and J. C. Spall. Optimal sensor configuration for complex systems. In *Proceedings of the Test Technology Symposium*. U.S. Army Test and Evaluation Command, June 1996. <http://www.atc.army.mil/~tecom/tts/proceed/optsenr.html>.
- [15] R. H. Smith, D. C. Chin, and N. L. Kleinman. Evaluation of real-time, non-network-model-based traffic signal control technique in the presence of long-term trends. In *The Proceedings of the 1995 Annual Meeting of ITS America. Intelligent Transportation: Serving the User Through Deployment*, volume 1, pages 633–638. Intelligent Transportation Society of America, 1995.
- [16] J. C. Spall. A stochastic approximation technique for generating maximum likelihood parameter estimates. *Proceedings of the American Control Conference*, pages 1161–1167, 1987.
- [17] J. C. Spall. A stochastic approximation algorithm for large-dimensional systems in the Kiefer-Wolfowitz setting. *Proceedings of the IEEE Conference on Decision and Control*, pages 1544–1548, 1988.
- [18] J. C. Spall. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control*, 37(3):332–341, 1992.
- [19] J. C. Spall and J. A. Cristion. Nonlinear adaptive control using neural networks: Estimation with a smoothed simultaneous perturbation gradient approximation. *Statistica Sinica*, 4(1):1–27, 1994.
- [20] P. B. M. Vranas, D. Bertsimas, and A. R. Odoni. Dynamic ground-holding policies for a network of airports. *Transportation Science*, 28(4):275–291, 1994.