Proceedings of the 2018 Winter Simulation Conference
M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, eds.

DESIGN AND IMPLEMENTATION OF DECISION SUPPORT FOR TRAFFIC MANAGEMENT AT MULTIPURPOSE PORT GATES

Ketki Kulkarni
Hoong Chuin Lau
Hai Wang
Sathyavarathan Sivabalasingam
Khiem Trong Tran

Fujitsu-SMU Urban Computing & Engineering Corp. Lab
Singapore Management University
71 Stamford Road
178895, SINGAPORE

ABSTRACT

Effective traffic management can help port operators gain a competitive edge in service level and efficient use of limited resources. One critical aspect of traffic management is gate operations management, ensuring a good customer experience to logistic carriers and considering the impact of congestion in and around the port. In this paper, we describe the design and implementation of a decision support tool to help gate operators plan for future scenarios with fluctuating demand and limited resources. We propose a simulation-optimization framework which incorporates theoretical results from queuing theory to approximate complex multi-lane multi-server systems. Our major contribution in this paper is the demonstration that the proposed design, when coupled with real data, can indeed help port operators improve their performances. To provide concrete real-world evidence that such technology has benefits, we have tested the system operationally since December 2017 and present the results and analysis in this paper.

1 INTRODUCTION

Port operations and logistics are integral parts of global supply chains. There are many subsystems of a port that may be optimized as smaller but decomposed units. For example, ship berthing and yard usage have been studied a great deal in literature (Bierwirth and Meisel 2015). An interesting, yet complex subsystem of the port is the gate system, which is made up of a combination of multiple gates in parallel or in sequence or both. The gates are a crucial link and a potential bottleneck that may lead to large queues, running up to a few kilometers either in or out of the port. On one hand the gate operators have to ensure a good service experience to their customers (logistics carriers); on the other hand they have to adhere to strict safety procedures perhaps even relevant to national security. For many Asian countries with busy ports, port infrastructure is a scarce resource. When facing high congestion under high demand, increasing the number of lanes at a gate or increasing the number of gates itself is not often a feasible solution due to land or budget constraints. In addition, the staffing of gates requires trained labor, which is also an expensive resource. Hence, optimization of gate operations, particularly lane management, is an important area to explore, considering its impact on overall port operations.

We first introduced this problem in our previous work (Kulkarni et al. 2017). The first part of this project, presented in the previous work, detailed our study to understand the gate operations through historical data analysis and field surveys. It explained how the discrete-event simulation model was built. Specifically, we described how a simulation-optimization heuristic approach was applied to solve the optimization problem of lane management. In this paper, we describe the second part of this project, that is, how the simulation model has been incorporated into a decision support tool, along with an optimization solver. We have developed a mathematical model that solves the problem exactly, using an Integer Linear Programming (ILP) solver. This exact formulation is part of the engine of the tool. We also describe how this tool was designed, implemented and finally tested under multiple scenarios at the port. The results from the field trials are analyzed to quantitatively demonstrate the value of the tool. We now provide a brief description of the system that is being modeled.

The port in our study uses different freight vehicles (trucks) to bring cargo to and from vessels (ships). There are multiple gates at different locations for entry and exit of these freight vehicles. Since this is a multipurpose port, the vehicles are varied. The services required by a vehicle depend on the type of cargo they carry and hence service time is a function of vehicle type. Each vehicle needs to undergo multiple kinds of services, such as driver identification and cargo verification. A checkpoint may provide only one kind of service at the gate and hence a vehicle is required to visit multiple checkpoints before entering or exiting the port. The choice of lane or gate does not affect the service time. The tool described in this paper has been developed for deployment at the main gate of the port, for managing the exit lanes. These lanes are used by vehicles to exit the port, carrying various types of cargo. There are multiple identical lanes in parallel, each with multiple successive checkpoints. Every vehicle needs to pass through these checkpoints in sequence. More details on the layout are presented in Section 3.1. The rest of the paper is organized as follows. The next section presents a literature review of related work. Section 3 describes the tool and its various components. Sections 4 and 5 present the details of the field trials along with the analysis of collected data from these trials. Section 6 concludes the paper.

2 RELATED WORK

Optimization of port operations and the tools for port planning are gaining importance in research. As Lee et al. (2012) describe in their study, use of Information Technology (IT) based tools and services can lead to a better customer experience, while ensuring that all regulations are complied with. For example, informing logistics carriers of expected latency time at gates while they are still inside the port, can help them plan their routes or breaks. The authors provide a useful conclusion based on empirical studies, that planning tools or e-Transformation can help ports in retaining competitiveness while ensuring customer satisfaction. In our previous work (Kulkarni et al. 2017), we conducted a literature review on diverse streams of research relevant to port operations and in particular, to port gate simulation and optimization. In this paper, we refer to other efforts on developing mathematical models along with simulation techniques to solve real problems. We are especially interested in cases where the research has been applied using actual data.

Melamed et al. (2016) describe a version of Little's law adapted for finite horizon scenarios. They discuss how the traditional form of the law, which is used under the assumption of long-term steady state behavior, can be adapted to analyze finite horizon data. Most real systems would require analyzing data over limited time horizons, since observation periods are finite. However, Melamed et al. (2016) do not apply the findings to a real case study. In our work, we do apply similar insights in analyzing the vehicles in the system over finite horizons. This is described in detail in Section 3.2.

Kim and Whitt (2013) provide a useful statistical approach in estimating sojourn times, given the queue length and arrival rate. They use output from a simulation model to improve estimates. They compare estimates obtained from various methods, and observe that there is strong merit in statistical analysis using simulation models. They also use data from a real system to compare different approaches. However, while this approach is useful for analyzing data, it cannot be directly incorporated into an optimization framework. This is because using the expressions in an optimization model requires expressing the parameters in terms of decision variables. Often, the resulting constraints would introduce non-linearity, so that the optimization problem is hard to solve using existing LP solvers. In this paper, we describe some approximations essential for ensuring that the model remains "solvable" in reasonable time.

3 TOOL DESIGN

This section discusses the design and implementation of our decision support tool for gate system operations. Figure 1 shows the various aspects of the tool at a glance.

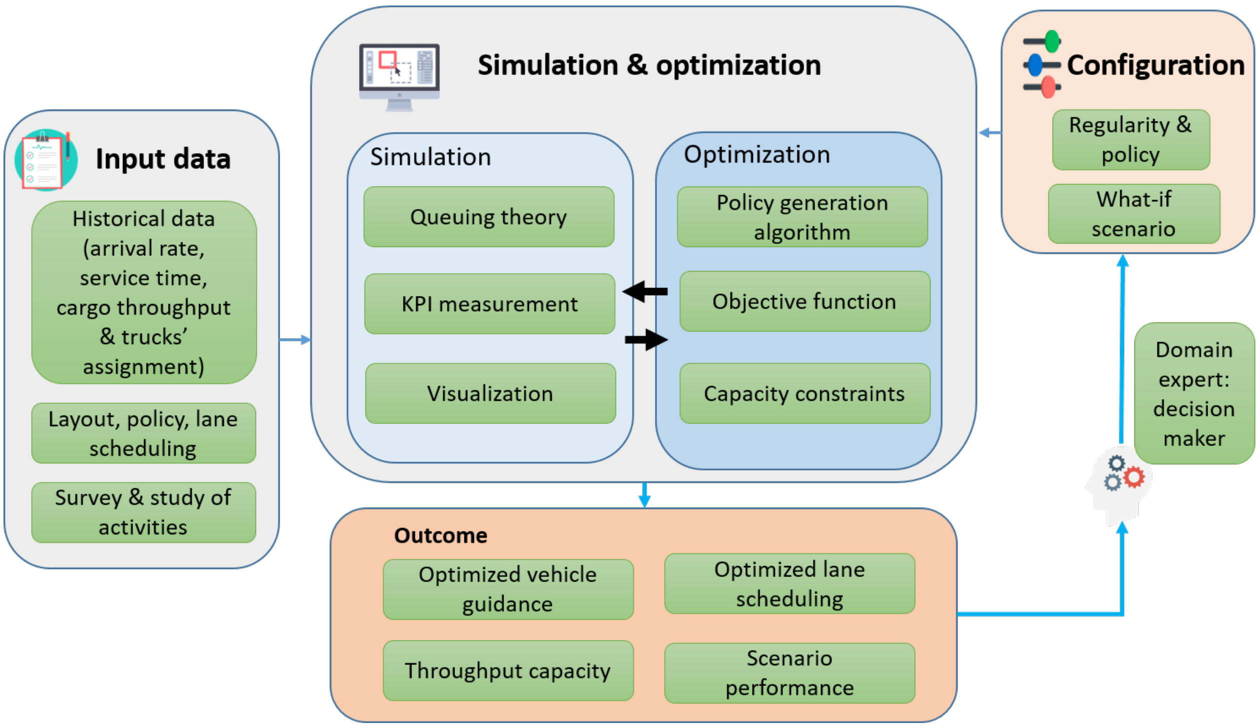

Figure 1: Tool design at a glance.

The inputs to the system include the layout of the gate, existing scheduling policies, historical data on vehicle arrivals and service rates, and other inputs from field studies. The engine contains two blocks: the simulation and the optimization units. The discrete-event simulator is built using Java libraries. The Key Performance Indicators (KPIs) are the average time spent in the system (sojourn time) and average waiting time. The optimization component generates a schedule that gives the number of lanes to be open in each hour of the day. It is coded in Java and run using the CPLEX solver. The objective function is to minimize the resources used, that is, the number of open lanes, while ensuring that the sojourn time is within acceptable limits. The output from the simulation and optimization units can be viewed using the tool's visualization panels, where the various KPIs are organized in different categories. The visualization has been developed using Bootstrap for user interface (UI) elements such as buttons, date picker and sliders, Plotly for drawing diagrams and jQuery for communicating with the engine (back-end) and updating UI dynamically. The tool also allows users for performing various what-if scenarios, by varying different input conditions manually. Multiple scenarios can be analyzed together for comparison. Each component of this tool is now described in detail.

3.1 Simulation Functionality

The simulation unit is a discrete-event simulator developed in Java using queuing libraries. The model is intended to capture the structure and operational practices at the port gate, along with stochastic elements such as vehicle arrivals and service rates. Figure 2a shows a typical gate structure. The gate consists of multiple identical lanes in parallel, each with several checkpoints in series. The choice of a lane by the vehicle does not affect the service rates. The simulation not only captures the current layout of the port gate, but is also flexible enough to incorporate future changes, such as the number of checkpoints or number of lanes or the order of the checkpoints. This model is capable of being extended to simulate a larger, more complex gate system as well, involving more than one gate. The multiple-gate system can be modeled as a network of gates, where vehicles may be required to enter another gate after exiting the previous gate, as shown in Figure 2b. This is useful for the port management to study and understand the impact of any planned infrastructural changes to the gate.

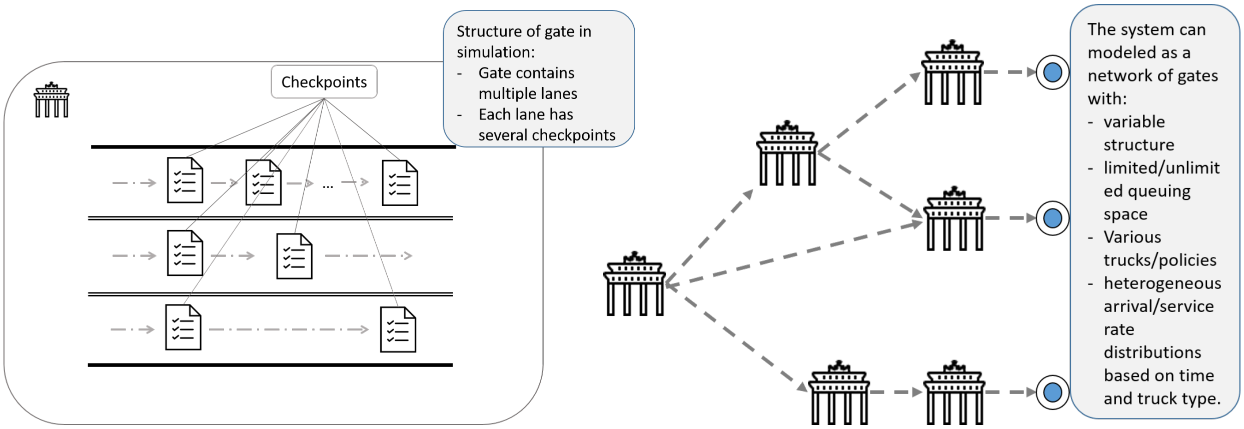

(a) Simulation structure. (b) Network of gates.

Figure 2: Gate structure and network of gates.

The simulator can work as a standalone functionality or in combination with the optimization unit, simulating the output of the optimizer. The optimizer has been integrated into the simulation framework to facilitate usage of common inputs of service and arrival parameters.

3.2 Optimization Functionality

The second main component of the tool is the optimizer which solves the lane schedule optimization problem for a period of 24 hours, generating the number of lanes required to be open in each hour. The complexity lies in the fact that it is a multi-period problem where vehicles that do not finish their service in a given time period continue their service in the next subsequent time period. This section presents some discussion on the mathematical formulation. Details of how this model has been developed are beyond the scope of this paper, which focuses on how the techniques have been implemented as a tool and tested at an actual port.

The system under consideration is stochastic in terms of the arrival of different types of vehicles in each hour of the day and the service rate at the various checkpoints in the system. While these aspects are quite easily captured by the simulation model, suitable approximations need to be made in order to operate with an exact optimization model. The planning horizon is divided into 24 periods and the vehicle arrivals are approximately Poisson processes in each hour. Since the system is stochastic, it is impossible to know the precise number of vehicles in each hour or the exact service time of each vehicle. Therefore, in order to have a "solvable" mathematical model, an approximation of these two parameters are required. As mentioned previously, the optimization is a plug-in in the simulation environment and has access to the common inputs. We use the Random Number Generator (RNG) in the simulator to generate realizations of random variables that represent the stochastic arrival and service processes, for any given time period t in the planning horizon. This means, the optimization model is solved for a given number of vehicles of different types arriving in each hour, where the number is obtained by sampling from the distribution representing the stochastic arrival process. The formulation is now described in detail, starting with the

Table 1: Notations and symbols.

Parameters		**Decision Variables**	
β_l	Cost of opening lane l	x_{tl}	1 when lane l is open in time t, and 0 otherwise
α_t	Penalty for unserved vehicle in time t	η_t	Number of vehicles not served at the end of time t
λ_v^t	Number of vehicles of type v arriving in time t (n)	y_{tlvn}	1 if vehicle n of type v is assigned to lane l in time t, and 0 otherwise
W^*	Acceptable sojourn time	Z_v^t	Sum of vehicles arriving in t and unserved vehicles from $t-1$
W_{lt}	Average sojourn time in lane l in time t		

notation and symbols described in Table 1. The set of all lanes is L, set of time periods is T and set of vehicles is V.

$$\min \sum_{t\in T}\sum_{l\in L}(\beta_l * x_{tl} + \alpha_t * \eta_t) \tag{1}$$

$$\text{subject to} \quad \sum_{l\in L} y_{tlvn} \le 1 \;\; \forall t\in T, v\in V, n=1,\ldots,Z_v^t \tag{2}$$

$$\sum_{l\in L} x_{tl} \ge 1 \;\; \forall t\in T, \tag{3}$$

$$\sum_{v\in V}\lambda_v^t - \sum_{l\in L, v\in V}\sum_{n=1}^{Z_v^t} y_{tlvn} = \eta_t \;\; \forall t\in T \tag{4}$$

$$Z_v^{t+1} = \lambda_v^{t+1} + \eta_t \;\; \forall v\in V, t\in T \tag{5}$$

$$W_{lt} \le W^* \;\; \forall t\in T, l\in L \tag{6}$$

$$\eta_t \ge 0 \;\; \forall t\in T \tag{7}$$

$$x_{tl}, y_{tlvn} \in \{0,1\} \forall t\in T, l\in L, v\in V \tag{8}$$

The objective function (1) minimizes the number of open lanes in each hour and penalizes the number of vehicles not served at the end of each time period. We consider all lanes to be identical and hence the operating cost of any lane is the same. However, to ensure that a minimum number of lanes are opened in each hour, we use β_l. For the first lane, β_1 is 1, which ensures that it is always open. For each subsequent lane, $\beta_l = 100 * \beta_{l-1}$. Therefore, the cost of opening an additional lane when one lane is already open, is quite high (for modelling purposes). Compared to β_l, the values of α_t are much smaller. This term penalizes any vehicles left without service in time period t. When the traffic is significantly low, for example, around early morning (1am-5am), the model is expected to open lesser lanes. During peak hour traffic, the model suggests a higher number of lanes. However, during intermediate traffic hours, there is a trade-off between opening new lanes versus allowing vehicles for carrying over into the next time period. The parameter α_t is useful to control the model behaviour during such hours.

Constraint (2) ensures that each vehicle is assigned to at most one lane. Constraint (3) ensures that at least one lane is open in any given hour of the day. Constraints (4) and (5) count the number of vehicles that arrived in time period t but cannot be served and carried over to the next subsequent time period. Constraint (6) enforces that the average observed sojourn times are less than the system threshold for sojourn time for each lane.

We will now briefly describe how the sojourn times are approximated for each lane. From Gautam (2012), the expression for average queue length for an M/G/1 queue is obtained as shown in Equation (9) where λ is the long run arrival rate, $\rho = \frac{\lambda}{\mu}$ is the utilization, mean service time is $\frac{1}{\mu}$ with variance σ^2.

$$L = \rho + \frac{\lambda^2}{2}\frac{\sigma^2 + \frac{1}{\mu^2}}{1-\rho} \tag{9}$$

From Little's Law, it is known that $W_{lt} = \frac{L_{lt}}{\lambda_{lt}}$ in the long run. Combining the two formulas, the average sojourn time can be expressed in terms of parameters λ and μ as shown below.

$$W_{lt} = \frac{1}{\mu} + \frac{\lambda_{lt}}{2}\frac{(\sigma^2 + \frac{1}{\mu^2})}{1-\rho_{lt}} \tag{10}$$

In Equation (10), λ_{lt} represents the number of vehicles assigned to lane l in time t and ρ_{lt} is the utilization for lane l. Note that unlike λ the parameter μ is not time-varying and is identical across lanes as well. Also, a single-server approximation is made to represent the multiple servers (checkpoints) in a lane. Since service time depends on vehicle type, the single-server approximation $\frac{1}{\mu}$ depends on the mix of vehicles assigned to the lane.

Lastly, we end this discussion on optimization by presenting the ranges and some likely values for the parameters mentioned in the model in Table 2. For the output and intermediate parameters like η_t and Z^t, the values mentioned are summarized observations obtained after running the model, for the given weeks of trial. The range mentioned for λ_v^t is across all vehicle types. The ranges for individual vehicle types are in between this range.

Table 2: Likely values for model parameters.

Sets		**Parameters**		**Intermediate ranges**	
L	1,...,6	W^*	5,...,10 min	λ_v^t	1,...,100 (for different v)
T	1,...,24	α_t	1,...,10	η_t	0,...,20
V	1,...,4			Z_v^t	10,...,250

3.3 Role of Queuing Theory and Simulation

While simulation is useful in performing what-if analyses, in order to get recommendations for lane schedules, an optimization framework is required. To capture and integrate the parameters related to waiting, service, and arrivals within the optimization model, we use queuing theory as described in the sections so far. To make this model computationally efficient, certain approximations are made, so that closed form queuing expressions can be used. This helps to ensure that the mathematical formulation remains linear, thereby increasing computational tractability. However, in reality, the system is not M/G/1. It consists of multiple checkpoints (servers) in each lane. Further, there is a finite space between the checkpoints, where vehicles may wait. Capturing a system with multi-class non-homogeneous Poisson arrivals and multiple sequential servers with spaces in between them using analytic expressions is extremely challenging and also results in highly non-linear expressions. These effects have been modeled with relative ease in the simulation framework. Hence the role of queuing theory is to efficiently find solutions (lane schedules) while making suitable approximations, and the role of simulation is to evaluate the efficacy of these solutions (as well as other scenarios) against the 'real' system.

The process of validating results using simulation is now briefly described. The lane schedule from the optimization model is tested by using it as an input in the simulation model along with the expected vehicle arrival patterns obtained from historical data inference. Given the various sources of stochasticity (service times at multiple checkpoints and arrival rates for different types of vehicles) as well as the finite spaces in between servers for waiting, the solution robustness is checked by running the simulation over multiple replications. The waiting times reported by the tool are averages obtained over these multiple replications on multiple runs. This analysis then allows the user for viewing expected hourly performance. In case the users are not satisfied by the expected waiting times in a particular hour, they have the option of making manual edits to the schedule in that hour. Following this, the modified schedule may be simulated again, to ensure that the desired performance is obtained in each hour.

The third main component of the tool is the KPI calculation and visualization. An analysis dashboard is developed for this purpose. This is described in the next subsection.

3.4 Analysis Dashboard

An important feature of any simulation tool is capturing the performance metrics for comparison and analysis. As mentioned earlier, the sojourn time of vehicles is an important metric. There are four main panes on the tool's panel. The first one gives the summary of vehicles that were simulated over 24 hours under this scenario. The vehicles are arranged by vehicle type. The second pane gives a visual display of sojourn times for each lane over 24 hours. A darker red or pink color indicates higher sojourn times. The third pane displays the lane operations schedule that was simulated in this scenario, indicating how many lanes were open in each hour. Finally, the last pane shows the sojourn times. The default view shows the average sojourn time in each hour for all vehicles in the system. The user has the option of selecting a particular vehicle type to view the average sojourn time experienced by vehicles of the particular type. For example, the tool allows the user for seeing the average time spent by vehicles of type T1 in the system. Examples of these data are shown as part of the analysis of field trials in the subsequent sections.

4 FIELD TRIAL

We conducted a field trial in a local port to verify our system. Firstly, the simulation module was tested, by observing how effectively the tool captured the system behavior, given the input data. Secondly, a lane management schedule, indicating the number of lanes open in any given hour, was generated by the optimization tool based on observed and simulated behavior for a week. The schedule was recommended to the planner for ground execution and the performance of the schedule was quantified based on the average time in system for different vehicle types, as well as the number of lanes used for the trial period. The trials were carried out in two phases. In the first phase, the users (gate planners) were assisted in operating the tool. The second phase of the trial involved independent use of the tool by the operators. Both these trials are described next.

4.1 Assisted Field Trial

This trial was carried out over four days of two successive weeks. The data from the first week were used to create the baseline scenario, based on which recommendations were generated using the optimization tool for the second week. The KPI were observed over both weeks to test the effectiveness of the recommendations. Two days (referred to as Day 1 and Day 2) of each week were selected as days of observation, since prior data analysis revealed higher traffic on these days of the week. After the data collection from survey for two days of week 1, the data were analyzed and compared with the historically available data for the same period. It was observed that the current demand patterns at the port were higher than the historical values. Particularly, there was a higher traffic of T1 vehicles than usual. The input to the simulation was based on the current traffic trend, not the historical values. No records are available for service times at certain checkpoints. The simulation tool works on the service times based on surveys conducted over different months in 2017. During the current trial, these values were observed again, to check for any substantial changes in service behavior. It was noticed that while most values remained close to previously noted values, the service times for T2 and T3 did show changes. These values were modified to reflect the current practices.

After updating the inputs to the simulation model and verifying that current demand patterns were represented as accurately as possible, the tool was then used to generate schedule recommendations. These recommendations were forwarded to the users, who implemented them for the second week of the trial. After running the optimization model, the tool gave recommendations for lane opening and closing. These recommendations were then simulated multiple times to ensure that the average time in system is less than or equal to 5 minutes for any hour of the day. The tool further allows for manually improvising on the recommended schedule to reduce waiting time for certain hours of the day or to incorporate any operational features that are not in the current model. For example, certain checkpoint operators prefer to have a partial lunch hour between 1pm and 2pm. This was manually incorporated into the optimization-recommended schedule and thoroughly tested to meet system requirements. The recommendations were only generated for the hours of 8am–4pm due to the limitation on availability of volunteers for observation.

4.2 Analysis of Assisted Field Trial

In the first part of this section, we present the results from paired t-testing. For the given study, there were two sessions of observations on each day: morning (M) and afternoon (A). Around 400 vehicles were surveyed on each day. Also, as previously mentioned, there are 4 vehicle types under consideration. The data are sorted into categories based on observation session and vehicle type. Since it is impossible to compare values for each individual vehicle, averages for each "vehicle type-session" are compared instead. The data for Day 1 and Day 2 of week 1 are paired with the respective Day 1 and Day 2 of week 2. Table 3 shows the sorting of data and pairing of the days for testing.

Table 3: Paired T-test: Sorting data for comparison.

Vehicle	**Average sojourn time (sec)**		**Difference**	**Average sojourn time (sec)**		**Difference**
and slot	***Day 1-Week 1***	***Day 2-Week 2***	$d_1 - d_2$	***Day 1-Week 2***	***Day 2-Week 2***	$d_1 - d_2$
T1 - M	31	31	0	25	30	-5
T1 - A	28	21	7	26	21	5
T2 - M	171	123	48	165	149	16
T2 - A	141	171	-30	199	199	0
T3 - M	335	252	83	289	208	81
T3 - A	184	179	5	227	167	60
T4 - M	294	218	76	163	181	-18
T4 - A	161	103	58	104	86	18

The null hypothesis for this test is that any difference in observations from the weeks can be attributed to chance. That is, the data in week 2 are not significantly different from the data in week 1. The level of significance (α) for this test is 0.05. The p-values obtained were higher than 0.05. Therefore, the conclusion is that any difference in observation is likely due to chance and the difference in data is **not** statistically significant. This test helps us establish that the reduction in lane-hours did **not** worsen the system performance. Table 4 shows the comparison of recommended schedules with default operating practices. From the discussion so far, we have obtained that the difference in average sojourn times was not statistically significant. We further see that the recommended schedules use fewer hours to obtain this performance, and hence are more efficient than the default.

Table 4: Tool recommended schedule.

Regular	**Recommended schedule from 8am-4pm**			
Lane-hours	***Day of week***	***Proposed lane-hours***	***% saving***	***Sojourn time difference***
24	Day 1	21	12.5	not significant
24	Day 2	21	12.5	not significant

Although the results of statistical testing indicate that the observations over two weeks are not statistically different, visual inspection of data hinted that the values for Day 2 in week 2 are less than the corresponding values in week 1. We ran simulation experiments to understand this change better. Apart from the reduced lane-hour schedule, the recommendation for Day 2 also included shifting the lunch hour from the usual 1–2pm slot to 12–1pm instead. The simulation runs were set up as follows. Each simulation run was for 30 virtual days and 10,000 replications. For each scenario, 30 such simulation runs were performed. The average sojourn time obtained at the end of these 30 runs is reported in this section. For a given schedule, two different scenarios were tested; one with the lunch hour from 1–2pm and the other with the lunch hour moved to 12–1pm. In Figures 3a and 3b, the darker columns indicate the average time in system for vehicles observed for the schedule without shifting the lunch hour (default).

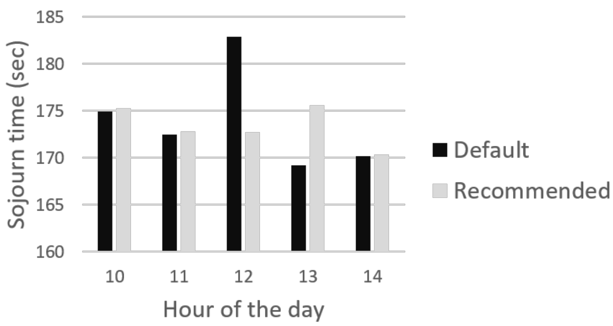

(a) T3 sojourn times.

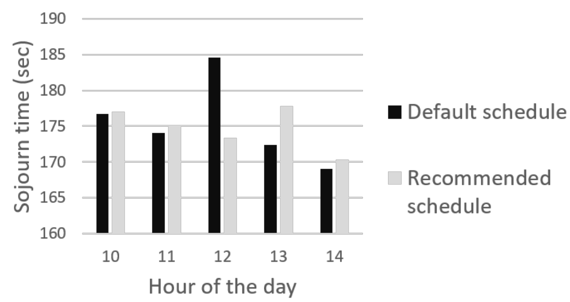

(b) T2 sojourn times.

Figure 3: Peak hour sojourn times.

The grey columns show the average time in system after moving the lunch hour (recommended). It is clear from the figures that minor shifts in lunch hours can bring down the average sojourn times around that time. This effect is higher for certain vehicle types. This was an important takeaway from the trial, since it showed that even if the total numbers of lane-hours are similar, impact can be obtained by rescheduling breaks.

5 INDEPENDENT FIELD TRIAL

In the second phase of the field trial, the tool was used by the user independently who had been trained on tool operations in generating schedules, importing new data, manipulating inputs, and using the analysis dashboard. This trial was conducted for a period of one week. The week of trial was a low-demand period for the port and hence expected traffic was less than normal. However, different days of the week allowed for experimenting with various scenarios. While three of the five days were regular days with normal operations, the gate operators decided to schedule planned maintenance activities on two days. Historical data analysis has shown that traffic is not uniform across the week; the two days chosen for maintenance were expected to have higher traffic as compared to other days. However, due to availability of maintenance staff, the downtime had to be scheduled on the relatively higher traffic days. This was useful in testing the tool's capabilities of helping understand impact of planned downtime.

5.1 Analysis of Independent Trial

We report the results obtained from the one week of trial. The baseline for comparison are the data from the previous week, since historical analysis and current trend study showed that these two weeks are likely to experience similar demand. The default lane schedules are the operational schedules that were used in the previous week. A discussion with the users confirmed that if not for the tool recommendations, they were likely to use the same schedules in the current week as well. While analyzing each day, three main metrics were compared between the two weeks: the number of open lanes, the vehicles arrivals, and the average sojourn time. Figures 4a and 4b show the analysis where operating schedules and vehicle arrivals are compared for the two Mondays. It is observed that the vehicle arrivals (or demand patterns) are fairly similar in both cases, validating the assumptions made by the user. The actual values have been re-scaled for confidentiality.

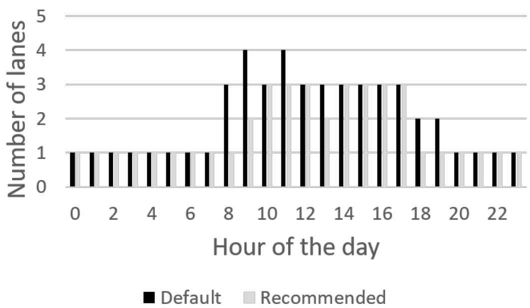

(a) Lane schedule comparison.

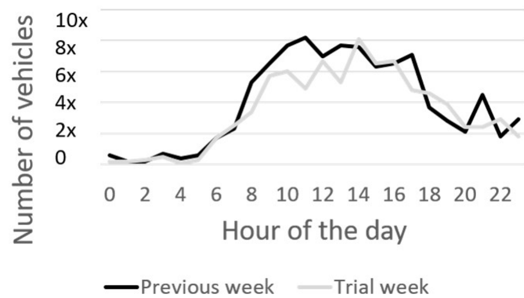

(b) Vehicle arrival comparison.

Figure 4: Comparison of simulation inputs.

Figures 5a and 5b show the comparison between the average sojourn times over 24 hours. It is observed visually that for the trial week, there are fewer peaks of sojourn time. However, this could also be attributed to the slight variation in demand patterns. Overall, the more reliable metric once again is that the difference in observations are not statistically significant. Similar analyses were performed for all days of the week and the tool was able to recommend schedules that were more efficient, without affecting service levels.

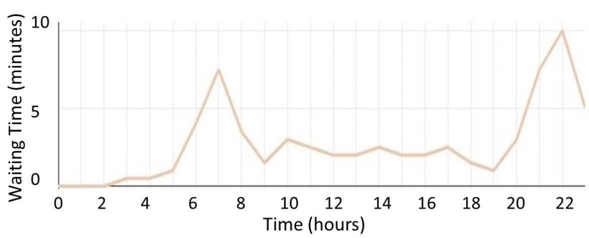

(a) Default schedule.

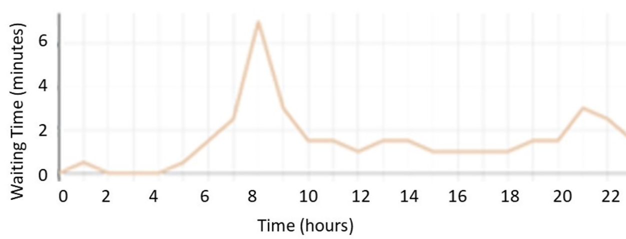

(b) Recommended schedule.

Figure 5: Analysis of average sojourn time from simulation.

Table 5 shows the summary of the schedule recommendations for the week of the trial. The days of

Table 5: Schedules generated by the user.

Regular Lane-hours	***Day of week***	***Proposed lane-hours***	***% saving***	***Type***	***Sojourn time difference***
48	Day 1	40	16.6	Regular	not significant
48	Day 2	42	12.5	Regular	not significant
44	Day 3	42	4.5	Planned maintenance	not significant
44	Day 4	42	4.5	Planned maintenance	not significant
52	Day 5	44	15.4	Regular	not significant

the week are numbered and presented out of order for confidentiality. Around 100 vehicles were surveyed during peak hours on each day. As is seen, the proposed schedules always recommended less lane-hours than the default operating practice. Day 1 recorded the highest resource saving of 16.6%.

5.2 Analysis of Planned Maintenance

This second phase of trial allowed for experimenting with planned downtime. This gave the user multiple options: either to manually enter a schedule and check for consequences without using optimization, or obtain an optimized solution and make manual adjustments to reflect the planned downtime. In either case, this is an iterative process, where the user fine-tunes the schedule. We had two such sessions where downtime was planned: Day 3 morning and Day 4 afternoon. We present the analysis of the Day 4 session. Figure 6 shows the average sojourn times evaluated by the tool after taking into account the planned downtime, along with the schedule and vehicle arrivals expected. The user had chosen to use the optimization feature and then manually close the lanes under maintenance. The planned downtime was between 1–4pm. The tool showed that high waiting times of 15–20 minutes may be expected between 2–3pm. It was observed that 13 vehicles experienced high waiting between 2–3pm (above 5 minutes).

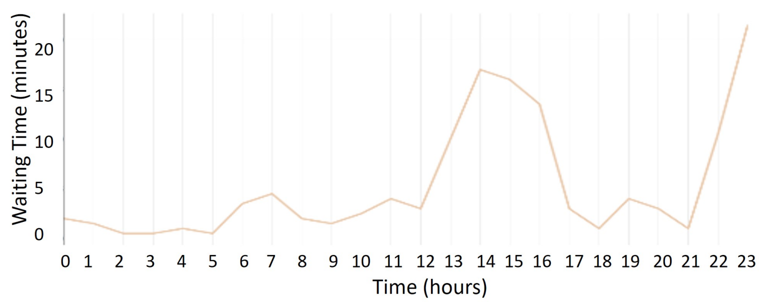

Figure 6: Average sojourn times for Day 4.

At least 5 vehicles waited for 10 or more minutes and the worst sojourn time observed was 20.43 minutes. The impact of the lane closure reduced significantly after 3pm, finally achieving required service levels around 3.30pm. This observed performance closely matched the scenario analysis by the tool and hence the user was able to prepare in advance for the situation. This was evident by the actions taken by the gate operators: knowing beforehand that vehicles could end up waiting 15–20 minutes during 2–3pm, they had prepared an additional lane in another area to facilitate service of certain vehicle types, so that long queue build could be tackled if needed.

6 CONCLUSION

The objective of this study was to develop a decision support tool for gate operators at a multi-purpose port. A simulation-optimization based tool was designed and implemented to suggest lane operating schedules based on demand and service patterns. Two rounds of field trials were conducted to gauge the impact and usability of the tool. The tool was able to provide recommendations that were more efficient than default practices, by suggesting less use of resources (lanes). Statistical testing ensured that service level was not adversely affected by reduction of resources. The tool also helped evaluate downtime scenarios, enabling the user to plan for unavoidable queues in advance.

7 FUTURE WORK

This tool and solution framework have been successfully adapted to suit the port operations. Work is ongoing to generalize the framework for other related applications such as airport immigration or fast food service counters. The exact mathematical model works reasonably well for small-scale systems. However, for potentially larger systems requiring a faster response, a heuristic-based approach is being worked on.

ACKNOWLEDGMENTS

This research is funded by the National Research Foundation Singapore under its Corp Lab @ University scheme and Fujitsu Limited as part of the A*STAR-Fujitsu-SMU Urban Computing and Engineering Centre of Excellence.

REFERENCES

Bierwirth, C., and F. Meisel. 2015. "A Follow-up Survey of Berth Allocation and Quay Crane Scheduling Problems in Container Terminals". *European Journal of Operational Research* 244(3):675 – 689.

Gautam, N. 2012. *Analysis of Queues: Methods and Applications*. Boca Raton, FL: CRC Press.

Kim, S.-H., and W. Whitt. 2013. "Statistical Analysis with Little's Law". *Operations Research* 61(4):1030–1045.

Kulkarni, K., K. T. Tran, H. Wang, and H. C. Lau. 2017. "Efficient Gate System Operations for a Multipurpose Port using Simulation-Optimization". In *Proceedings of the 2017 Winter Simulation Conference*, edited by W. K. V. Chan et al., 3090–3101. Piscataway, New Jersey: IEEE.

Lee, S.-Y., J. L. Tongzon, Y. Kim, G. Logistics, and C. Cheiljedang. 2012. "E-transformation in Port Management: An Empirical Investigation". In *Proceedings of the Second Annual International Workshop on Port Economics and Policy*, December 10^{th}-11^{th}, National University of Singapore.

Melamed, B., R. Leuschner, W. Chen, D. S. Rogers, and M. Cao. 2016. "Inventory Turns and Finite-Horizon Little's Laws". *Annals of Operations Research*:1–18. Published online doi: https://doi.org/10.1007/s10479-016-2157-9.

AUTHOR BIOGRAPHIES

KETKI KULKARNI is a Research Fellow of Fujitsu-SMU Urban Computing and Engineering Corporate Lab. She holds a Ph.D. in Industrial Engineering and Operations Research from the Indian Institute of Technology Bombay, India. Her research interests include discrete-event simulation and applied operations research. Her e-mail address is ketkivk@smu.edu.sg.

HOONG CHUIN LAU is Professor of Information Systems and Director of the Fujitsu-SMU Urban Computing and Engineering Corp Lab at the Singapore Management University. He currently serves on the editorial board of the IEEE Transactions on Automation Science and Engineering. His e-mail address is hclau@smu.edu.sg.

HAI WANG is an Assistant Professor of Information Systems at the Singapore Management University. His research interests include stochastic analysis and optimization of complex systems, especially in urban transportation and logistics context, high dimensional data inference, and health-care analytics. His email address is haiwang@smu.edu.sg.

SATHYAVARATHAN SIVABALASINGAM is a Software Engineer at Fujitsu-SMU Urban Computing and Engineering Corporate Lab at the Singapore Management University, Singapore. His research interests include system design and integration. His email address is ssathya@smu.edu.sg.

KHIEM TRONG TRAN is a Software Engineer in Wego, Singapore. His email address is khiem@wego.com.