

DECISION SCIENCES INSTITUTE

A Simulation Approach for Evaluating Urban Snow and Ice Removal Planning and Operations

(Full Paper Submission)

Vitali Volovoi
LMI, Independent Contractor
vvolovoi@lmi.org

David K. Peterson
LMI, Senior Consultant
dapeterson@lmi.org

ABSTRACT

This paper presents a graphical discrete-event simulation (DES) model for the planning and operation of snow and ice removal in urban areas. The modeling is developed based on the scenarios previously described in Alprin (1975) using a new version of stochastic Petri nets (SPN) called Abridged Petri Nets (APN). The importance of flexible vehicle assignment to different snow removal routes is demonstrated. The model includes constraints with respect to limited vehicle and infrastructure resources as well as allowing for the possibility of mid-mission vehicle failures.

KEYWORDS: Discrete Event Simulation, Stochastic Petri Nets, Urban Planning, Decision Theory, Vehicle Assignment Problem

INTRODUCTION

The winter of 2013-2014 was an especially rough season for many parts of the country with record-setting snow totals in many cities and regions. In fact, in the U.S., this year's winter weather has been so severe, that it's becoming commonplace to name the experience. For example, "Snowpocalypse" in Atlanta and "Chiberia" in Chicago.

Unfortunately, this severe winter weather also poses significant public safety problems. "Motor vehicle accidents involving wintry conditions and other hazardous weather claim the lives of more than 4,000 people in the United States and injure several hundred thousand each year" (NCAR, 2013). The severe weather also poses significant fiscal challenges to state and local governments, with many agencies exceeding their snow removal budgets and forcing them to tap into future infrastructure maintenance funding (Vock, 2014). For example, as of March 2014, Virginia's Department of Transportation (VDOT) had spent about \$156 million (budgeted \$63 million) for snow removal and Maryland's State Highway Administration had spent \$123 million (budgeted \$46 million) (WTOP, 2014).

These safety and fiscal challenges highlight the importance of governmental transportation managers employing the class of analytical capabilities generally categorized as "business analytics." These analytical techniques include "the use of data, information technology, statistical analysis, quantitative methods, and mathematical or computer-based models to help managers gain improved insight about their business operations and make better, fact-based decisions" (Evans, 2012).

Given that governmental transportation agencies need to quantify, project, and defend snow removal budgets and resources, more analytically-oriented operational planning can help

transportation managers make better resourcing decisions while also helping these same agencies better address the myriad resourcing challenges they face in deploying and sustaining snow removal equipment within the context of the agency's charters.

BACKGROUND AND LITERATURE REVIEW

In this section we first review some key concepts regarding urban snow and ice removal with an emphasis upon emerging technological trends and their implications for urban planners and their business analytics approaches. We then summarize route optimization research with respect to snow removal operations and highlight an apparent gap in analytical capabilities currently available to governmental transportation planners.

Urban Snow and Ice Removal Operations

Across the U.S., every winter an army of "snow warriors" works around-the-clock to keep the nation's roadways open. As noted by Virginia's Transportation Secretary "this winter has put a strain on our maintenance budget.... However, we have sufficient resources to ensure our roads are well kept and we will continue to be good stewards of public funds.... we're committed to delivering a safe and reliable transportation system" (Rollison, 2014). In recent years, the best snow removal strategy that many transportation agencies recommended wasn't "for snowplows to make it possible for people to drive to work and school – it's for people to refrain from driving to work and school so that snowplows can get on with it" (Pearlstein, 2014).

In general, when weather forecasts indicate the possibility of snow or icing conditions, transportation departments begin preparing and prepositioning their equipment (DSNY, 2014). Salt spreaders are deployed once the snowfall begins; however, when the accumulation exceeds roughly three inches or the temperature drops too far (exceeding salt's effective temperature range), snowplowing operations commence. Depending on the snow totals, these snowplow operations can become multi-day efforts. For example, in the northern Virginia region, this season one eight-inch snowfall required 4,000 trucks and 36 hours to complete the snowplowing plowing operations (Pearlstein, 2014).

Technological Trends in Urban Snow and Ice Removal

Like so many facets of modern life, sophisticated information-based applications have brought greater efficiency and effectiveness to snowplowing operations.

For example, in the past, transportation officials have often lacked critical information about roadway conditions. They relied on ground-based observing stations that were spaced miles apart, requiring them to interpolate the intermediate weather conditions (NCAR, 2013). In addition, intervening micro-climate effects on roadways created highly variable and unpredictable conditions which contributed to ineffective and inefficient snowplowing operations. "If officials dispatch snowplows unnecessarily, or treat roads ... when not needed, they risk wasting money and harming the environment. If they do not treat the roads, however, drivers may face treacherous conditions" (NCAR, 2013). To better manage this risk, the U.S. Department of Transportation launched a test program "deploying hundreds of plows [across three states] with custom-designed sensors that continually measure road and weather conditions.... It gives road crews ... [a] detailed, mile-by-mile view of road conditions" (NCAR, 2013). Global Positioning System (GPS) technology is used to locate and time-stamp these measurements which are transmitted via Internet, radio or cellular networks to NCAR centralized databases and then shared with the states.

Another rising information-based capability is the use of geographic information system (GIS) technology for snowplow tracking. A recent project in Columbus, OH is typical of these efforts. Columbus' project involved monitoring over 100 snow removal vehicles across approximately 2,000 miles of roadway (Scott, 2013). "This new GIS technology will internally monitor both real-time and historical performance of the city's snow and ice removal activities (Scott, 2013)." Similar GPS-based capabilities have been deployed in other U.S. cities, such as Chicago and New York City where PlowNYC "tracks the location of more than 1,700 plows and spreaders ... working to clear the streets and roadways of the 5 boroughs" (Lewis, 2013).

A final example of the impact that information-based technologies are having upon winter snow removal operations and driving safety is Maryland's Coordinated Highways Action Response Team (CHART). The CHARTWeb website provides a variety of atmospheric and surface level metrics, including: air temperature; wind speed and direction; precipitation type, rate and accumulation; and pavement temperature (CHART, 2014). This latter metric is of particular interest as it provides the "Temperature range of the pavement sensor roughly 3 mm ... below the surface of the sensor" and was used by Washington DC news outlets to warn motorists of possible winter roadway icing conditions.

These and many other new and upcoming information-based capabilities are providing transportation managers with insights into their snow removal operations that were previously unimaginable. Clearly the era of "big data" has arrived for transportation operations, and technology is providing ever-more robust data sources which can serve as key enablers for governmental transportation managers to routinely employ business analytics.

Route Optimization Research

Perrier, Langevin and Campbell (2007) surveyed optimization models and solution algorithms for spreading operations vehicle routing. They classified these vehicle routing solution methods into three broad classes: optimization methods, rule-based methods, and heuristics (which was further subdivided into constructive methods, composite methods and metaheuristics).

Based upon their survey, Perrier, Langevin and Campbell (2007) observed that "Vehicle routing problems in winter road maintenance are the most studied of any winter road maintenance problems. Because of the inherent difficulties of these problems, most solution methods that have been proposed are heuristics." They further noted that, while "Early attempts to apply simple heuristics... produced nice results from simulation studies, [they] ... were rarely implemented and used in spreading operations. Recent models are solved with more sophisticated local search techniques... [and] are showing much promise to assist planners in making routing decisions for spreading operations in practice."

Perrier, Langevin and Campbell (2007) conclude that "considerable work remains to be accomplished on the design of fast heuristic algorithms that produce good approximate solutions, and on the development of more comprehensive models that address the integration of depot location models with spreader routing decisions."

Route Evaluation Needs

In their conclusions, Perrier, Langevin and Campbell (2007) suggest where research should be focused with regard to the spreader routing problem. However, they do not describe how to determine what makes for a "good approximate solution." Alprin (1975) used a simulation to test the performance under uncertainty of two alternative heuristic solution methods under "Various combinations of resources used in snow removal, such as the number and location of salt stockpiles or the number and type of salt spreader trucks...."

This thesis research was formally published by Cook and Alprin (1976) and their paper is frequently cited in the literature review of many subsequent papers researching the snowplow/spreader vehicle routing problem. Unfortunately, it appears that the simulation originally developed by Alprin (1975) was not regularly employed by these latter research efforts. However, if his simulation could be reprised and implemented in a modern, visually-oriented simulation language, it would be of great benefit for evaluating the relative performance of alternative snowplow route selection solution techniques.

THEORETICAL AND MODEL DEVELOPMENT

Given the clear need to arm urban planners with modern modeling and simulation tools for evaluating their snow and ice removal plans, we next lay out the analytical and visualization capabilities of a modern discrete event simulation. From this basis, we replicate Alprin's (1975) snow and ice removal simulation; specifically, the two cases employed in his analysis – assigned routes and shortest distance route assignments.

Simulation and Visualization

In broad terms, there are three distinct established frameworks for modeling dynamic interactions among multiple entities: discrete event simulation (DES), system dynamics (SD), and agent-based modeling (ABM). Each framework provides distinct capabilities. As discussed in Brailsford (2008) and the references provided therein, DES provides a more flexible means for quantitative analysis and the resulting models are quite detailed and associated with the significant set-up costs both in terms of time and money. In contrast, SD models are more useful for understanding the “big-picture”, with an emphasis on qualitative analysis as SD diagrams are deemed to be more easily understood by the clients and managers. SD models are usually easier to construct than DES, but they lack the quantitative flexibility of DES.

ABMs are relative newcomers to the simulation world and they provide even more flexibility than DES. The main reason for this flexibility is ABM's reliance on local (distributed) rules for the simulation sequence, and the possibility to include agents with continuously changing state-space. In contrast, DES (as the name implies) is focused on discrete events (i.e., discrete changes to the system's state) and it generally relies on a centralized list of events that has to be maintained in chronological order during the simulation. The flexibility of ABMs comes at a price, as the resulting models are usually even more detailed and less structured than DES, and therefore are more computationally intensive and more difficult to create. At the same time, the ability of implementing locally distributed rules for updating an agent's state can reduce the need for computational resources if the level of abstraction is chosen judiciously (Yu et al. 2011). As commercial tools evolve to satisfy their analytical customer's needs, the boundaries between these modeling frameworks are often blurred.

For historical reasons, most DES research models are based on concepts derived from queuing theory with activities, resources, and queues providing the fundamental building blocks of the model. Stochastic Petri nets (SPNs) provide an alternative means to create DES models. Therein the system is comprised of individual components that can change their states and can trigger or prevent the state changes of other components. As a result, while an activity can be interpreted as a time delay between the state transitions for a given component, the concepts of resources and queues are derived properties, rather than fundamental properties. SPNs and their parent framework, Petri nets, are popular tools in computer science as they provide a well-structured and visual means for analyzing complex interactions among components such as synchronization and concurrent operations. However, their commercial applications in the context of DES are very restricted to date, as practicing engineers find them too abstract and

difficult to understand. Different flavors of SPNs (developed mainly in academia) have come and gone for the past 30 years, and commercial tools are mainly used internally by the more analytically-oriented companies (such as Siemens or Total).

In this paper a new version of SPNs, Abridged Petri Nets (APNs), is utilized. APN is a new graphical framework for modeling the stochastic behavior of complex systems that consist of multiple interacting components (Volovoi, 2013). This framework can be considered as a derivative of SPNs (Marsan, 1990) that aims at retaining SPN's versatility in terms of modeling power, while streamlining the choice of the modeling building blocks. The visual clutter and often confusing choices that are often perceived as the major obstacle to the larger success of SPNs are reduced (Bowden 2000), resulting in simpler and more transparent models that can be built using only a graphical interface.

The following essential properties of APNs can be identified:

- An APN is defined as a network of places (denoted as hollow large circles) that are connected by directed arcs (transitions). Changes in the system's state are modeled by a transition firing: i.e., moving a token from the transition's input place to its output place. The combined position of tokens in the net at any given moment represents marking of the net and fully specifies the modeled system.
- Each transition has no more than a single input and single output place (a transition can also have no input place, providing a source of new tokens every time it fires, or it can have no output place, providing a sink for tokens; upon the firing of such a transition, a token is removed from the net).
- Each token can have a discrete label (color) that can change when the token moves. In addition, tokens have continuous labels (ages) that can change both when tokens move, and with the progression of time while a token stays in the same place (Volovoi 2004).
- A transition is enabled or disabled based on the combined marking of the input places of the associated triggers (inhibitors and enablers). Inhibitors are depicted as arcs originating at a place and terminating at a transition with a hollow circle. An inhibitor of multiplicity k disables a transition that it terminates at if the number of tokens in its input place is at least k . An enabler (depicted as an arc originating at a place and terminating at a transition with a filled circle) is the opposite to an inhibitor: a transition is disabled unless an enabler of multiplicity k has at least k tokens in its input place. Triggers can be color-specific (and therefore enable only tokens in a place of a certain color, or act only if there are a specified number of tokens of a given color in their respective input places, or both).
- Transitions have color- and age-dependent policies that specify the delay between the moment when the token is enabled and when it is fired.
- If a token-transition pair is enabled, a firing delay is specified based on the combination of token and transition properties. If the token stays enabled throughout the delay, after this delay expires the token is fired. If there are multiple enabled tokens in the same place, they can all participate in the firing "race" in parallel. Similarly, the same token can be involved in a race with several transitions. If a token-transition pair is disabled, the firing is preempted (however, the aging label of the token can change as a result of being enabled for a finite amount of time).
- The delays can be deterministic (including zero delay) or follow any specified random distributions.
- The performance of the system is based on the statistical properties of marking in the system, and can be evaluated using discrete-event simulation or differential equations (including, but not limited, to finite-difference solutions). "Sensors" at each place can evaluate the chances and the number of times a given threshold of the

- number of tokens is crossed, or evaluate the time-averaged number of tokens at a given place. In the latter case, the correlation matrix for all results can be evaluated as well, providing the mechanism for calculating the variances (in addition to the mean values) of global metrics that aggregate the readings of individual sensors.
- Hierarchical constructions for combining multiple subnets are used to model large-scale systems.
 - Fusing places, commonly used in hierarchical Petri nets (see for example, (Jensen, 1993)) are employed to connect different parts of the model. Fused places appear as distinct graphical entities during the model construction, but represent the same entity in simulation.

Alprin's Simulation Parameterization

The following problem is considered following Alprin (1975). There are 30 routes divided into multiple segments. Each segment is defined as a single unit of a street that can be fully treated in both directions by a single truckload of salt (four tons). There are also eight-ton trucks that can cover two segments at a time before reloading. For regular streets, a single segment is two miles long, and for highways it is one mile (as the highways require twice as much salt).

The time to complete a segment is considered to be the same for regular streets and highways. There is a single salt pile location where trucks need to travel to reload. Table 1 shows the number of segments and average travel time to/from each location to a given route.

Table 1: Transit time and number of segments for each route

Route	Transit Time Minutes	N of segments
1	1.5	2
2	4.8	4
3	7.0	3
4	11.0	5
5	16.5	6
6	21.5	4
7	25.5	4
8	29.5	4
9	33.5	4
10	37.5	4
11	41.5	4
12	45.5	4
13	49.5	4
14	53.0	3
15	57.0	5
16	62.0	5
17	66.5	4
18	70.5	4
19	75.0	5
20	79.5	4
21	83.5	4
22	88.0	5
23	92.5	4
24	97.5	6
25	114.5	4
26	119.0	5
27	124.0	5
28	102.0	6
29	106.0	10
30	110.5	8

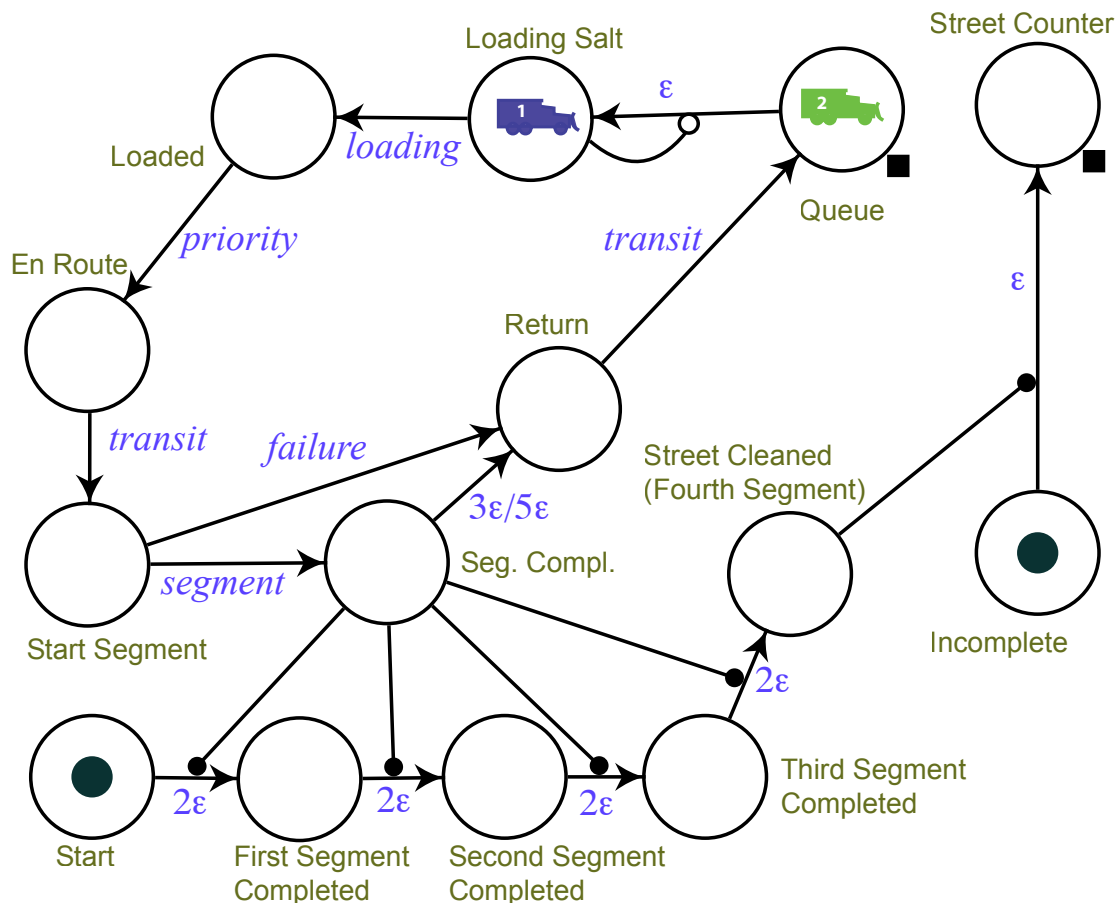
There are 27 four-ton and three eight-ton trucks available. The loading time for each truck and duration of each segment are parameters of the models. As discussed in the next sections, various operational logics for vehicle utilization are considered. In addition, the possibility of failure of the truck is included in the model (with the failure rate as a model parameter).

Replicating Alprin’s Simulation Logic

Two different APN models are created reflecting different vehicle utilization logic. Case 1 corresponds to a rigid assignment of vehicles to each route. In contrast, Case 2 corresponds to flexible vehicle utilization where trucks are assigned to different routes dynamically as needed.

A portion of the Case 1 APN model is shown in Figure 1. Only one street is shown with the full model consisting of 15 pages with two streets per page. Sub-models are connected by means of fusing three places: “Loaded”, “Queue”, and “Street Counter”. As a result, tokens representing each vehicle are merged at the “Queue” place, indicating that there is a common loading facility. When the treatment of a street is complete, a token is deposited into “Street Counter,” so associating a sensor with that place enables one to monitor the overall progress of operations (see the results section). In Case 1 each vehicle is represented by a token with a unique color (i.e., an integer label $i=0, \dots, 29$). The “priority” transition is color sensitive (so that only tokens with the matching color can move along the transition corresponding to a given route).

Figure 1: Single route portion of APN model (dedicated vehicle allocation)

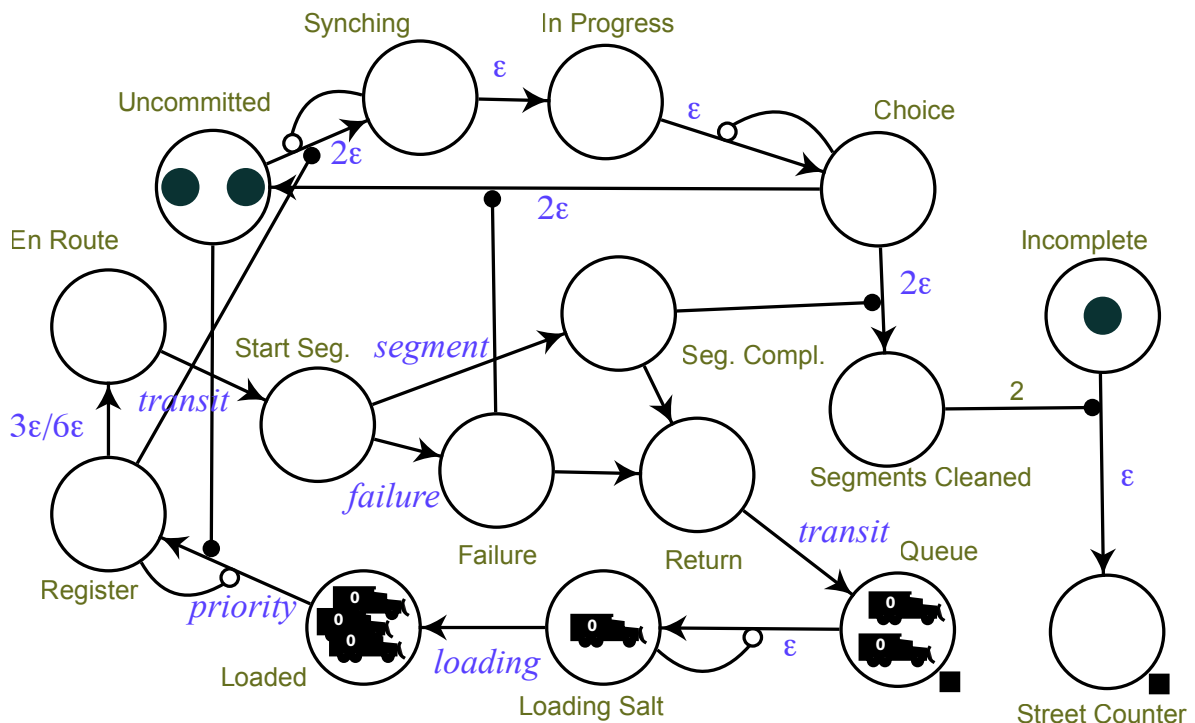


Once the vehicle is committed to a route the corresponding token moves to the “Start Segment” place. When the segment is completed, the token moves to the “Seg. Compl.” place. The token stays there just long enough to facilitate a move of the token representing the progress along a given route (depicted along the bottom of the diagram in Figure 1, note the corresponding enablers).

For the routes where four-ton trucks are employed, the token stays in the “Seg. Compl.” place 3ϵ while it takes 2ϵ to move the progress token one spot to the right. For the routes with eight-ton trucks the token stays in the “Seg. Compl.” for 5ϵ units of time, ensuring that two segments of the route are accounted for. Here ϵ is a small fixed delay (in the implemented model $1e^{-7}$ minutes is used). The model depicted has four street segments to clean/treat, so for each route the number of segments (and the number of places that lead to the “Street Cleaned” place) varies as shown in Table 1. If a vehicle fails during the segment cleaning, the vehicle returns to the base (the salt pile location) and no progress for that segment is recorded (it is assumed that the problem is fixed at the base and it takes as much time to fix the problem as to fill the truck with salt). Alternative repair times for failed trucks can be easily implemented.

Figure 2 shows the APN model for Case 2. As in the Case 1, this diagram shows one route out of 30 modeled. The main difference is that in this model several vehicles can attend to the same route. The number of these vehicles is controlled by the presence of tokens in the “Uncommitted” place. The priority transition controls the priority for selecting routes. At any given moment the number of vehicles committed to a given route can be less than or equal to the number of untreated segments on that route. When a token representing a vehicle moves to the “Register” place it triggers an outflow of tokens from the “Uncommitted” place. In the considered model there are two colors of tokens, where the color represents the type of vehicle (four- vs. eight-tons).

Figure 2 Single route portion of APN model (flexible vehicle allocation)



The fixed delay for transition from the “Register” to the “En Route” place is color dependent: for regular four-ton trucks (color 0) the delay is 3ε , so that only one segment is committed, and only one token will be moved from the “Uncommitted” to the “Synching” place, for tokens of color 1 (representing eight-ton trucks) this delay will be 6ε enough to move two tokens out of the “Uncommitted” place). A similar color-dependence is implemented to account for either one or two tokens when the outcome of covering a given segment is determined (see the “Choice” place): if the segment is successfully treated one token for color 0 and two tokens for color 1 are moved to the “Segments Cleaned” place. On the other hand, if a vehicle failure occurs then those tokens move back to the “Uncommitted” place. The model shown in Figure 2 is for two segments, but unlike the APN model for Case 1, adding segments does not require any addition of places and transitions: for k segments there should be k tokens in the “Uncommitted” place and the enabler originating from the “Segments Cleaned” should have multiplicity k .

RESULTS

A sample of the results is discussed below. Here each street segment is completed on average in 15 minutes with a standard deviation of 2.186 minutes (see Alprin, 1975) and a lognormal distribution is used. There are no vehicle failures and the salt pile loading time follows a lognormal distribution with mean value of 1.5 minutes and standard deviation of 1.5 minutes.

Case 1 – Assigned Routes

Figure 3: Results for Case 1 – assigned routes

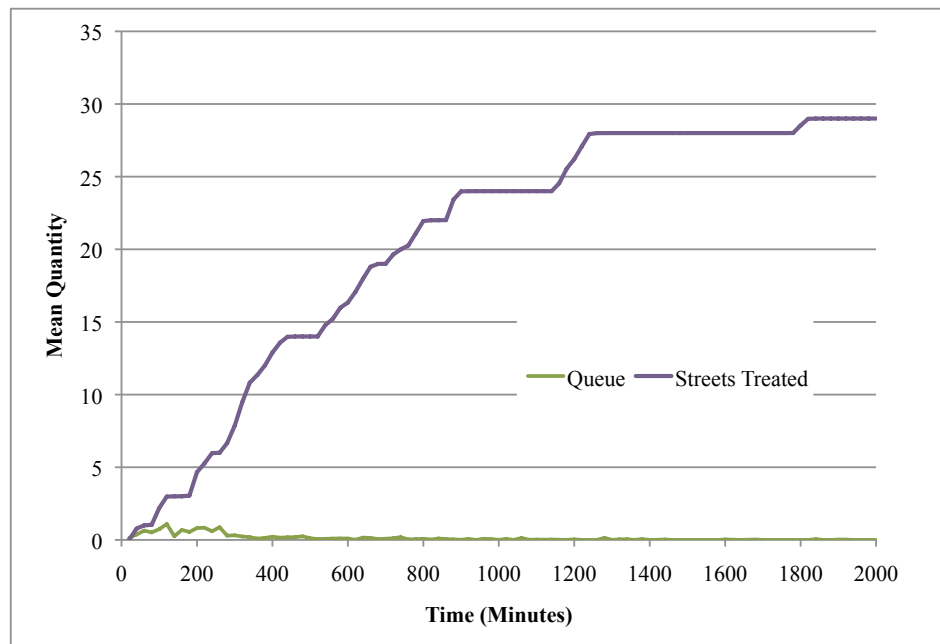


Figure 3 shows the expected number of completed routes as a function of time as well the mean amount of time vehicles spend in the queue for loading. One can observe that for the parameters specified the loading and waiting time has a negligible effect. One street (route 29) has not been completed even after 2000 minutes (i.e. 33 hours). This is expected as using deterministic durations for this route and ignoring the queuing effects for salt loading leads to an

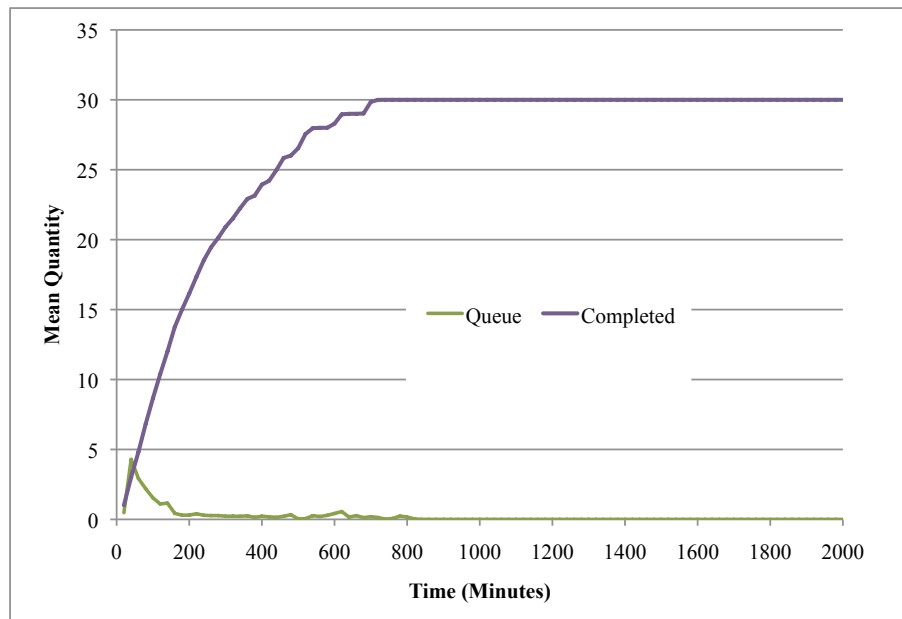
estimate of 2270 minutes for the completion of the route. Shown here and below the results of one million Monte Carlo simulation runs, which takes under 10 minutes on a Mac Air (2 MHz i7).

Case 2 – shortest distance route assignment

Figure 4 shows the expected number of completed routes as a function of time as well the mean time vehicles spend in the salt pile queue for loading. In this case, all streets are expected to be treated after about 720 minutes (12 hours). Again, these results match quite well the deterministic calculation that ignores the effect of queuing or any other inefficiencies, and thus provides an effective lower bound. This calculation simply adds up all deterministic times for travelling and dividing them among available trucks, yielding 625 minutes for the completion of all tasks.

In comparing the APN simulation results to those presented in Alprin (1975) it is clear that while the ratio of improvement is similar (or even higher), the results obtained indicate that it will take significantly longer to complete the treatment of the streets than the original simulation conducted by Alprin. The 12 experiments conducted by Alprin resulted in 8:08 hours and 5:14 hours average times for cases 1 and 2, respectively. The sources of these differences are likely to reside with the Alprin's deployment of input parameters and can be easily resolved for specific applications.

Figure 4: Case 2 - dynamic vehicle routing



Extensions

Next, let us consider sensitivity of the results of Case 2 to the loading time. Figure 5 compares the expected number of vehicles in the queue for loading as a function of time. The baseline parameters are compared to when the mean is doubled while keeping the standard deviation as the baseline, and also to the case where both mean and standard deviation are doubled. One can readily observe that the mean is significantly more important than the standard deviation.

Figure 5: Considering the loading time impact on Case 2 queue length

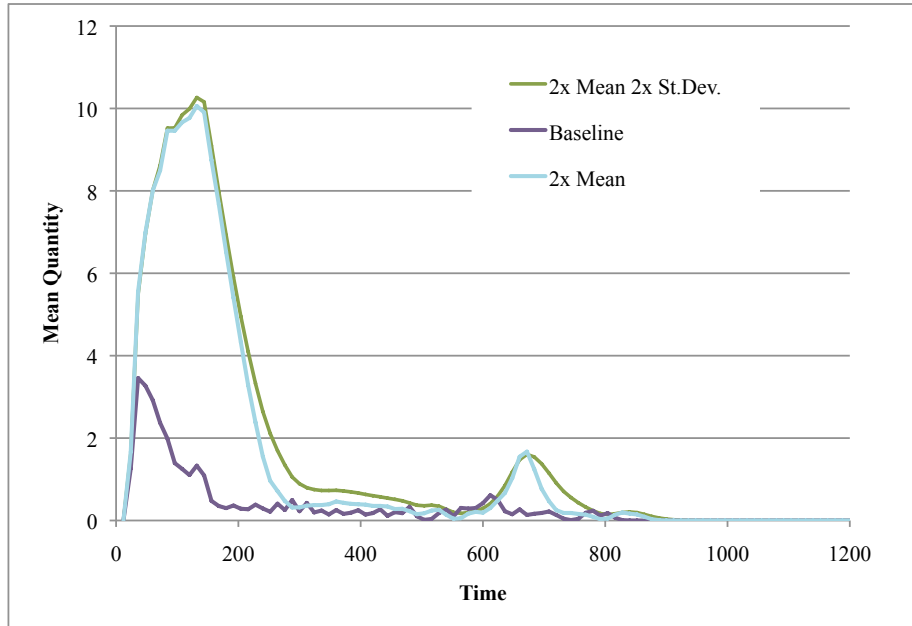



Figure 6 shows the impact of those differences on the number of routes completed as a function of time.

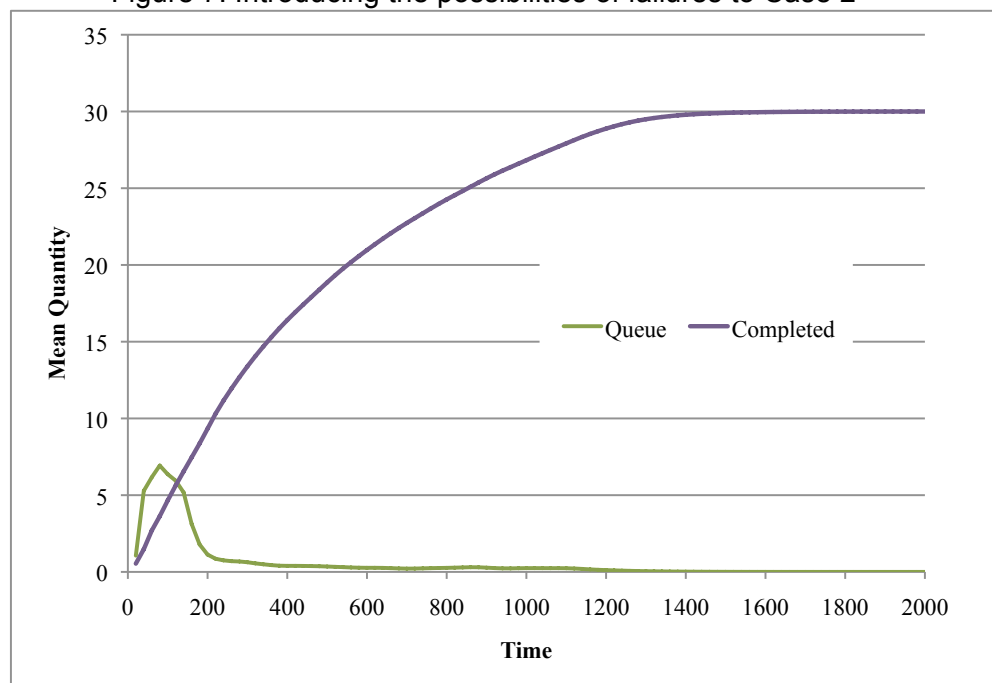
Figure 6: Considering the loading time impact on Case 2 route completion



Finally, let us consider the possibility of vehicle failures for the shortest distance route assignment scenario. Figure 7 shows the results when vehicle failures follow an exponential distribution with mean value of 30 minutes. One can observe the following three effects:

- Despite retaining deterministic transit times, the introduction of random vehicle failures provides the “mixing” to the associated process, smoothing out any deterministic effects
- Since the failure rate is fairly high, the impact on the completion time is also very significant (the time effectively doubles)
- Although queuing remains relatively insignificant it too has noticeably increased.

Figure 7: Introducing the possibilities of failures to Case 2



DISCUSSION AND CONCLUSIONS

The APN simulation model replicates Alprin’s (1975) original snow removal simulation and quantifies the performance of a fleet of snowplows. As shown above, the simulation results mirror the general findings and conclusions that Alprin presented in his thesis, while the graphical nature of the current model allows exploration of various “what-if” scenarios on the fly.

From a broader perspective, given these results, it is clear that modern discrete event simulations can offer governmental transportation planners a great deal in terms of effectively and efficiently employing their snow removal assets in the wake of a snow and/or ice event. Furthermore, the visualization capabilities of modern discrete event simulations enable vehicle managers to rapidly uncover where bottlenecks are likely to occur (e.g., at the salt pile) as well as the tyranny of distance in terms of getting even the most remote street segments cleared and safe for passage in a timely manner.

Clearly, this study has barely scratched the surface of the possibilities that modern climatological, GPS and GIS data systems, used in conjunction with state-of-the-art simulation

capabilities offer urban planners. In the course of this research, a number of additional recommendations for future research became apparent:

1. The snowplow simulation, like any model, is very data dependent. While this research employed Alprin's parameterization (to the degree that the parameters were documented), it would be very useful to replicate his data from another locale and repeat this analysis.
2. It would improve the model's validity if the historic data collection from step one could be associated with the actual time required to clear/treat roadways in the locale under study. This would be essential for truly gauging the accuracy of the model as it currently exists, as well as helping to identify shortfalls in the model's design.
3. The general APN snow removal model is very robust and could also serve future researchers well as an unbiased "sandbox" where alternative vehicle routing schemes could be evaluated.
4. Finally, this model could be adapted to assist urban planners in terms of making strategic, tactical and operational resourcing decisions. Employing a simulation such as this would bring a level of consistency and reproducibility to local government resource managers that promises an even better return on taxpayers' dollars.

REFERENCES

- Alprin, C. (1975). "A Simulation Approach to Solving the Snow and Ice Removal Problem in an Urban Area." Master's Thesis, Department of Management and Marketing, The University of Tulsa, Tulsa, OK.
- Bowden, F. (2000). "A Brief Survey and Synthesis of the Roles of Time in Petri Nets". *Mathematical and Computer Modelling* 21:55–68.
- Brailsford, S. C. (2008). "System Dynamics: What's in it for healthcare simulation modelers". In *Proceedings of the 2008 Winter Simulation Conference*, edited by S. J. Mason, R. R. Hill, L. Monch, O. Rose, T. Jefferson, and J. W. Fowler.
- Cook, T. M. and B. S. Alprin. (1976). "Snow and Ice Removal in an Urban Environment." *Management Science*, 23:227-234.
- Coordinated Highways Action Response Team (CHART). (2014). "Atmospheric Data Definitions." Maryland Department of Transportation. <http://www.chart.state.md.us/default.asp>.
- Department of Sanitation – New York City (DSNY). (2014). "About PlowNYC." New York, NY. <http://www.nyc.gov/html/dsny/html/special/DSNYSnowcategory.shtml>.
- Evans, J. R. (2012). "Business Analytics: The Next Frontier for Decision Sciences," *Decision Line*, Vol. 43, No. 2 (p. 4).
- Jensen, K. (1993). *Coloured Petri Nets with Time Stamps*. Computer Science Department, Aarhus University, Denmark.
- Lewis, A. (2013). "Track New York City's snow plows." Fox 5. New York City, NY. <http://www.myfoxny.com/story/24217904/track-new-york-citys-snow-plows>.
- Marsan, M. A. (1990). Stochastic Petri nets: An elementary introduction, Volume 424 of Lecture Notes in Computer Science, 1–29. Springer.

-
- National Center for Atmospheric Research (NCAR). (2013). "New Technology Targets Slick Roads." Boulder, CO. <https://www2.ucar.edu/atmosnews/news/10799/new-technology-targets-slick-winter-driving>.
- Pearlstein, S. (2014). "You're all too quick to surrender to snow: The case for digging out." *The Washington Post*. Washington, DC.
- Perrier, N., A. Langevin, and J. F. Campbell. (2007). "A survey of models and algorithms for winter road maintenance. Part III: Vehicle routing and depot location for spreading." *Computers and Operations Research*. 34:211-257.
- Rollison, T. (2014). "Winter Weather Impacts VDOT's Snow Budget." Virginia Department of Transportation. Richmond, VA. http://www.virginiadot.org/newsroom/statewide/2014/winter_weather_impacts_vdots71445.asp
- Scott, D. M. (2013). "Managing GIS Operations for Snow Removal for the City of Columbus-Warrior Style." ESRI ArcNews, Redlands, CA, <http://www.esri.com/esri-news/arcnews/winter1314articles/managing-gis-operations-for-snow-removal-for-the-city-of-columbus-warrior-style>
- Vock, D. C. (2014). "Potholes and Snow Plowing Punish State Budgets." PEW Charitable Trust, Washington, DC. <http://www.pewstates.org/projects/stateline/headlines/potholes-and-snow-plowing-punish-state-budgets-85899542334>
- Volovoi, V. (2013). "Abridged Petri Nets". ArXiv:arXiv:1312.2865.
- Volovoi, V. V. (2004). "Modeling of System Reliability Using Petri Nets with Aging Tokens". *Reliability Engineering and System Safety* 84 (2): 149–161.
- WTOP Radio. (2014). "Winter blows up local snow budgets." WTOP Radio. Washington, DC. <http://www.wtop.com/41/3574828/Winter-blows-up-local-snow-budgets>
- Yu, T.-T., J. P. Scanlan, R. M. Crowder, and G. B. Wills. (2011). "Approaches to Modeling the Gas Turbine Maintenance Process". *Journal of Computing and Information Science in Engineering* 12 (1).