Proceedings of the 1995 Winter Simulation Conference
ed. C. Alexopoulos, K. Kang, W. R. Lilegdon, and D. Goldsman

EVALUATION OF AN ADAPTIVE TRAFFIC CONTROL TECHNIQUE WITH UNDERLYING SYSTEM CHANGES

Richard H. Smith
Daniel C. Chin

The Johns Hopkins University Applied Physics Laboratory
Johns Hopkins Road
Laurel, Maryland 20723-6099, U.S.A.

ABSTRACT

A key problem in traffic engineering is the optimization of the flow of vehicles through a given road network. Improving the timing of the traffic signals at intersections in the network is generally the most powerful and cost-effective means of achieving this goal. Recent efforts have resulted in the development of a fundamentally different approach for optimal centralized signal timing control that eliminates the need for an open-loop model of the traffic network dynamics. The approach is based on a neural network (NN) serving as the basis for the control law, with the internal NN weight estimation occurring real-time in closed-loop mode via the simultaneous perturbation stochastic approximation (SPSA) algorithm. This paper investigates the application of such a non-network-model-based approach and illustrates the approach through a simulation on a nine-intersection, mid-Manhattan, New York network. The simulated traffic network contains varying short and long-term congestion behavior and short-term stochastic, nonlinear effects. The approach results in a net 10% reduction in vehicle wait time relative to the performance of the existing, in-place strategy.

1 INTRODUCTION

A major component of advanced traffic management for complex road systems is the signal phase timing for the signal-controlled intersections. This is an extremely challenging control problem for most realistic settings. Modern traffic control (i.e., signal timing) algorithms that have been implemented for centralized control of complex networks require an open-loop model for the traffic dynamics which is then used to determine the signal control parameters. In order to accommodate complex network dynamics, this model may take the form of a set of differential/difference equations (Papageorgiou 1990 and Smith and Ghali 1990) or a neural network (Nataksuji and Kaku 1991) or a fuzzy logic or rule-based expert system (Kelsey and Bisset 1993 and Ritchie 1990). Whatever the type of model used, it is serving as a representation of the effect of the current signal timings on the traffic flow in the network. Such an approach, however, will produce suboptimal, centralized control in realistic road networks of many intersections and for time periods of many months duration, since there are numerous unmodelable interactions and seasonally changing effects. In particular: "To develop a 'general theory' for the stochastic behavior of a traffic system is out of the question. Even if it were possible, such a theory would be so complex as to be of no practical value." (Newell 1989, p. 258). Hence, a reduced or non-network-model-based approach to centralized traffic control provides an attractive alternative and has recently been developed in Spall (1992), Spall and Cristion (1994), and Spall and Chin (1994). Spall and Chin (1994) provides a more technical presentation of the algorithm for traffic control while Spall (1992) provides the mathematical basis and Spall and Cristion (1994) develops the integration of the algorithm with neural networks.

The unique aspect of the control strategy in this approach is that it does not require a mathematical (or other) model of the traffic network dynamics (which is typically constructed "off-line" using past traffic data in its determination of the centralized controller signal parameters). It is based on a neural network function approximator for use in the control function, and this processor can obtain its structure directly from traffic system observations rather than from a system network model. This feature eliminates the problem of non-robustness of system model-based controls to operational traffic situations that differ significantly from situations represented in the data used to build the system model (this non-robustness can sometimes lead to unstable system behavior). Further, the non-reliance on network modeling will of course simplify the adaptation of the controller to modifications in the underlying measure of effectiveness (MOE).

The NN/SPSA complex works as a coordinated process to control the traffic. The NN represents a function that transforms recorded traffic conditions into the appropriated traffic signal timing parameters whereas the SPSA algorithm optimizes the weights used in the NN (see Section 2 for details). The NN responds to the short-term variations in traffic and interpolates, or extrapolates from previous responses if necessary, to produce appropriate timing parameters for nontypical traffic incidents. The SPSA algorithm smooths over stochastic variations in traffic during its NN weight estimation, which occurs over a period of several days, and tolerates inaccuracies in MOE estimation since it accounts for noise in traffic flow data (Spall 1992).

In addition to the above considerations, true intelligent control requires that the controller automatically adapt to the inevitable long-term (say, month-to-month) changes in the system. This is a formidable requirement for the current model-based controllers as these long-term changes encompass difficult-to-model aspects such as seasonal variations in flow patterns, long-term construction blockages, changes in the number of residences and/or businesses in the system, etc. In fact, in the context of the Los Angeles traffic system, Rowe (1991) notes that the difficulty in adapting to long-term changes is a major limitation of current traffic control strategies. The non-network-model-based approach, however, is able to produce a controller that converges toward optimal, centralized, instantaneous (cycle-by-cycle) signal timings while automatically adapting to long-term (month-to-month) system changes. Such an approach could thus be incorporated by advanced traffic control planners to reduce the cost and need for frequent traffic control strategy upgrades in a traffic system.

The control strategy here (like any other demand-responsive controller) requires real-time sensor data related to the traffic flow. In some cases, the MOE of interest can be formulated directly in terms of the sensor data, e.g., an MOE measuring vehicles/unit time passing through the network intersections can be calculated directly from common loop detectors at the intersections that provide vehicle counts. In other cases, the MOE may involve quantities not directly related to the available sensors, e.g., an MOE that reflects total vehicle wait time near intersections cannot be determined straightforwardly if only upstream loop detector data are available. In such cases, some very localized (decentralized, link-specific) modeling would be required to relate the sensor data to the MOE (this requirement, of course, applies to any control technique). This modeling, however, is usually much simpler than attempting to model the underlying traffic dynamics that relate the <u>signal timings</u> to the MOE at a network-wide level (as discussed earlier). The reason for this relative simplicity is that the relationship between the sensor data and MOE is typically much more direct, short-term, and localized to specific traffic queues than the effect of a set signal timings from multiple signals on the network-wide traffic flow and associated MOE. For example, loop detectors near an intersection can provide data for reliable estimation of vehicle wait time at the intersection (see, e.g., Wallace, Courage, and Hadi 1991, Section 4.3.4.1, or Tsay, Kang, and Hsiao 1991); these estimated wait times can then be summed to provide the estimated network-wide wait time. Although the means of evaluating MOEs must be addressed, the approach presented here is not specific to a particular MOE. Our approach is beneficial in the sense of reducing the modeling effort, if required, to only the sensor-MOE relationship at most to accomplish system-wide traffic control.

2 OVERVIEW OF CONTROL STRATEGY

The fundamental non-network-model-based and system-wide control strategy develops, at the outset, a general mathematical function that takes any current data values on the state of the traffic conditions (e.g., instantaneous upstream traffic counts and perhaps speed) and produces a set of signal timings (length of green phase or split, cycle length, etc., for the subsequent light cycle of all controlled signals) in such a way as to optimize the performance of the system over a period of several hours. This control function is implemented in this approach by a neural network (NN), which is a powerful technique for approximating complex nonlinear functions such as the "true," but unknown, optimal control for signal timings. Essential to the performance of the NN are the values of the connection weights in the NN. The major unique feature of the non-network-model-based approach is the simultaneous perturbation stochastic approximation (SPSA) method by which these NN-controller weights are estimated ("trained") by observing the effect on the traffic of small NN weight perturbations and the resulting small signal timing changes. Figure 1 illustrates the overall relationship between the NN control, the traffic system to be controlled, and the SPSA training process. The SPSA method is specifically designed to accommodate stochastic traffic fluctuations encountered during the training process.

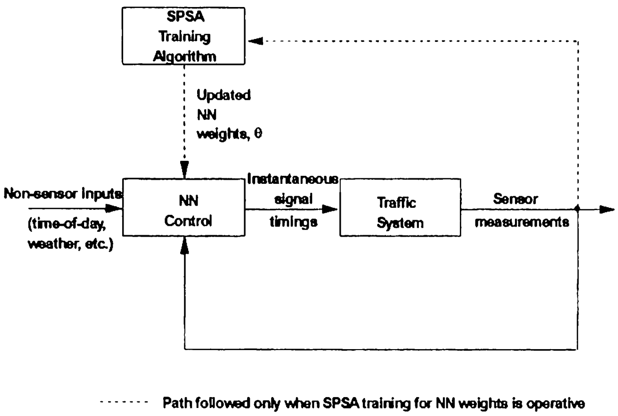

Figure 1: Overall Relationship Between Traffic System and NN/SPSA Controller

It is sufficient to view the NN as a certain collection of nonlinear functions with unknown coefficients (called connection weights) that must be estimated. After the weights are specified, the NN acts as a function by taking a set of inputs and producing a set of outputs (analogous to a polynomial with specified coefficients). Theory given in, say, Funahashi (1989) shows that any reasonable mathematical function can be approximated to a high level of accuracy by a NN if (and only if) the weights are properly estimated. In the present case, the NN is being used to approximate the (unknown) optimal control function for the signal timings. Thus, although there is a wealth of literature on NNs in many types of applications, the focus here is simply in using them as a function approximator. The SPSA methodology, to be summarized in the next section, then proceeds to train and continually and adaptively adjust the NN weights. This is done in such a way that the resulting system-wide signal controls are responsive and converge toward optimality for the variations of traffic flow patterns encountered during the immediate daily period of operation.

3 ALGORITHM FOR DETERMINING NN WEIGHTS IN CONTROLLER

As discussed above, the NN-based control depends on a set of weight coefficients, which must be estimated. Once these weights are properly specified, there will be a fully defined function that will take state information on traffic conditions (e.g., traffic counts and perhaps speed throughout the traffic system) at any given time of day and collectively produce optimized signal timings for the subsequent cycle of each traffic signal. It is within these weights that information about the optimal control strategy is embedded. To reflect reality, it is important that the weights contain short-term, time-of-day information and that they be able to evolve in the long-term (month-to-month) in accordance with the inevitable changes in the transportation system. The SPSA algorithm for weight estimation is based on observing traffic patterns when slight changes are made to the weights and thus to the current signal timing settings; these "slight changes" are done in a special way that is at the heart of the algorithm and that maximizes the amount of information available for optimization while minimizing disruptions to the traffic system.

A step-by-step summary of how the SPSA algorithm would be implemented, to achieve optimal traffic control, will now be given. A more detailed discussion of the mathematical basis and procedure for implementing the SPSA technique for traffic control is presented in Spall and Chin (1994) and in Spall and Cristion (1994) for more general control applications.

1. An initial weight vector for the NN controller is constructed first by implementing the NN in an open-loop and so-called "back-propagation" mode using several days of traffic count and current (suboptimal) traffic signal controller data. The data can be supplied either by a traffic engineering database or by a realistic computer simulation of the traffic network to be controlled. The simulation should employ a reasonable (though suboptimal) control strategy for all of the signals in the traffic network being addressed. After this initialization phase, the NN controller should be able to approximately emulate the initial real or prescribed control strategies. We should emphasize here that the computer simulation suggested above is for initialization purposes only, and its operation and employed strategies are not critical to the optimal strategies to be derived subsequently by the SPSA algorithm.
2. Given the current weight vector estimate, change all values slightly in the manner described in Spall and Chin (1994) and Spall and Cristion (1994).
3. Monitor the traffic system throughout the control time period (e.g., several hours) and form the sample loss function (sum of prescribed MOEs) based on the observed system behavior.

4. During the same control time period on the following like day (e.g., possibly weekday-by-weekday), repeat steps 2 and 3 with a complementary slight change in the weight vector estimate, as described in Spall and Chin (1994) and Spall and Cristion (1994).
5. With information from steps 3 and 4 on separately calculated loss functions, take one iteration of the SPSA algorithm, which updates the values of the elements of the weight vector.
6. Repeat steps 2-5 with the new weight values until the traffic flow is optimized (i.e., shows convergence to a maximum or minimum as desired) based on the chosen loss function.

There are several practical aspects of the above procedure that are worth noting. First, since each iteration of SPSA requires two days, we would expect that adequate convergence to the optimal weights for the traffic system would take a month or two. While this real-time adaptive optimization or training is taking place, the controls will not, of course, be optimal. Nevertheless, by initializing the weight vector at a value that is able to produce the initial signal timings actually in the system (or in a reliable simulation), the algorithm will tend to produce signal timings that are between the initial and optimal timings while it is in the training phase. Hence, there should be no significant control-induced disruption in the traffic system during the training phase. After the converged weight values have been obtained, we will have derived a NN-based control processor that produces optimized light timings for any given intra-light cycle traffic flow conditions and for all signals in the traffic system. In order to adapt to the inevitable long-term changes in the underlying traffic flow patterns, the controller may be allowed to continue in its training mode operation indefinitely.

4 EXAMPLE OF SPSA-BASED SYSTEM-WIDE ADAPTIVE TRAFFIC CONTROL APPROACH

4.1 Introduction

This section illustrates by simulation an application of the SPSA real-time adaptive control approach for system-wide traffic signal control described above. The six-step training process outlined in Section 3 is employed here to construct a NN-based controller. In particular, we are considering control for one four-hour time period and are estimating, across days, the NN weights for the collective set of traffic signal responses to instantaneous traffic conditions during this four-hour period. The software used here is a modified version of SPSA control software that was originally designed for a smaller scale traffic system problem (see Spall and Chin (1994)); the simulation was conducted on an IBM 386 PC; and the software is written in the programming language C++. The traffic dynamics were simulated using state-space flow equations similar to those in Papageorgiou (1990) or Nataksuji and Kaku (1991) with Poisson-distributed vehicle arrivals at input nodes. Of course, as described throughout this paper, the controller does not have knowledge of the equations being used to generate the simulated traffic flows, but it is able to adapt to the system by efficient use of small system changes and observation of resulting system performance.

4.2 The Simulated Traffic Flow and Form for NN Controller

The studies conducted here are based on the simulation and test case treated in Chin and Smith (1994) (i.e., a mid-Manhattan, NY business sector). Two studies were conducted for a simulated 90-day period: one which implemented steady, long-term growth in Poisson arrival rates over the total period, and another which introduced a 10% step increase in arrival rates at day 10 during the total period. Both are realistic and difficult control scenarios; the long-term change might represent a new office complex development and the step change might represent a single business opening. For the long-term change, all input queue arrival rates experience a net increase of .08% of the original level per day for each of the 90 days in the study or a total increase of 7.2% after 90 days. In both studies the simulated traffic network runs between 55th and 57th Streets (North and South) and from 6th Avenue to Madison Avenue (East and West) and therefore includes nine intersections with 5th Avenue as the central artery. Figure 2 depicts the scenario. The time of control covers the four-hour period, from 3:30 p.m. to 7:30 p.m., which represents evening rush hour. The technique could obviously be applied to any other period during the day as well. In the four-hour control period several streets have their traffic levels gradually rising and then falling. Their traffic arrival rates increase linearly from non-rush hour rates starting at 3:30 p.m. The rates peak at 5:30 p.m. to a rush hour saturated flow condition and then subside linearly until 7:30 p.m. Back-up occurs during rush hour. Nonlinear, flow-dependent driver behavioral aspects are embedded in the simulation. Some streets have unchanging traffic statistics during the total time period while others have inflow rates from garage-generated egress at the end of office hours from 4:30 p.m. to 5:30 p.m. The simulation has been extensively tested to ensure that it produces traffic volumes that correspond to actual recorded data for the Manhattan traffic sector. (A complete discussion of the development and testing of the baseline simulation and the details of its operation are given in Chin and Smith (1994)).

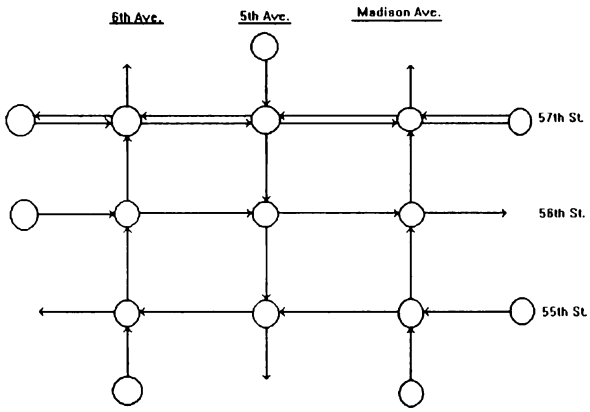

Figure 2: Traffic Simulation Control Area (Mid-Manhattan)

For the controller, we used a two-hidden-layer, feed-forward NN with 42 input nodes. The 42 NN inputs were (i) the queue levels, at each cycle termination, for the 21 traffic queues in the simulation, (ii) the per-cycle vehicle arrivals at the 11 external nodes in the system, (iii) the time from the start of the simulation, and (iv) the 9 outputs from the previous control solution. The output layer had 9 nodes, one for each signal light's green/red split. The two hidden layers had 12 and 10 nodes, respectively. For this NN, there were a total of 745 NN weights that must be estimated as part of the control algorithm.

The controller is designed to map, in a centralized fashion, the above listed information into optimal values for the green-time/total-cycle-time for the succeeding cycle of each of the nine signals in the traffic network. The controller operates in a real-time adaptive mode in which its cycle-by-cycle responses to traffic fluctuations are gradually improved, over a period of several days or weeks, based on an MOE consisting of the calculated total traffic system wait time over the daily four-hour period. Note that since the underlying MOE for the NN controller weight estimation is based on system-wide traffic data (i.e., data downstream from each traffic signal as well as upstream) over a several-hour time period, the effect of signal settings, turning movements, etc. several minutes into the future, after each cycle, and the future accumulation of traffic at internal queues is factored into the formation of the controller function.

Although not included in the current study, the calculation of the MOE in an actual implementation would require upstream traffic sensor data and a simple model that relates the upstream data to the downstream queue development and vehicle wait time for each queue, independently. As discussed in the Introduction, this is viewed as a much simpler requirement than modeling sensor/queue interactions collectively for the entire traffic network, which the SPSA methodology does not require. Of course, the placement of additional sensors at intersections could eliminate the modeling effort. The current example demonstrates the capability of the SPSA algorithm to collectively and effectively deliver an adaptive control of signals for an entire traffic network with nonlinear and stochastic traffic flows (and either slowly changing or step increasing demand) using a network-wide battery of point-based traffic data inputs. Consistent with the SPSA methodology, the controller output timings change continuously (cycle-to-cycle) as a function of the instantaneous input traffic flow data (i.e., resulting from Poisson samples) while the underlying NN weights that define the control function are changed on a day-to-day basis in a gradually adaptive training process. The adaptive process was operated continuously over the 90-day span of the simulation while the Poisson arrival rates of the traffic in the network were undergoing the added feature of net gradual changes as described above.

4.3 Results

Figure 3 presents the results of our simulation study of the system-wide traffic control algorithm for long-term changes. Figure 4 shows similar results for the step increase case. In order to show true learning effects (and not just random chance as from a single realization) the curves in Figures 3 and 4 are based on an average of 100 statistically independent simulations. In Figure 3, the middle curve displays the normalized average values of total system (nine intersection) wait time over the four-hour period of each day for the 100 simulations, and the upper and lower curves are the bounds within which 90% of the sample data lies. The values in the curves are normalized by the average response of an equivalent traffic network controlled by a fixed signal split strategy. The fixed strategy assumed a green-time/total-cycle-time value of .55 for all signals along N-S arteries. This was the prior strategy in-place in the Manhattan sector during the recording of actual data. As evident in Figure 3, 30 days of the 90-day period were reserved as optional "evaluation days" to demonstrate improved values of the MOE. However, only data from the other 60 "training days" were used in the SPSA algorithm; thus, the adaptive training period could have been reduced to only 60 days. In this case, the improvement would have been more pronounced since the long-term change would have had less time to evolve.

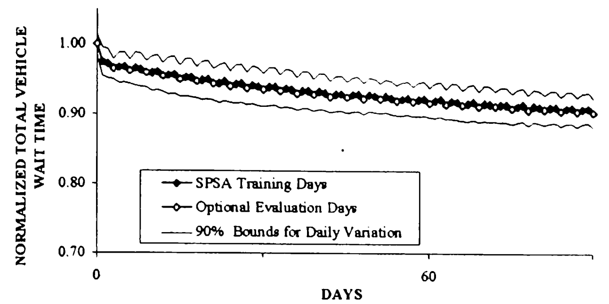

Figure 3: Reduction of Total Wait Time for NN/SPSA Control of Traffic Signal Timing with Long-Term Trend

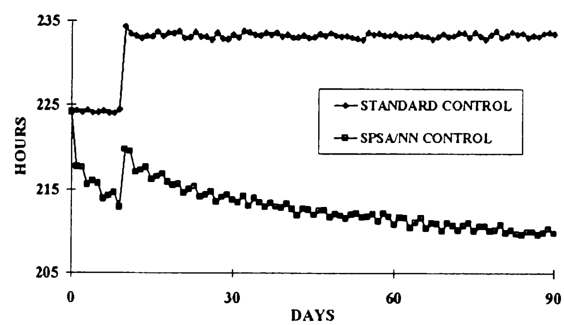

Figure 4: Reduction of Total Time for NN/SPSA Control of Signal Timing with Step Increase in Flow

With the fixed-time strategy in use in the baseline system, the total system wait time experienced a 5.4% approximately linear increase over the 90-day period, on the average, as a result of the 7.2% total linear increase in traffic arrival rates. For the same system under control of the SPSA-based algorithm, the total system wait time showed a decaying exponential decrease in spite of the linearly increasing arrival rates. After the 90-day period, the SPSA-controlled system showed an average decrease in total system wait time of 4.6%. Normalization of the SPSA-controlled system data by the fixed-time strategy data resulted (see Figure 3) in a net improvement for the SPSA-controlled system of 10% relative to the fixed-strategy-controlled system. This reduction in total wait time represents a reasonably large savings with a relatively small investment, particularly for high traffic density sectors. In comparison, major construction changes to achieve a net improvement in traffic flow of 10% in a well-developed area, such as for the traffic system in mid-Manhattan, would be enormously expensive.

In the step increase case, Figure 4 shows a corresponding step increase in total system wait time under the fixed-time strategy. For the traffic system under SPSA control, a step increase also occurred in total system wait time, but the wait time continued to decrease without any transient behavior subsequent to this phenomenon, and a 10% improvement is evident after the 90-day test period. This result demonstrates the broad robustness properties of the SPSA technique. Furthermore, the total system wait time in both studies displays a generally downward trend; given more time, the system is expected to show continuing gradual improvement.

5. CONCLUDING REMARKS

It was shown that the SPSA optimization technique can be used for practical real-time, long-term adaptive, centralized, non-network-model-based control for a moderate-sized traffic system. This was achieved by using the SPSA algorithm to determine the weights for a neural network controller that produces the optimized signal timings collectively for each signal cycle within a given time period based on observed intra-cycle traffic conditions. Thus, this controller will operate in real-time and make signal timing adjustments for multiple signals in a traffic network to accommodate short-term conditions such as congestion, business openings, nearby accidents, brief construction blockages, adverse weather, etc. It also has the ability to automatically accommodate long-term system changes (such as seasonal traffic variations, new residences or businesses, long-term construction projects, etc.) without the cumbersome and expensive off-line remodeling process that has been customary in traffic control. The NN training process gradually adapts to long-term changes in a manner that is essentially invisible to the vehicle operators in the system and produces an optimized control on a system-wide basis.

ACKNOWLEDGEMENTS

This work was partially supported by The Johns Hopkins University Applied Physics Laboratory Independent Research and Development Program.

REFERENCES

Chin, D. C., and R. H. Smith. 1994. A traffic simulation for mid-Manhattan with model-free adaptive signal control. In *Proceedings of the 1994 Summer Computer Simulation Conference*, ed. D. K. Pace and A. Fayek, 296-301. The Society for Computer Simulation, San Diego, California.

Funahashi, K. I. 1989. On the approximate realization of continuous mappings by neural networks. *Neural Nets* 2:183-192.

Kelsey, R. L., and K. R. Bisset. 1993. Simulation of traffic flow and control using fuzzy and conventional methods. In *Fuzzy logic and control*, ed. M. Jamshidi, et al., Chapter 12. Englewood Cliffs, NJ: Prentis Hall.

Nataksuji, T., and T. Kaku. 1991. Development of a self-organizing traffic control system using neural network models. *Transportation Research Record* 1324, 137-145. TRB, National Research Council, Washington, D.C.

Newell, G. F. 1989. *Theory of highway traffic signals.* Institute of Transportation Studies, Univ. of California, Berkeley.

Papageorgiou, M. 1990. Dynamic modeling, assignment and route guidance in traffic networks. *Transportation Research-B* 24B:471-495.

Ritchie, S. G. 1990. A knowledge-based decision support architecture for advanced traffic management. *Transportation Research-A* 24A:27-37.

Rowe, E. 1991. The Los Angeles automatic traffic surveillance and control system. *IEEE Transactions on Vehicular Technology* 40:16-20.

Smith, M. J. and M. Ghali. 1990. The dynamics of traffic assignment and traffic control: A theoretical study. *Transportation Research-B* 24B:409-422.

Spall, J. C. 1992. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control* 37:332-341.

Spall, J. C. and J. A. Cristion. 1994. Nonlinear adaptive control using neural networks: estimation with a smoothed simultaneous perturbation gradient approximation. *Statistica Sinica* 4:1-27.

Spall, J. C. and D. C. Chin. 1994. A model-free approach to optimal signal timing for system-wide traffic control. In *Proceedings of the 1994 IEEE Conference on Decision and Control*, ed. M. Peshkin, 1868-1875. Institute of Electrical and Electronics Engineers, Lake Buena Vista, Florida.

Tsay, H-S, J-F Kang, and C-H Hsiao. 1991. Algorithm for estimating queue lengths and stop delays at signalized intersections. *Transportation Research Record* 1324, 123-129. TRB, National Research Council,Washington, D.C.

Wallace, C. E., K. G. Courage, and M. A. Hadi. 1991. *TRANSYT-7F user's guide (Methodology for optimizing signal timing, MOST 4),* Gainesville, Florida: Courage and Wallace.

AUTHOR BIOGRAPHIES

RICHARD H. SMITH is a member of the Principal Professional Staff in the Strategic Systems Department of The Johns Hopkins University Applied Physics Laboratory. He received his Ph.D. in electrical engineering from The Johns Hopkins University. His research interests include optimization, modern control systems, statistical analysis, modeling, and simulation. He is a member of IEEE, Tau Beta Pi, and an affiliate of the Transportation Research Board.

DANIEL C. CHIN is a Senior Staff Mathematician in the Strategic Systems Department of The Johns Hopkins University Applied Physics Laboratory. He received his M.S. in mathematics from Northern Illinois University and has additional, extensive graduate study and experience in computer science and simulation. His research interests include stochastic approximation, neural networks, optimization of stochastic systems and traffic control and simulation. He is a member of Sigma Xi.