

Risk-Averse Group Elevator Scheduling

Matthew Brand, Daniel Nikovski

TR2004-045 April 27, 2004

Abstract

We introduce a novel group elevator scheduler based on an explicit decision-theoretic calculation for the expectation of any polynomial function of passengers' waiting times, marginalized over sources of uncertainty in the system state. The same framework can be used to identify risky assignments that have low expected costs but may cause some individuals to be "stranded" with excessive waits for service. The resulting scheduler avoids risky assignments while minimizing average passenger waits. Experimental results demonstrate that the method can reduce the variance of waiting times and the fraction of passengers waiting excessively in comparison with risk-neutral schedulers, while still achieving better waiting times than reference ETA controllers.

International Congress on Vertical Transportation Technologies (ELEVCON) 2004

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Risk-Averse Group Elevator Scheduling

Matthew Brand and Daniel Nikovski
Mitsubishi Electric Research Laboratories, USA

Key Words: group elevator control, decision-theoretic scheduling, Markov chains, dynamic programming

ABSTRACT

We introduce a novel group elevator scheduler based on an explicit decision-theoretic calculation for the expectation of any polynomial function of passengers' waiting times, marginalized over sources of uncertainty in the system state. The same framework can be used to identify risky assignments that have low expected costs but may cause some individuals to be "stranded" with excessive waits for service. The resulting scheduler avoids risky assignments while minimizing average passenger waits. Experimental results demonstrate that the method can reduce the variance of waiting times and the fraction of passengers waiting excessively in comparison with risk-neutral schedulers, while still achieving better waiting times than reference ETA controllers.

1. INTRODUCTION

Elevator passengers expect efficient service within acceptable time limits, and it is the job of a group elevator scheduler to provide such service when an elevator bank is installed in a building. Exactly what is understood by "efficient" service is open to interpretation. Ultimately it reduces to a matter of subjective preferences of elevator passengers (and of the building developer). This paper examines mathematical models for passengers' aversion to excessively long waits, and develops computationally efficient methods that can be employed to lower the risk of such waits while maintaining an overall high quality of service.

The average waiting time (AWT) of passengers is by far the most common optimization criterion for quality of service in group elevator control (GEC) (Barney 2003). The implicit assumption behind minimizing AWT is that passengers' (dis)satisfaction from service is proportional to the AWT: e.g., if passengers are made to wait twice as long, they will be twice as dissatisfied. Related criteria such as the average journey time (AJT) and the average system time ($AST=AWT+AJT$) are also occasionally used.

Methods for reducing excessively long waits usually revolve around manipulating some statistics that characterize the distribution of waits. It is easy to see that minimizing AWT has also the implicit effect of minimizing the fraction of people waiting more than a specified threshold, *if* shape of the distribution of waiting times does not change. For example, if this

distribution were exponential, as has been suggested for classical traffic control systems (Halpern 1992), then decreasing its mean while preserving its exponential nature would reduce the fraction of long waits, because the shape of the distribution is wholly determined by the mean.

In practice the distribution of waiting times is more complicated, and must be characterized by a mean, a variance, and possibly higher-order moments. This presents many ways to shift the distribution to reduce the fraction $f(\theta)$ of passengers who wait more than some threshold θ , for example, holding the mean constant and reducing the variance (Halpern 1992; Barney 2003), and directly minimizing the second moment (SSWT; sum of squared waiting times).

Minimizing SSWT is the aim of the DLB, HUFF, and FIM algorithms due to Bao et al. (1994) for the case of down-peak traffic. The reinforcement-learning algorithm due to Crites and Barto (1998) also used quadratic waits, but exponentially discounted in the future. These algorithms are based on the observation that it is much more plausible that passengers' dissatisfaction with elevator service is not directly proportional to AWT (respectively, AJT or AST), but can be represented better by some other (possibly non-linear) function of it. In particular, user dissatisfaction is likely to grow *supralinearly* in AWT, perhaps exponentially or as some power of AWT greater than 1. Such optimization criteria have the direct effect of penalizing long waits much more than short waits, thus implicitly discouraging long waits and tilting the distribution of waits towards lower values.

The use of these two strategies – minimization of variance and the use of supralinear optimization criteria – is known under the common name of *risk-averse control* (Bertsekas 2000). Section 2 analyzes the specific application of risk-averse control (RAC) to the field of GEC, and discusses the significant computational difficulties encountered upon attempting to modify one of the known scheduling methods for the purposes of RAC. This section also shows that many instances of risk-averse control can be reduced to computing higher statistical moments of the AWT with respect to the uncertainty in passenger destinations. Section 3 presents the main contribution of this paper: an efficient decision-theoretic algorithm for the computation of these moments. It also provides a solution to the related risk-averse control problem of minimizing AWT under the (soft) constraint that no passenger's expected waiting time should exceed a pre-specified value. Section 4 shows some experimental results from the execution of these algorithms in a detailed simulator, and discusses their performance with respect to widely known risk-neutral benchmark controllers. Section 5 closes the paper with some of the remaining open questions and challenges in risk-averse group elevator scheduling.

2. RISK-AVERSE CONTROL OF ELEVATOR BANKS

As discussed above, two of the main approaches to risk-averse control are based on minimization of the variance of waiting time and/or minimization of a superlinear function of AWT. This section focuses on the practical difficulties arising in the computation of variances and non-linear functions of waiting times, and demonstrates how both approaches reduce to the computation of the statistical moments of AWT.

Variance-Based Strategies

There are three often used strategies for variance optimization (Puterman 1994). One option is to make decisions that minimize variance subject to the constraint that the mean cost be no greater than a specified threshold. The converse idea is to make decisions that minimize cost subject to the constraint that the variance be no greater than a specified threshold. Finally, a third idea is to form a new decision criterion that is a weighted average of mean and variance, with user-specified weights.

In GEC, the cost to be minimized is usually AWT. Halpern (1993) described a scheduler based on the last strategy, but replaced the waiting times variance estimate with an estimate of the longest hall-call waiting time, on the presumption that estimating the variance directly is computationally infeasible. Indeed, the difficulty of estimating the waiting times variance has been a major impediment to the more common use of risk-averse controllers — providing a solution to this problem is one of the main objectives of this paper.

The reason why it is hard to estimate variances is that traditional GEC computational methods almost universally ignore the probabilistic nature of passengers' waiting times, which is equivalent to assuming that the variance of these times is zero—a completely wrong and unintended consequence. To illustrate it, consider the decision-making process in a well-known scheduling algorithm such as Estimated Time to Arrival (ETA), many variants of which have been employed in the elevator industry since the algorithm's inception in the 1970s. The objective of ETA is to assign a car to a new hall call in a manner that minimizes the AWT of all existing and future passengers from the current moment on. In order to deal with the uncertainty in destination floors arising from the use of two-button hall-call panels, ETA makes a major simplification and assumes that the elevator car will follow a single, deterministic trajectory. In reality, the actual path followed by a car (given its current passenger pick-up commitments) is stochastic, because the intended destination floors of each waiting passenger are unknown random variables. In general, the number of possible future car paths is exponential in the number of passengers waiting in the halls or building floors, whichever is less. In order to combat this combinatorial explosion of paths to be considered, the ETA algorithm analyzes only a single path, assuming that it will happen with probability one, and ignores other possible paths. For definiteness, each waiting passenger is assigned a proxy destination floor (heuristics vary), and the waiting times of known hall calls are then computed along the resulting proxy single path. The average such waiting time is taken as a surrogate measure of the expected waiting time of all passengers to be transported by this car.

Aside from the very questionable accuracy of such a simplification, another unintended effect is that the assumption for a single car trajectory entails that the variance of the individual waiting times is zero. Indeed, if there is no variability in the time necessary to reach each floor, the variance of waiting times should be zero – an obviously incorrect conclusion. What is necessary, then, is a method that does not ignore blindly the variability in car trajectories, but accounts properly for their probabilistic nature. In practice, what would suffice is an algorithm to compute the statistical moments of waiting times, i.e., the mathematical expectation of the first and second powers of individual waiting times with respect to the uncertainty in passengers destinations. More formally, we are considering a particular car i and a set of hall calls h_{ij} , $j = 1, N_i$ currently existing at floors a_{ij} and assigned to this car, plus another a set of existing car calls already constraining the motion of the car in a known

manner. Let the corresponding (unknown) destination floors of calls h_{ij} be denoted by the variables d_{ij} , and let the vector variable $D_i = [d_{i1}, d_{i1}, \dots, d_{iN_i}]$ denote a particular valid set of possible destinations for all hall calls assigned to car i . Once the destinations of passengers are fixed, the waiting time to answer call h_{ij} can be determined precisely. In order to reflect the dependency of waiting time on D_i , we will denote it by $w_{ij}(D_i)$. When w_{ij} is used without D_i , we will assume it to be a random variable denoting the waiting time for hall call h_{ij} .

The objective is to estimate the *expected* waiting time $\mu_{ij} = E_{D_i}[w_{ij}]$ until the car reaches each of the floors a_{ij} , counted from the time hall call h_{ij} actually occurred, and the *expected* squared waiting time $s_{ij} = E_{D_i}[w_{ij}^2]$. The expectation $E_{D_i}[\cdot]$ is taken with respect to all possible sets of destinations D_i , weighted according to their respective probabilities $\Pr(D_i)$. Note that the expected squared waiting time s_{ij} is *not* equal to the square of the expected waiting time μ_{ij}^2 , because expectation and raising to the second power do not commute. Instead, their difference $\sigma_{ij}^2 = s_{ij} - \mu_{ij}^2$ is exactly the variance of the waiting time for call h_{ij} , since the variance of any random variable is equal to the difference between its second uncentered moment and the square of its first moment (the mean). If μ_{ij} and s_{ij} can be found for each car i for a particular assignment of calls to cars, the estimated AWT μ for answering all outstanding hall calls can then be computed as $\mu = \frac{1}{N} \sum_{i=1}^C \sum_{j=1}^{N_i} \mu_{ij}$, where C is the total number of cars, and $N = \sum_i^C N_i$ is the total number of outstanding hall calls. If desired, scheduling decisions can be risk-neutral and can be based only on the AWT μ resulting from each possible scheduling decision, or the individual variances of waiting times σ_{ij} can be included in the optimization criterion as well, to encourage risk-averse decisions.

Supralinear Optimization Functions

The other described method for risk-averse control, that of optimizing superlinear functions of waiting time, can also be reduced to the computation of the statistical moments of waiting time. When the optimization function is the expected squared waiting time $S = \frac{1}{N} \sum_{i=1}^C \sum_{j=1}^{N_i} s_{ij}$, the connection is obvious.. Notice that the alternative optimization criterion $V = \frac{1}{N} \sum_{i=1}^C \sum_{j=1}^{N_i} \mu_{ij}^2$ (squared expected waiting times) *does not* optimize the expected squared waiting time, though it is often used in hopes of penalizing and suppressing excessive waits. Because S equals V plus the variances, minimizing V puts no direct pressure on the variance of waiting times—in principle one can reduce V and yet incur some spectacularly long waits. Yet many algorithms use V instead of S , e.g. those in (Bao et al., 1994), because it is easy to make an (admittedly imprecise) estimate of V from the surrogate estimates of individual waiting times μ_{ij} produced by ETA on a single path. In contrast, there has been no efficient way to compute S until now.

Computing and minimizing other non-linear functions of waiting time is more complicated, but only a little more so. In general, any continuous differentiable nonlinear function $g(w)$ of waiting time w on a fixed interval can be approximated arbitrarily well by means of polynomials. A Taylor-series expansion around a particular point, for example zero, is one of the simplest approximations. It is also well known that Chebyshev polynomials of degree n provide the best approximation accuracy in a least-squares sense on a fixed interval among all polynomials of degree n , and Chebyshev coefficients are easy to compute (Press et al. 1992). Provided that such an approximation has been performed, the function can be represented as

$g(w) \approx \sum_{k=0}^n b_k w^k$. The actual optimization criterion then is:

$$G = E_D \left[\sum_{i=1}^C \sum_{j=1}^{N_i} g(w_{ij}) \right] = E_D \left[\sum_{i=1}^C \sum_{j=1}^{N_i} \sum_{k=0}^n b_k w_{ij}^k \right] = \sum_{k=0}^n b_k \sum_{i=1}^C \sum_{j=1}^{N_i} E_{D_i} [w_{ij}^k] = \sum_{k=0}^n b_k \sum_{i=1}^C \sum_{j=1}^{N_i} m_{ij}^{(k)} = \sum_{k=0}^n b_k m^{(k)},$$

where the expectation is taken with respect to the uncertainty in the system (hall-call destinations) for *all* shafts, represented by $D = [D_1, D_2, \dots, D_C]$, and $m_{ij}^{(k)} = E_{D_i} [w_{ij}^k]$ is the k -th raw (uncentered) moment of the waiting time of hall call h_{ij} , with $m_{ij}^{(1)} = \mu_{ij}$ being the first moment, $m_{ij}^{(2)} = s_{ij}$ being the second, etc.

3. AN ALGORITHM FOR COMPUTING MOMENTS OF WAITING TIME

The previous section demonstrated that the optimization of common performance objectives for the purposes of risk-averse GEC can be reduced to the computation of several statistical moments of waiting time. In short, we need to compute the expected values of the first few powers of waiting time with respect to the uncertainty in the system, then make decisions that minimize their weighted sum. The central problem, then, becomes how to compute these moments efficiently, avoiding the combinatorial explosion that would result from an explicit enumeration of all possible paths an elevator car can take in order to serve the known origin and unknown destination floors of the passengers assigned to it.

The key to solving this problem efficiently is to divide the trajectory of the car into stages and use the method of dynamic programming to compute the cost-to-go for each stage. In the case of expected waiting time, the cost-to-go is the increase in waiting times incurred from any particular stage onwards. The key observation is that there may be many different system paths to the same stage, but the cost-to-go is invariant to how the system got there. This is called the Markov property. A trivial example in the elevator domain is the point at which a car has emptied out and changed direction—how it got to be that way is immaterial to the remaining wait times of the people it is about to pick up. The essence of dynamic programming is that these cost-to-go values can be cached and reused so efficiently that the exponentially large set of possible system trajectories can be evaluated in linear time. Nikovski and Brand (2003) devised such an algorithm to embed the elevator system in a Markov chain and compute expected waiting times considering all possible passenger destinations and car paths. We refer readers to the original paper for details; here we develop an extension to compute arbitrary higher moments of waiting time.

The essential step of the DP algorithm, called a Bellman back-up, computes the cost-to-go of a particular state as the expected value of the sum of the direct costs on all outgoing

transitions and the costs-to-go of the successor states at the end of these transitions. For example, when the objective is to compute the expected waiting time, the cost-to-go $W(s)$ of a state s has the meaning of residual cumulative waiting time from that state on for all passengers h among those in the set $H(s)$ who are still waiting in s : $W(s) = \sum_{h \in H(s)} E[w_h]$. In this

case, the Bellman back-up has the form $W(s) = \sum_{s'} \Pr(s, s') [C(s, s') + W(s')]$, where $W(s)$ is

the cost-to-go of the current state s , the sum ranges over all possible successor states s' that the system transitions to with respective probabilities $\Pr(s, s')$ and immediate costs $C(s, s')$. The probabilities are determined by the number of people inside the car and their likelihood to stop at the next floor vs. going on. The immediate cost of a transition (segment) is simply the duration $\Delta w(s, s')$ of this segment in seconds, times the number $|H(s)|$ of people still waiting at the current state s . After the costs-to-go of all states are backed-up as described, $W(s_0)$ of the initial state s_0 of the Markov chain is equal to the expected cumulative waiting time of all passengers assigned to the car.

Now we can extend the algorithm to compute higher moments of the waiting time, for the purposes of variance minimization or optimization of nonlinear performance functions. When the objective is to compute the expected *squared* waiting time, the meaning of the cost-to-go of a state changes to the residual cumulative expected *squared* wait from that state on. We

will denote this function by $W^{(2)}(s) = \sum_{h \in H(s)} W_h^{(2)}(s) = \sum_{h \in H(s)} E[w_h^2(s)]$, where we have broken it

down by the squared waiting times of *individual* passengers $h \in H(s)$. The cost-to-go of a particular passenger h who is still waiting at state s can easily be expressed by means of the binomial expansion:

$$\begin{aligned} W_h^{(2)}(s) &= E[w_h^2(s)] = E\{[\Delta w(s, s') + w_h(s')]^2\} = E[\Delta w^2(s, s') + 2\Delta w(s, s')w_h(s') + w_h^2(s')] \\ &= \sum_{s'} \Pr(s, s') [\Delta w^2(s, s') + 2\Delta w(s, s')W_h^{(1)}(s') + W_h^{(2)}(s')] \end{aligned}$$

In this equation, everything is known at the time of back-up: the immediate costs (segment durations) of all outgoing transitions are constants, $W_h^{(1)}(s')$ is the expected waiting time of passenger h from state s' on, and has been found in the basic version of the DP algorithm (performed separately for each waiting passenger), and finally, $W_h^{(2)}(s')$ is the residual *squared* wait of successor states s' , and should be known if Bellman back-ups are performed in the correct order, namely backwards from the terminal state to initial one. The above equation defines the form of the Bellman back-up for the computation of the second moment of waiting times. As can be seen, it uses partial estimates (costs-to-go) obtained during the computation of the first moment, so moments should be computed in increasing order, since they “bootstrap” each other. Third and higher moments are also easy to compute, following similar binomial expansions, and reusing partial estimates from lower moments.

Another possible extension of the algorithm allowed by the DP method is the computation of the longest possible waiting time, so that the scheduler could attempt to avoid it in case it exceeds a pre-specified threshold. This can easily be done by defining the cost-to-go $W_h^{\max}(s)$ of state s to be the longest possible wait of a passenger from the current state on, again computed individually for each passenger. Instead of finding the expectation of waiting time, it would suffice for the Bellman back-up to find its maximum:

$W_h^{\max}(s) = \max_{s'}[\Delta w(s, s') + W_h^{\max}(s')]$. Then, $W_h^{\max}(s_0)$ of the initial state s_0 would be equal to the longest wait possible for passenger h .

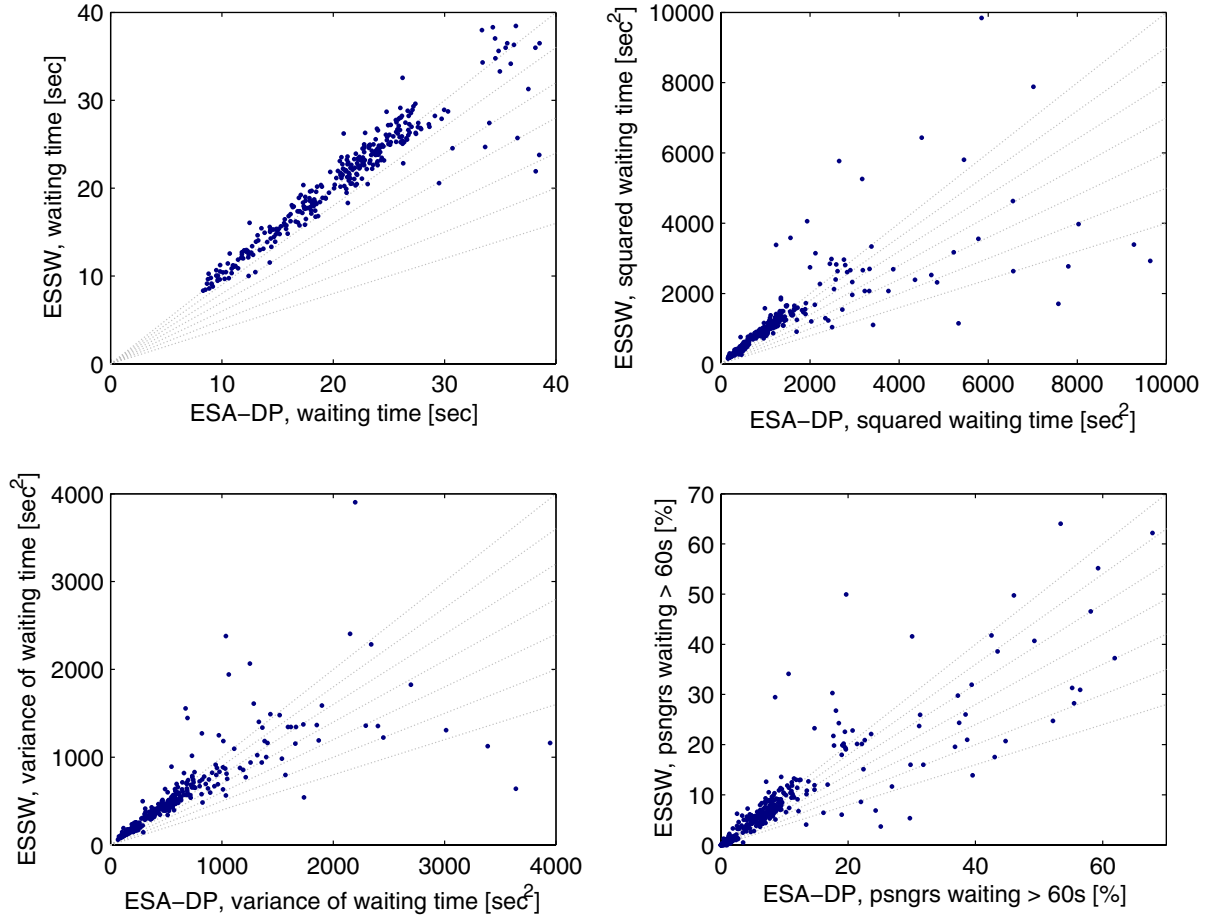


Figure 1. Comparison between two decision-theoretic schedulers minimizing different moments of waiting time. ESA-DP minimizes expected waiting time, while ESSW minimizes the expected square of waiting time, and indirectly its variance.

4. EXPERIMENTAL RESULTS

A risk-averse scheduler, ESSW, implementing the described method for computation and minimization of the second moment (square) of waiting time was compared experimentally with an identical decision-theoretic scheduler, ESA-DP (Nikovski and Brand 2003), which minimized simple expected AWT and was thus risk-neutral. The comparison was performed in ELEVATE 5.0, supplied by Peters Research, Ltd., and covered three buildings (10 floors/3 shafts, 15 floors/6 shafts, and 20 floors/8 shafts), in four different regimes (up-peak, down-peak, interfloor, and uniform), with three arrival rates and ten different arrival streams, resulting in a total of 360 simulated hours per scheduler. Each of the cars had a rated speed of 3 m/s, rated acceleration of 1 m/s² and jerk 1.13 m/s³. The results are shown in Figure 1, where each point is the average value of the respective parameter over all passengers in a one-hour simulation. Table 1 summarizes the averages over all 360 experimental runs. It can be seen that although the risk-averse scheduler (ESSW) is not necessarily much superior to the risk-neutral one (ESA-DP) in terms of pure AWT, it is still able to achieve lower SSWT and variance of waiting times, which in its turn leads to a lower fraction of excessive waits.

Table 1. Comparison between risk-neutral and risk-averse scheduling.

Algorithm	AWT	SSWT	Variance	% > 60s
ESA-DP, risk-neutral	24.78	1737.07	878.03	9.04
ESSW, risk-averse	23.62	1273.05	589.39	7.77

5. CONCLUSION

We have presented a decision-theoretic algorithm for computing estimates of any statistical moment of waiting time, and a method for using these moments to optimize performance criteria corresponding to different user attitudes towards long waits. The main simplifying assumption of ETA, that elevator cars move deterministically, has been eliminated, which opens the possibility for much more precise estimates of arbitrary functions of waiting times specified by elevator designers. The probabilistic approach to elevator control, combined with the significantly increased computational capabilities of modern embedded controllers, could possibly lead to a new generation of decision-theoretic group elevator schedulers.

BIBLIOGRAPHY

- Bao, G.; Cassandras, C. G.; Djafaris, T. E.; Gandhi, A. D.; and Looze, D. P. (1994). *Elevator dispatchers for down-peak traffic*. Technical report, University of Massachusetts, Department of Electrical and Computer Engineering, Amherst, Massachusetts.
- Barney, G. C. (2003). *Elevator Traffic Handbook*. London: Spon Press.
- Bertsekas, D. P. (2000). *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific.
- Crites, R. H., and Barto, A. G. (1998). Elevator group control using multiple reinforcement learning agents. *Machine Learning* Vol 33(235), pp. 1-33.
- Halpern, J. B. (1992). Variance analysis: a new way of evaluating elevator dispatching systems. *Elevator World*, September.
- Halpern, J. B. (1993). Variance analysis of hall call response time: trading mean hall call response time for increased variance performance. In G.C.Barney., ed., *Proceedings of ELEVCON'93*, Vol. 5, pp. 98-106. Bramhall, England: IAEE.
- Nikovski, D., and Brand, M. (2003). Decision-theoretic group elevator scheduling. *Proceedings of the 13th International Conference on Automated Planning and Scheduling ICAPS'03*, Menlo Park: AAAI Press, pp. 133-142.
- Press, W. H.; Teukolsky, S. A.; Vetterling, W. T.; and Flannery, B. P. (1992). *Numerical Recipes in C*, 2nd. Edition, Cambridge University Press.
- Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, NY: JohnWiley & Sons, Inc.

BIOGRAPHICAL INFORMATION

Matthew Brand is a Senior Research Scientist at Mitsubishi Electric Research Laboratories (MERL) in Cambridge, USA. His main research interest is in data modeling. Daniel Nikovski is a Research Scientist at MERL, and his research is focused on probabilistic methods for planning and learning.