

SIMULATING MARKET EFFECTS ON BOUNDEDLY RATIONAL AGENTS IN CONTROL OF THE DYNAMIC DISPATCHING OF ACTORS IN NETWORK-BASED OPERATIONS

James D. Brooks
David Mendonça

Industrial and Systems Engineering
Rensselaer Polytechnic Institute
Troy, NY 12180, USA

ABSTRACT

This work investigates the effect of market structure on the performance of actors who dispatch resources in network-based organizations. As digital tracking of these actors increases in scope and specificity, it is becoming feasible to extract computational models of dispatcher decision making which can be used to simulate market influences on these decisions. This work presents (1) a method for extracting computational decision models from transaction data and (2) a sensitivity analysis of the effect of market design (i.e., payment structure) and predictability of dispatcher decisions on system performance. The approach is illustrated for the case of a recent debris removal mission following extensive tornadoes in Alabama in 2011 which involved daily dispatching of an average of 28 hauling vehicles to an average of 6 work locations in the region of interest. Results of model validation against these historical tracking data are presented, along with implications for future work.

1 INTRODUCTION

There are many systems in which resources need to be allocated to a network including mining, logging, and debris removal following natural disasters. Recently, detailed transactional data from these systems is becoming available. However, these data are likely to be incomplete from the perspective of a system modeler. Thus, novel techniques are required in order to generate and validate a reasonable system model.

The resources flowing in these systems are commonly under the control of a much smaller number of dispatcher agents. These agents update the allocation of the resources to paths in the network in response to system performance. Additionally, the dispatcher agents are compensated according to some market context within which the performance of their entities is rewarded.

While much work has explored the performance of equilibrium flows under control of rational agents in congestion systems (see, e.g., Roughgarden 2003), there is no known work which considers the performance of such systems under the control of boundedly rational agents. Further, validation of such complex systems with limited data is uncommon. The work presents both a validation method and a simulation study in which the system performance under different market contexts and varying degrees of dispatcher agent rationality (i.e., decision determinism) are compared.

Background and related work are presented in Section 2. The model implementation details are discussed in Section 3. The case study data and model parameter estimation and validation are described in Section 4. Finally the simulation study and results are shown in Section 5. The paper ends with conclusions and comments about future work.

2 BACKGROUND & RELATED WORK

Pricing strategies may be used to induce behavior in congestion networks assuming rational agents. An agent decision model is presumed, and prices (or price ranges) are determined which induce behavior to be optimum in some sense (see e.g., Naor 1969, Stidham 1985, and Hassin and Haviv 2003). However, none of this work considers whether these decision models are representative of systems in which human actors are involved or whether the “optimal” pricing schemes lead to improved performance when the decisions made do not completely conform to the presumed deterministically rational model.

Modeling frameworks for dynamic decision making include, in rough order of decreasing cognitive plausibility, various cognitive architectures (e.g., Lee, Ravinder, and Johnston 2005 and Brooks, Wilson, and Sun 2012), beliefs desires intentions (e.g., Lee, Son, and Jin 2010), simple production systems (e.g., Hattori, Nakajima, and Ishida 2011), regression equations (e.g., Crowder, Robinson, Hughes, and Sim 2012), and Kalman filters (e.g., Pattipati, Kleinman, and Ephrath 1983). These models are typically generated through methods such as task analysis, interviews, questionnaires, or detailed behavior tracking.

Validation has been noted as a critical but often overlooked enterprise in modeling of dynamic systems. Several general validation methods have been proposed depending on the available system data (Kleijnen 1995). Of interest for the present work is *trace-driven* validation in which historical input data is used to generate comparable simulated output data. For queueing networks, comparison of mean throughput measures using estimated parameters from available historical data is a common approach (see e.g., Chen, Harrison, Mandelbaum, Ackere, and Wein 1988).

The network of interest has many parallel cycles in which several types of material are delivered from many remote sites to a fewer number of processing sites by some fixed number of hauling vehicles. Each processing site is assumed to handle some subset of the material types. The number of hauling vehicles working one any team which delivers material from each remote site is under the control of the dispatcher as shown in Figure 1. Team and remote site will be used interchangeably throughout. The task of the dispatcher is to dynamically allocate hauling vehicles (e.g., trucks) to these teams with the goal of maximizing productivity.




Figure 1: Example parallel-cycle queueing network under control by a single dispatcher.

3 MODEL

The desired model captures two elements: the dynamics of the work structure (i.e., the queueing behavior) and the effects of the market context and performance feedback on the dispatch decision maker. These modules will be discussed in turn in the following sections.

3.1 Queueing Network

The network of interest is one in which the nodes shown in Figure 1 are single-server queues. A standard discrete-event approach is taken for modeling queueing behavior of arriving customers at each of these nodes (i.e., remote sites and processing centers) without reneging, jockeying, or other non-standard customer behavior. The customers cycling in the queueing network represent hauling vehicles which are assumed

to be non-autonomous in the sense that they obey the dispatcher's orders. While the remote sites are determined by the dispatcher, the processing center is selected for each load according to the material type. It is assumed that the material type is randomly selected for each load according to some distribution. Each queue is further assumed to have an exponential service distribution with constant mean throughout. Finally, travel times between nodes in the network are assumed to be constant as well.

3.2 Dispatcher

The dispatcher agent model captures the dynamics of observed allocation decisions. The dispatchers are assumed to be economic, rational agents in the sense that their only consideration is to maximize their profit for a given market context. For example, if it takes \bar{W}_{ij} hours on average to make a single delivery from remote site i to processing center j (one complete cycle) for which the dispatcher receives a payment p_{ij} , the best decision at a given point in time is to move hauling vehicles from cycles with a lower average benefit-cost ratio, or *earning rate*, p_{ij}/\bar{W}_{ij} , to those with higher ratio.

It should be noted that there are likely many other considerations a dispatcher considers when allocating trucks (e.g., equipment failures, safety concerns, other ground conditions). However, the available data do not contain any information regarding these other factors. Further, because the dispatchers, and information sources (i.e., team members), are cognitively constrained, variability in selecting the best decision is to be expected. This variation is captured in the model through the use of a Boltzmann distribution (Luce's choice axiom) in which the probability of selecting any one option is proportional to the value of the choice (Luce 1977). The model assumes that for each decision, one hauling vehicle is transferred from one team to another. Thus, if there are n teams, there are $n(n-1)$ possible decision options. The value of these options is quantified by the log ratio of the earning rates, denoted a_k , for decision $k \in [1, n(n-1)]$. The probability of making decision k is then given by

$$p(a_k) = \frac{e^{a_k/T}}{\sum_i e^{a_k/T}} \quad (1)$$

where the amount of stochastic variation allowed is controlled by the parameter T . A low parameter value gives highly rational (i.e., deterministic, optimal) behavior, while a high parameter value gives completely stochastic selection (i.e., all choices equally likely).

The mean cycle times, \bar{W}_{ij} , are generally unknown to the dispatcher beforehand. Rather, only particular instances (i.e., trips), W_{ij} , are known. As a result, the dispatcher is required to make continual estimates of cycle times, \hat{W}_{ij} , as trips are made over time. The model assumes that these estimators are exponentially-weighted moving average (EWMA) filters of all past cycle times for that day which is given by

$$\hat{W}_{ij}^+ = \alpha W_{ij} + (1 - \alpha) \hat{W}_{ij} \quad (2)$$

where α is the forgetting factor parameter which indicates the amount of averaging in the filter.

Further, it is assumed that decisions are made at random points in time, herein called decision points, and that each instance results in one and only one truck being transferred as discussed previously. The dispatcher model execution then follows the flow shown in Figure 2. The decision points are chosen randomly according to some inter-decision distribution. At each point, the estimates of cycle times and the known payments are then considered to make the decision.

Thus, the question remains: can a pricing scheme which presumes a rational agent still provide performance benefits even when the agents act in a boundedly rational way with respect to the presumed model? The next section briefly describes the study data taken in the context of debris removal following natural disasters, and the method in which parameter estimates and model validation were performed. The simulation study designed to explore this question is then described in Section 5.




Figure 2: Flow chart of dispatcher decision model execution.

4 CASE STUDY

This section provides background for the case study and the parameter estimation and validation method.

4.1 Background

The passage of an extensive tornado storm in April, 2011 resulted in widespread damage throughout the northern half of the state of Alabama and neighboring regions. The storm is estimated to have generated nearly 10.5 million cubic yards of debris in Alabama alone (Paepcke 2012). The U.S. Army Corps of Engineers (USACE) coordinated the mission, for which there was a single prime contractor and two tiers of subcontractors. The first debris load for this mission was delivered on 8 May 2011 and the last on 20 October 2011.

In order to clear this debris, many debris removal teams are at work throughout the affected area (i.e., at various remote sites). Hauling vehicles are assigned to these teams by dispatchers to transport material (either vegetative, construction and demolition (C&D), or hazardous material) to the temporary debris storage and reduction (TDSR) sites (i.e., processing sites). A schematic of the work flow for one representative loop of activity is shown in Figure 3.

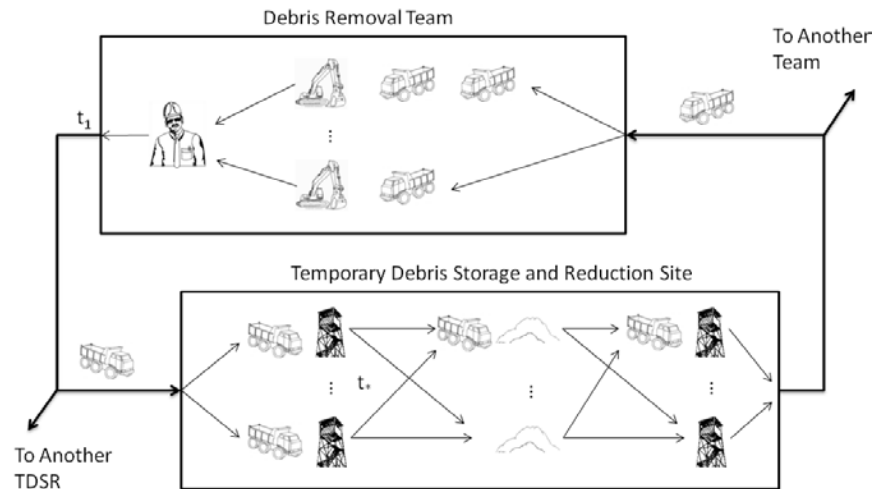


Figure 3: Overview of the workflow for debris removal showing time stamps available in the transaction data as t_1 and t_* in the diagram.

Data Sources All load ticket data were obtained from the USACE in machine-readable format, then cleaned and verified via automated and manual methods. For example, automated queries were used to identify potential anomalies and errors in the data (e.g., AM/PM time errors, date errors, out of range distances, invalid time sequences between tickets). This cleaning resulted in nearly 130,000 load tickets

over five months, during which the organization consists of one prime contractor and 33 sub-contractors, operating within 33 municipalities, with a total of 49 TDSRs.

These load tickets provide pickup and drop-off identifiers and one time stamp at each location for every trip made in the system by all entities – leaving time at remote site and arriving time at processing center. Thus, the system structure and allocation can be inferred from these data. An average of approximately 66 debris removal teams (i.e., remote sites) were active in the field on each day, and a total of 2271 different trucks took part in the mission. Team sizes ranged from two to thirty two trucks (mean of 4.0), with 81.6% of teams consisting of five or fewer trucks.

An isolated region of the data is used for model validation and parameter estimates. The selected restricted set contains 7,440 loads hauled by 101 trucks to 2 TDSR sites (i.e., processing sites – one for vegetative material and the other for C&D material) over 51 days. These trucks were organized into 4 debris removal teams on average and 99.5% of the loads were hauled by a single subcontractor, thus avoiding the possibility of any strategic behavior due to competition with other subcontractors in the region. To further avoid the confound of system configuration, only days for which 6 remote sites were active are used. A summary of this data is shown in Table 1.

Table 1: Summary of validation field data. Mean values are shown with ranges in parentheses.

<i>Day</i>	<i># Trucks</i>	<i>Team Size</i>	<i>Trips Made</i>
May 16	12	2.00 (3)	9.83 (25)
May 20	24	4.00 (7)	27.33 (58)
May 22	22	3.67 (7)	30.50 (68)
May 26	22	3.67 (6)	16.67 (36)
May 31	23	3.83 (9)	31.17 (93)
June 5	21	3.50 (7)	16.67 (43)
June 12	13	2.17 (1)	14.50 (17)
June 14	24	4.00 (10)	26.17 (63)
June 15	16	2.67 (7)	17.83 (53)
June 25	22	3.67 (3)	28.17 (46)

4.2 Model Implementation & Parameter Estimation

The model was implemented in ARENA according to the earlier discussion. The system configuration and parameters for each day is determined by the procedures discussed here. As described previously, some field data on the operation of this representative system is available. While the data are incomplete, much information about the system parameters can be extracted. The methods employed to do so are described in this section.

Processing Center Service Rates Several assumption are made to estimate these parameters from the limited available data. The first is that waiting time at the processing center prior to beginning service is of the same magnitude as the wait time at the remote sites (i.e., equal queueing wait times at inspection tower and remote site on average). Additionally, all entities and servers are assumed to be working at all times (e.g., no lunch breaks). The unmeasured return travel time is furthermore assumed to be the same at the measured delivery travel time. Travel times between each remote-processing pair are estimated from the data using the mean of the time stamp differences. For entities which complete another trip from the same remote site, an estimate of the waiting time at the processing center is made using the difference between subsequent departures at the remote site. Twice the measured travel time is subtracted from the difference between remote site departures to give an estimate of that trip’s wait time at the processing site. Since all entities share processing sites, these wait time measures are then averaged for each processing

site each day to provide an estimate of the mean waiting time, W_j . The mean arrival rate to processing site j , λ_j , is simply the sum of arrivals divided by the duration of the working day (12 hours). Assuming the processing center is an independent M/M/1 queue (Gross, Shortle, Thompson, and Harris 2008), the estimated average service rate for a given day is then

$$\mu_j = \frac{1}{W_j} + \lambda_j.$$

These measures and the resulting estimates are shown in Table 2. The model service rate is then chosen to be fixed as the mean of these daily rate estimates for each central processing site.

Table 2: Arrival rates, wait time estimates, and calculated service rate estimates for both processing centers each day. The parameter values used in the model are in shown in boldface.

Day	Mean Arrival Rate		Wait Time Estimate		Service Rate Estimate	
	Veg	C&D	Veg	C&D	Veg	C&D
May 16	–	4.92	–	0.75	–	6.26
May 20	9.42	4.25	0.42	0.52	11.78	6.18
May 22	13.25	4.17	0.49	0.40	15.30	6.69
May 26	3.25	5.08	0.39	0.35	5.80	7.94
May 31	12.17	6.08	0.58	0.45	13.89	8.32
June 05	3.58	5.08	0.32	0.44	6.70	7.36
June 12	5.50	1.75	0.44	0.66	7.78	3.26
June 14	7.00	6.08	0.60	1.01	8.68	7.07
June 15	10.42	3.92	0.54	0.55	12.26	5.73
June 25	8.17	5.92	0.52	0.66	10.10	7.44
Mean (Std.Dev.):	8.08 (3.57)	4.73 (1.31)	0.48 (0.09)	0.58 (0.20)	10.26 (3.28)	6.62 (1.43)

Material Type The load trip distribution between processing centers in the field data is assumed to be representative of the material makeup and outside of the control of the entities (i.e., material type is stochastic). The routing probabilities from each remote site to the processing centers is determined by the proportion of trips made in the field data.

Remote Site Processing Rates These rates cannot be estimated from the data. As a result, they are assumed to be equal and this common rate is determined in the validation method as described in the next section.

Information Integration EWMA Filter The forgetting factor for the filter in Equation (2) is chosen from the set (0,1) by maximizing the number of observed decisions which are *logical* (i.e., $a_k > 0$) by numerically evaluating the observed decisions with filters taking various parameters on this interval. It is found that factor which best matches the field data (i.e., the one which maximizes the number of logical decisions) is $\alpha = 0.94$ (see Figure 4a). This suggests that the dispatchers are generally very reactive to the current state (i.e., the most recent trip made) rather than averaging over a substantial number of past trips.

Level of Rationality The value for the parameter in Equation (1) is chosen to be the one which maximizes the total likelihood of the observed decisions while using the forgetting factor chosen previously. The parameter value which maximizes the sum likelihood of the observed decisions is $T = 0.73$ (likelihood = 6.77; see Figure 4b). This represents a mean likelihood for any observed decision of only 5.4%. This suggests that the current dispatching decisions are highly stochastic (common parameter values in cognitive modelling literature are 0.1–0.25; see e.g., Gonzalez 2003 and Sun, Merrill, and Peterson 2001) and not often economically optimal.

Several other potential rational models were also evaluated which used haul distance, team size, and travel time estimates, but none provided a higher maximum total likelihood. Also note that a purely random




Figure 4: Parameter estimation plots for (a) number of logical decisions by EWMA forgetting factor parameter and (b) total likelihood by Boltzmann parameter.

model would give a total likelihood of 4.2. Thus, the proposed model gives 62% greater likelihood than the null (purely random) model.

Inter-Decision Times The field data suggests that the intervals between decisions are exponentially distributed with a mean of 0.44 hours.

4.3 Validation

It is critical that the network and decision model represent the system of interest so that any subsequent insight into market influences with boundedly rational agents is valid for the case study context. The basic method used is similar to trace-based method (Kleijnen 1995) in which some available data is used to setup the simulation. In this case, the available data include system structure (i.e., number of remote sites, their location, mean travel time estimate, and number of entities in each cycle), and performance (i.e., number of trips made) for each day. A range of values for the final remaining model parameter, the common remote site processing rate, is considered and that with the smallest error from the observed data traces is selected.

It should be highlighted that because the validation data contains several resource allocation decisions the number of hauling vehicles in each team was not constant. In order to validate the system model separately from the decision model, the number of vehicles in each team was extracted over time and the mean value is used for the validation (cf. effective cars in Green and Kolesar 1989).

The validation performance measure used is the Chi-square measure where the expected outcome are the trips made in the field data and observed are those from the simulation using the same configuration:

$$V(\mu) = \sum_{i \in \{\text{remote}\}} \sum_{j \in \{\text{processing}\}} \sum_{k \in \{\text{days}\}} \frac{(O_{ijk}(\mu) - E_{ijk})^2}{E_{ijk}}$$

The common remote site service rate is then chosen as the one which minimizes this validation measure as given by

$$\mu^* = \text{argmin } V(\mu).$$

Each of the 10 days of validation data along with several possible remote site service rates was setup as a scenario in ARENA's Process Analyzer for this procedure. Figure 5 shows the validation performance measure, V , as a function of the common processing rate. The field data is best reproduced using a common remote site service rate of $\mu^* = 6$. Note that the error curve is relatively flat as the processing rate is




Figure 5: Validation performance measure as a function of common remote site service rate.

further increased. This indicates that the amount of queuing at the remote sites is likely minimal in the system.

5 SIMULATION STUDY

The goal of the simulation study is two-fold: first to explore the effect of dispatcher rationality on system performance generally and secondly to assess the performance of different market contexts. To do so, the following experimental conditions are considered. The study compares system performance (i.e., total throughput) under two levels of rationality and three pricing strategies — uniform, proportional to distance, and proportional to one-way travel time (as this is independent of system congestion).

The two levels of dispatcher rationality — as determined by the parameter in the Boltzmann distribution described earlier — are chosen such that the first is that found to maximize the total likelihood of the decisions observed in the field data ($T = 0.73$) while the other is substantially more rational ($T = 0.2$).

The *uniform* pricing strategy awards one unit for each trip. For the *proportional to distance* pricing strategy, the prices are the distance divided by the average distance so that the remote site which is an average distance away from the processing centers gets compensated one unit per trip. Similarly for one-way travel times for the *proportional to travel time* pricing strategy.

Notice that for each team there are two travel times and distances (one for each processing center). The weighted average is taken according to the material debris proportions to calculate a single expected distance and time to be used. The team properties and resulting prices are shown in Table 3. It should be noted that no claims are made regarding optimality of any of these schemes. Rather, they are representative of the current pricing policy in a typical debris removal system (distance-based) and a comparable alternative (time-based). This study merely evaluates them as an interesting first step toward the evaluation of the degree of influence of pricing policies (i.e., market context) with boundedly rational decision makers.

The system configuration used (i.e., travel times, debris percentages) is that from May 22nd. This day was chosen because it is the only one which contained loads delivered to both processing centers for all six teams. Finally, the simulations are all initialized with a uniform assignment (i.e., 4 hauling vehicles assigned to each remote site). Each simulation is executed for one working week (5 days — 60 hours) with 100 replications for each experimental condition.

Table 3: Team properties and resulting prices used for each team in the simulation study.

Team	Material Proportion		Properties		Prices by Strategy		
	Veg	C&D	Avg. Dist	Avg. Travel Time	Uniform	Time	Distance
1	0.33	0.67	12.17	0.47	1	0.85	1.7
2	0.60	0.40	5.32	0.38	1	0.69	0.74
3	0.92	0.08	2.09	0.50	1	0.91	0.29
4	0.20	0.80	13.80	0.65	1	1.18	1.93
5	0.92	0.08	2.01	0.47	1	0.85	0.28
6	0.43	0.57	7.61	0.85	1	1.53	1.06




Figure 6: Total system relative throughput and revenue for all six conditions.

5.1 Results

The mean system throughput results are shown in Figure 6a. Here all values are scaled to the uniform pricing strategy with the observed decision parameter (high stochasticity). Two observations are immediately evident. First is that the uniform pricing strategy produces the best performance under both rationality conditions. And secondly, that increasing rationality does not lead to better performance under all pricing strategies. While this may initially appear counter-intuitive, notice that increased rationality is only expected to improve system performance when the pricing strategy is optimal. For non-optimal pricing strategies, one the other hand, increased rationality is expected to decrease performance.

Additionally, at the system level it is clear that the mean total revenue earned by the dispatcher increases with both non-uniform pricing strategies (see Figure 6b). While this is good for the dispatcher, it is unfortunate for the system stakeholders in that this added cost has not resulted in increased performance. Further, it is interesting to note that the variation in cost is less for the more rational condition.

Turning now to individual team performance, the observed scaled mean throughputs for each remote site and processing center are shown in Figures 7 and 8, respectively. It is easy to see that the pricing strategies influence individual team performance differently in line with their relative travel time and distance properties. The effect of the level of rationality is often not significant at the team level (only 9 of 36 confidence intervals are non-overlapping).

Finally, these results show that despite the stochastic nature of the observed decisions, the allocation behavior and resulting performance is shown to be influenced by the pricing strategy in the expected way.

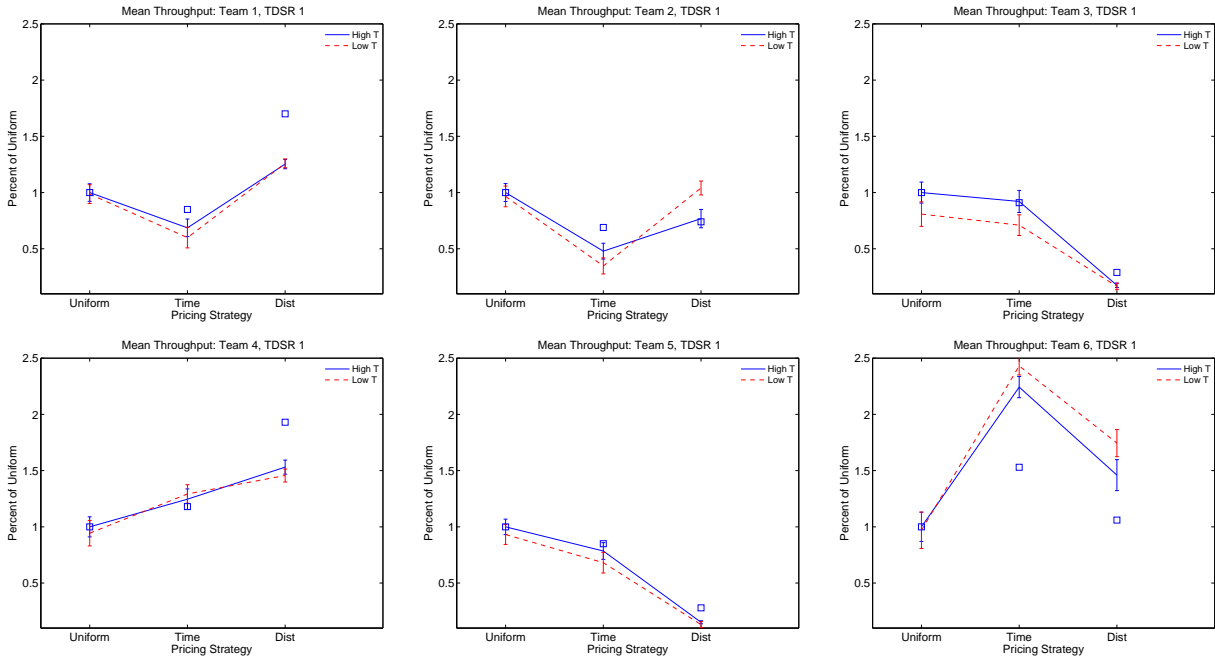


Figure 7: Average relative throughput with 95% confidence intervals by team for first TDSR. Team prices are also shown as boxes in the plots.

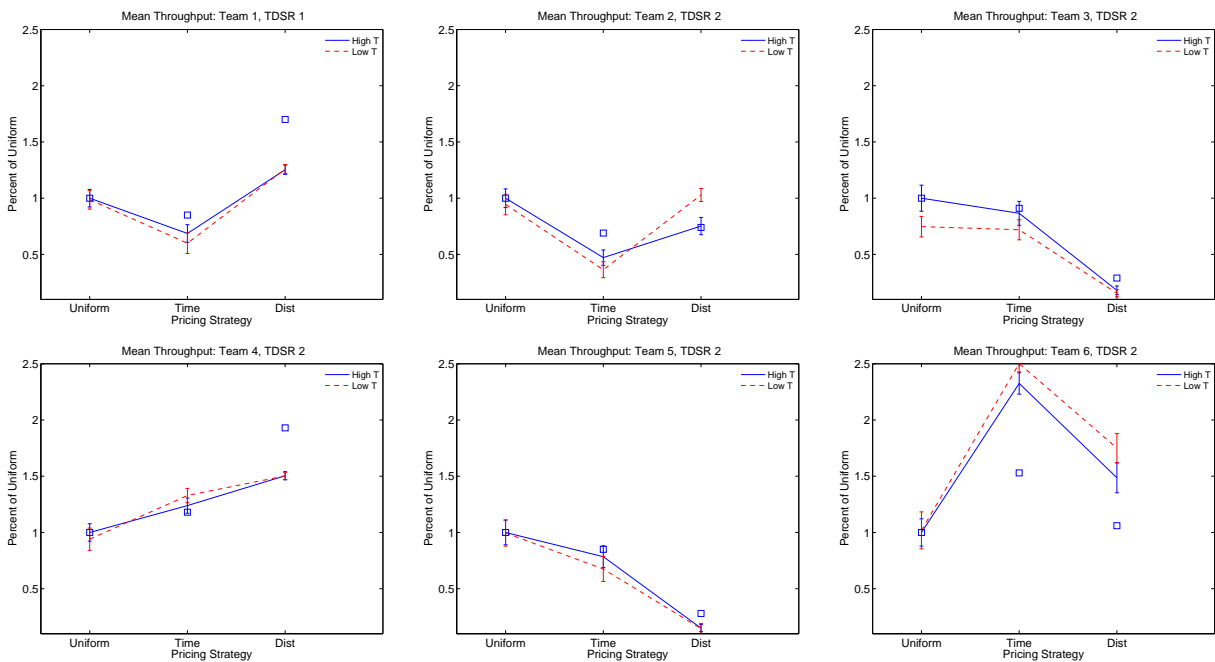


Figure 8: Average relative throughput with 95% confidence intervals by team for second TDSR. Team prices are also shown as boxes in the plots.

6 CONCLUSIONS

This work has presented a queueing network model for simulating the effects of market context on resource allocation behavior in networks. A novel validation technique was presented which allows approximate parameter estimation of queueing network processing rates and decision model parameters which maximize the total likelihood of observed decisions with limited transaction data. Further, specific validation for a recent debris removal mission was presented. A simulation study was also presented in which dispatcher rationality and market context were varied.

The results show that market context can still effectively influence allocation behavior of dispatchers even when the observed decisions are highly stochastic with respect to the idealized rational model. Further, the results underscore the importance of properly designing the market context for these systems. In particular, the particular configuration under study in this paper exhibits decreased performance and increased revenue (i.e., cost to system owner) for both the time- and distance-based policies explored.

It should also be noted that there appears to be some evidence of Parkinson's Law (i.e., work expands to the available time) in that the processing center service rates generally increase with arrival rates. A similar effect was recently reported in other human-based queueing systems (Marin, Drury, Batta, and Lin 2007). However, it can not be determined whether this effect was due to reconfiguration of the processing centers or adaptive behavior within the same configuration.

Future work includes an investigation of additional factors of interest to dispatchers in making these allocation decisions. Work is also underway to determine the optimal pricing strategy for these kinds of network systems. Additionally, other performance effects due to switching team membership are expected which can be included in the model and market context.

ACKNOWLEDGEMENTS

Thanks is due to several individuals at the US Army Corps of Engineers for assistance in procuring the data and helpful discussions: Dr. Stephen Diaz, Matt Tate, Jim Garner, Cecil Jernigan, and Marcella Denton. Many thanks also to Olivia von Nieda for implementing the simulation model and generating results. This material is based upon work supported by the National Science Foundation under Grant No. 1313589.

REFERENCES

- Brooks, J. D., N. Wilson, and R. Sun. 2012. "The effects of performance motivation: A computational exploration of a dynamic decision making task". In *Proceedings of the First International Conference on Brain-Mind*, 7–14. East Lansing, Michigan: BMI Press.
- Chen, H., J. M. Harrison, A. Mandelbaum, A. V. Ackere, and L. M. Wein. 1988. "Empirical evaluation of a queueing network model for semiconductor wafer fabrication". *Operations Research* 36 (2): 202–215.
- Crowder, R. M., M. a. Robinson, H. P. N. Hughes, and Y.-W. Sim. 2012. "The development of an agent-based modeling framework for simulating engineering team work". *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans* 42 (6): 1425–1439.
- Gonzalez, C. 2003. "Instance-based learning in dynamic decision making". *Cognitive Science* 27 (4): 591–635.
- Green, L., and P. Kolesar. 1989. "Testing the validity of a queueing model of police patrol". *Management Science* 35 (2): 127–148.
- Gross, D., J. F. Shortle, J. M. Thompson, and C. M. Harris. 2008. *Fundamentals of Queueing Theory*. John Wiley & Sons.
- Hassin, R., and M. Haviv. 2003. *To queue or not to queue: Equilibrium behavior in queueing systems*. Norwell, Massachusetts: Kluwer.
- Hattori, H., Y. Nakajima, and T. Ishida. 2011. "Learning from humans: Agent modeling with individual human behaviors". *IEEE Transactions on Systems, Man and Cybernetics, Part A* 41 (1): 1–9.

- Kleijnen, J. C. P. 1995. "Verification and validation of simulation models". *European Journal of Operational Research* 82:145–162.
- Lee, S., Y.-J. Son, and J. Jin. 2010. "An integrated human decision making model for evacuation scenarios under a BDI framework". *ACM Transactions on Modeling and Computer Simulation* 20 (4): 1–24.
- Lee, S. M., U. Ravinder, and J. C. Johnston. 2005. "Developing an agent model of human performance in air traffic control operations using apex cognitive architecture". In *Proceedings of the 2005 Winter Simulation Conference*, 979–987.
- Luce, R. 1977. "The choice axiom after twenty years". *Journal of Mathematical Psychology* 15 (3): 215–233.
- Marin, C. V., C. G. Drury, R. Batta, and L. Lin. 2007. "Human factors contributes to queueing theory: Parkinson's law and security screening". In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 602–606.
- Naor, P. 1969. "The regulation of queue size by levying tolls". *Econometrica* 37 (1): 15–24.
- Paepcke, J. 2012. "Statewide April 27th tornado debris rankings, by county and city". Available via <http://www.alabamas13.com/story/21406257/statewide-april-27th-tornado-debris-rankings-by-county-and-city> [accessed 18 February, 2013].
- Pattipati, K. R., D. L. Kleinman, and A. R. Ephrath. 1983. "A dynamic decision model of human task selection performance". *IEEE Transactions on Systems, Man and Cybernetics* 13 (3): 145–166.
- Roughgarden, T. 2003. "The price of anarchy is independent of the network topology". *Journal of Computer and System Sciences* 67 (2): 341–364.
- Stidham, S. 1985. "Optimal control of admission to a queueing system". *IEEE Transactions on Automatic Control* 30 (8): 705–713.
- Sun, R., E. Merrill, and T. Peterson. 2001. "From implicit skills to explicit knowledge : A bottom-up model of skill learning". *Cognitive Science* 25:203–244.

AUTHOR BIOGRAPHIES

JAMES D. BROOKS is a Ph.D. candidate in the Department of Industrial and Systems Engineering at Rensselaer Polytechnic Institute. His research interests generally lie in methods for distributed control of complex human-machine systems. He is a member of INFORMS and IEEE. His email address is brookj7@rpi.edu and his web page is <http://homepages.rpi.edu/~brookj7>.

DAVID MENDONÇA is an Associate Professor in the Industrial and Systems Engineering Department at Rensselaer Polytechnic Institute. He has been a visiting scholar at Delft University of Technology (The Netherlands) and the University of Lisbon (Portugal). His research examines the cognitive processes underlying individual and group decision making in high stakes, time-pressured conditions, particularly in the domains of post-disaster emergency response and infrastructure restoration. His email address is mendod@rpi.edu and his web page is <http://homepages.rpi.edu/~mendod>.