Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.J. Son, eds.

SIMULATION-BASED OPTIMIZATION TOOL FOR FIELD SERVICE PLANNING

Gabriel G. Castañé
Helmut Simonis
Kenneth N. Brown

Insight Centre for Data Analytics
Department of Computer Science
University College Cork
College Road
Cork, IRELAND

Yiqing Lin

United Technologies Research Center
411 Silver Lane
East Hartford, USA

Cemalettin Ozturk
Michele Garraffa

United Technologies Research Center Ireland
4th Floor, Penrose Business Center, Penrose Wharf
Cork City, Co. Cork, IRELAND

Mark Antunes

United Technologies Corporation
203 Sherbrook Street, Winnipeg, Manitoba,
CANADA

ABSTRACT

Many companies that deliver units to customer premises need to provide periodical maintenance and services on request by their field service technicians. A common challenge is to evaluate different design choices, related to staffing decisions, technician scheduling strategies, and technological improvements in order to make the system more efficient. This work provides a simulation-based optimization tool to support decision makers in tackling this challenging problem. The proposed framework relies on an optimization engine for the generation of the daily plans. A simulation component is used to evaluate the applicability of such plans by taking into account the stochastic factors. Furthermore, an interface manages the communication between these two components and allows a feedback loop between the simulator and the optimizer to achieve more robust plans. The applicability of the framework is demonstrated through a business case to evaluate different staffing decisions.

1 INTRODUCTION

In this paper, we consider the field service problem that involves dispatching service technicians to perform tasks on geographically distributed equipment such as copiers, telecommunication equipment, and heavy machinery. On-site service tasks can be routine preventive maintenance or repairs in response to equipment failures. Routine preventive maintenance tasks usually have due dates that are known well in advance. Therefore, those tasks can be scheduled and planned on a monthly and daily basis. On the other hand, repair requests are usually unexpected events. Some urgent requests may need to be addressed immediately. If a technician is responsible for both routine work and repair requests, then he may need to preempt a routine task for a repair request, causing an interruption of the schedule for routine work.

Field technician dispatching problems have been studied extensively in the past few decades (Dantzig and Ramser 1959) and are still relevant for companies and public bodies (Cheong et al. 2015; Kuo et al. 2014; Petrakis et al. 2012; Holland et al. 2017). Some common methods include optimization algorithms for vehicle routing problems and various heuristic approaches. Robust optimization approaches have also been applied to the problem taking into account stochastic variation in service time and travel time between customer sites. The goal is to generate a robust schedule that minimizes travel time and maximizes productivity while satisfying task due date requirements. However, those approaches usually do not consider unexpected repair requests. In the presence of those unexpected events, the pre-determined schedules are very often interrupted and cannot be completed as planned. The impacts of those unexpected events are difficult to evaluate using analytical approaches.

Discrete-event simulation (DES) has been widely used to model complex processes, especially in the presence of stochastic events, where analytical methods are not sufficient to provide accurate descriptions. Field service operations have been studied with DES methods in order to take into account stochastic events including variations in service time and travel time, as well as unexpected service requests. However, the DES models usually employ simplified technician dispatching decisions, most often heuristic approaches embedded in the simulation model, as it is non-trivial to adjust optimization solutions in response to unexpected events, such as urgent repair requests.

In this paper, we develop an integrated optimization and simulation framework that is capable of analyzing field service operations with advanced dispatching strategies, adapting to stochastic events. Field service tasks include both routine preventive maintenance tasks as well as unexpected repair requests. The main purpose of the work is to evaluate alternative staffing decisions as well as dispatching strategies before their actual implementation in the field. Compared with analytical approaches, the integrated framework provides a more accurate model for the dynamic environment of field service operations. On the other hand, compared with more commonly used simulation models with simplified dispatching decisions, our work describes a more intelligent field service operation with optimized dispatching decisions that are capable of adapting during the simulation runs to unexpected stochastic events.

The paper is organized as follows. Section 2 describes the motivation of the work and a literature survey. A detailed description of the integrated simulation-based optimization framework is provided in Section 3. Section 4 discusses use cases and experimental results. Section 5 provides some conclusions and future developments for this work.

2 MOTIVATION AND LITERATURE SEARCH

Based on the APICS (2010) definition, the term "field service" covers all functions of installing and maintaining a product for a customer. As these functions are mainly performed by field service technicians, planning and scheduling decisions related to field service operations fall into the human resource allocation and mobile workforce scheduling domain (Bouajaja and Dridi 2017).

Decision problems related to field service operations represent a hierarchical structure from strategic (i.e., grouping customer sites and defining service offices) to tactical (i.e., assigning portfolio of units to be maintained to technicians in each service office) and operational (i.e., daily dispatching of maintenance activities to each technician) level problems. While the strategic level decisions show characteristics of a facility layout and clustering problem (Antunes et al. 2018), tactical (planning) and operational (scheduling) problems are similar to vehicle routing problems with time windows. Readers can refer to Vössing (2017) for the taxonomy of field service planning problems in contrast to vehicle routing problems.

While strategic and tactical level field service problems can be cast as deterministic decision problems, uncertainty is an inevitable part of operational level decisions such as service times, travel times, occurrence of unplanned service requests, impact of traffic/weather conditions, and unavailability of technicians. Although it is theoretically possible to formulate all these factors as a closed form stochastic programming model, it is not practical for real industrial-size applications due to the very large search space, which leads to intractable models. Hence, most recent work from the stochastic programming community only consider the uncertainty in service times (Souyris et al. 2013). Meanwhile, simulation has been widely used for evaluating alternative scenarios of field service planning with a variety of sources of uncertainty (Olson and Overstreet 2015).

The most comprehensive model describing modules of a field service simulation and their interactions are depicted by Lin et al. (2002). According to their work, the main components are: (1) an equipment model used to simulate the usage of the service units, their conditions and failures (i.e., callback) over time, (2) a maintenance planner module used to generate preventive maintenance activities regarding to service and usage history, (3) a condition-based maintenance (CBM) planning module which monitors the health status of service units and generates a CBM task when needed, (4) a scheduling module that assigns and schedules generated tasks, and finally (5) a field service module that simulates the realization of generated tasks as well as technician activities such as traveling to service locations, lunch breaks and their availability. Lin et al. (2002) used this model to evaluate the impact of utilizing a CBM monitoring system on the number of machine breakdowns and showed its significant effect on increasing customer service level.

In contrast to Lin et al. (2002), the existing literature on simulating field service operations mostly focuses on evaluating staffing decisions. As an example, the pioneering paper of Dear and Sherif (2000) evaluates the impact of allowing technicians to roam between regions and of adopting several dispatching policies, in a field service system where only random callbacks are considered with non-stationary travel times. Similarly, Visser and Howes (2007) proposed a simulation analysis to evaluate the impact of the number of maintenance technicians on the completion of service requests. However, their approach does not take into account planned maintenance activities and lacks an optimization component for dispatching technicians. Alwadood and Rani (2010) also consider basic callback service requests, but in addition they model the skill sets of the maintenance staff and evaluate alternative dispatching decisions for these technicians.

Cortes et al. (2011) also covers the technician fleet size problem in only random-callback scenarios. However, their study is the first in the literature that performs a clustering study for grouping units and uses an external deterministic optimizer for dispatching service technicians. Other previous studies use heuristic dispatching rules internally in the simulation model.

Hertz et al. (2014) presented a unique simulation-based decision support system covering various aspects of the industrial field service planning problem, including grouping service units, locating inventory points for service spare parts as well as simulation of dispatching decisions. Although the developed decision support system provides the flexibility to investigate alternative designs in all these aspects, it lacks optimization-based tools for the strategic and tactical level decisions. Recently, Fuchs (2016) proposed a simulation tool to support evaluation of strategic level field service designs in terms of financial indicators.

In summary, there are several studies considering different aspects of the problem, but there is no research on combining optimization of strategic, tactical and operation level decisions in a holistic manner with evaluation via simulation. In this paper, a modular simulation-based optimization framework to evaluate the impact of various scenarios is proposed. The framework is based on the field service environment proposed by Lin et al. (2002). Finally, an industrial field service staffing problem is used to demonstrate applicability of the framework.

3 SIMULATION-BASED OPTIMIZATION FRAMEWORK

This section presents the framework used for scheduling a workforce using simulation and optimization to produce robust monthly schedules. Figure 1 shows the main components and the flow of data in the optimization and simulation framework. As the framework is not self-managed, an operator is shown responsible for adding and configuring some of the parameters and constraints that will be added to the optimizer to produce the schedules. Once the end user starts the process of generating the schedule, the *Optimizer* generates a set of configurations and schedules that will be given to the simulator through the communicator component. The latter is in charge of receiving the model, re-adapting it to any need of the simulation framework and triggering the simulation processes in the *Simulator*. Once the simulation processes end, the communicator receives the results including the list of finished tasks, the list of planned tasks that are not finished due to unexpected delays, and a list of newly generated tasks due to unit failures. This process is repeated iteratively for a given time period (hereafter assumed to be a month), after which a report is produced and given to the end user.

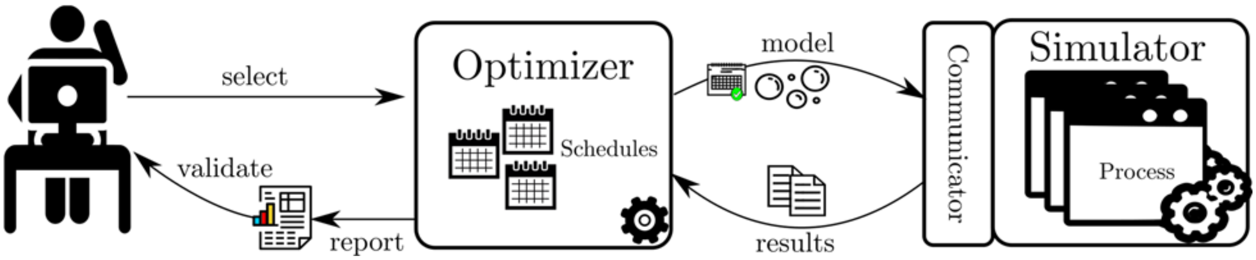

Figure 1: High-level overview of the framework.

3.1 Optimization Component Overview

In a complete workflow, an automated procedure should produce the detailed daily assignment of technicians to sites, considering the long-term workload while achieving the least-cost solutions. A decomposition mechanism, shown in Figure 2, is used as proposed in Antunes et al. (2018). Due to page limitations, readers are referred to Antunes et al. (2018) for details.

Figure 2: Overall problem decomposition.

The main components in the decomposition include:

Clustering Visits: As a first step, sites in close proximity are clustered together. It ensures that technicians visit these sites in the same trip where possible.

Route Generation: The second step is to partition sites into sets, which are called routes in the field service industry. The partition produces assignments with balanced workloads to technicians.

Monthly Schedule: In the next step, a monthly schedule is created for each technician, which takes into account all mandatory and optional tasks that should be performed in this time period.

Daily Rescheduling: During the day of operation, the pre-defined monthly schedule may be modified due to unexpected events.

3.2 Simulation Component

In Figure 3, a conceptual model describing the workflow for a field service technician is presented. The list of tasks for the technician consists of a set of planned maintenance tasks generated by the optimizer, and a set of unplanned tasks including call-backs and condition-based maintenance activities. All tasks are combined into one queue for each technician, and are ordered according to their priorities. When a technician is available, he/she first picks the activity at the front of the task queue. Then, the technician travels to the location of the selected activity and performs the required maintenance actions. Finally, he picks the next task in the queue.

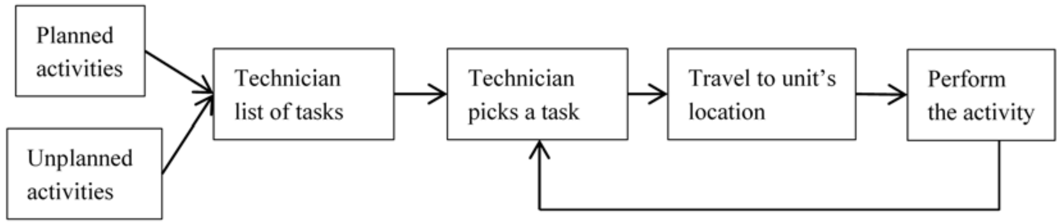

Figure 3: Conceptual simulation model of field service technicians.

This conceptual model has been implemented in AnyLogic (Borshchev, 2013), which allows a seamless integration with external applications and supports hybrid simulation methods. In the following subsections, AnyLogic implementation is presented along with the operational details.

3.2.1 Units, Callbacks and CBMs

Figure 4 depicts the behavior of a generic unit/component.

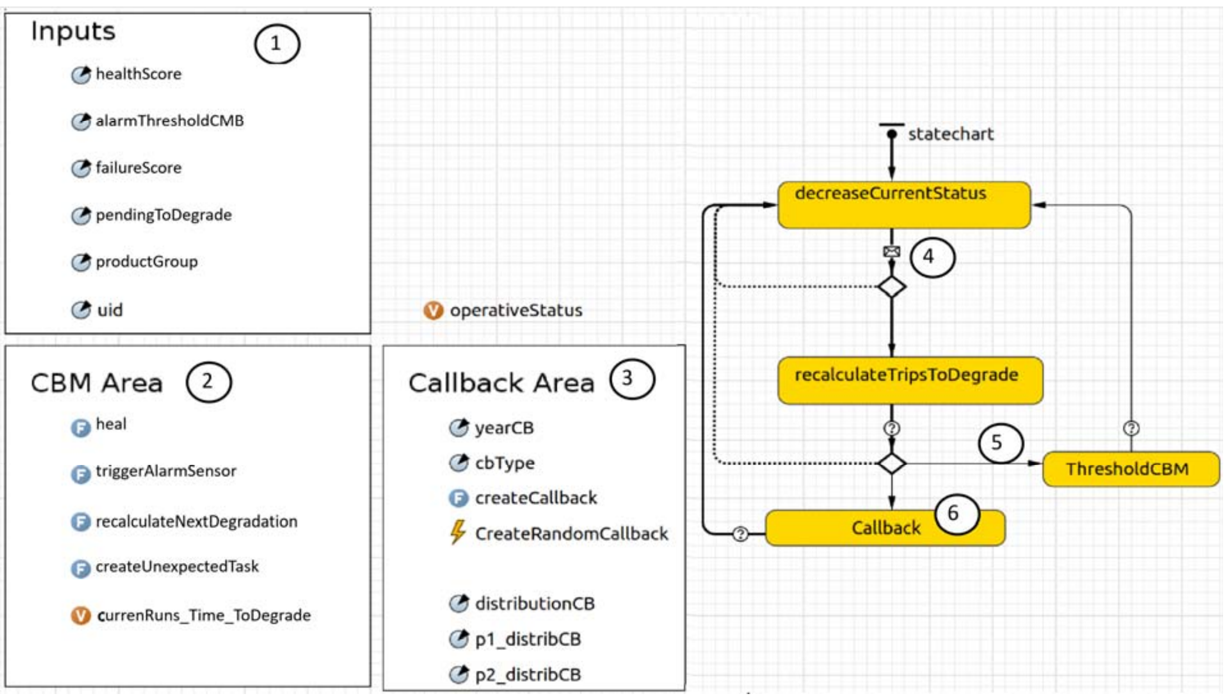

Figure 4: Agent-based implementation for units, their degradation, sensorization alarms, and the generation of unexpected tasks due to failures.

1. **Inputs** stores input data given by the optimizer, used to identify the units, components, and their features.
2. **CBM Area** models functions related to unit health including component degradation, health condition improvement after maintenance, and the creation of a failure resulting from an unattended CBM task.
3. **Callback Area** models random failures that are not correlated with unit health condition. The failures are generated from a distribution parametrized by the number of failures per year per callback type. The parameters are provided as input to the model.

The right part of the Figure shows a flowchart describing the agent-based simulation for the Units/Components. When a unit is in the *operative* state, its health score is gradually decreased due to its usage. If this score is reduced below the *ThresholdCBM,* then a CBM task is generated for the unit. If the CBM task is not performed by a technician, then the score continues to decrease until a callback is generated, which turns the unit into a *not operative* state. While the unit is in the *operative* status, the process continues from the *decreaseCurrentStatus* state. When a technician performs a maintenance task on the unit, the corresponding health score is restored.

3.2.2 Technicians and Task Executions

Technicians and tasks executions are composed of two parts: an agent-based modeling shown in Figure 5, and a process modeling shown in Figure 6. In the left part of Figure 5, a set of inputs are collected that describe start and end times, coordinates for the start and end locations, and whether the first and last trip should be taken into account as working time.

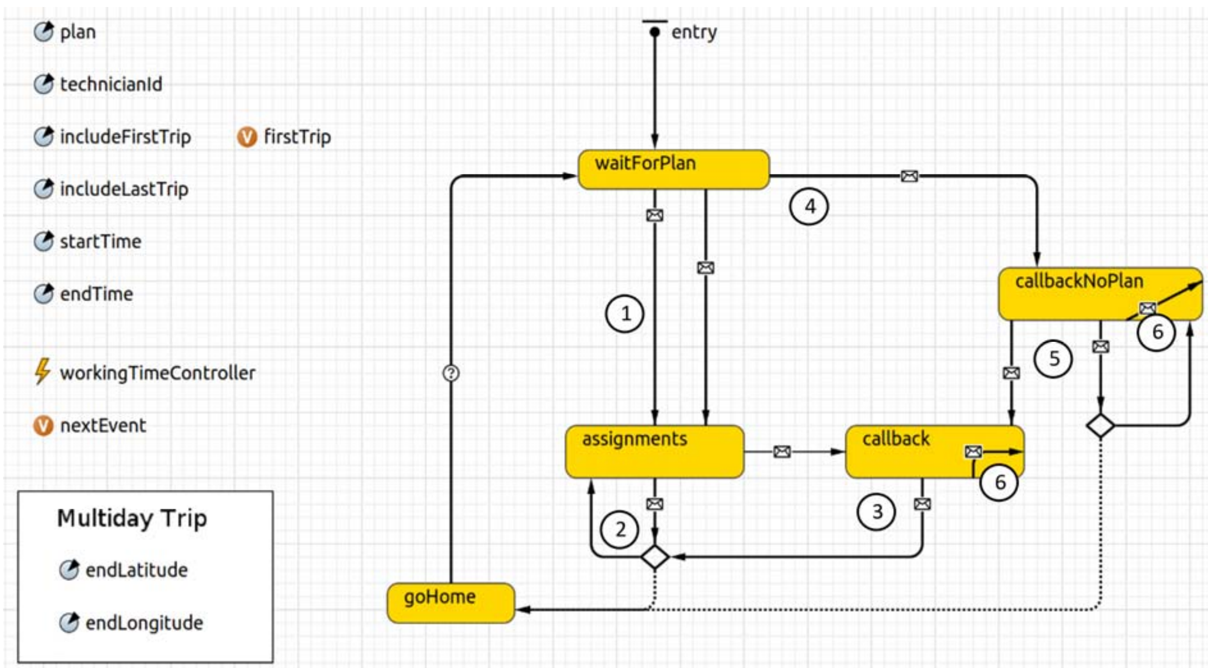

Figure 5: Agent-based implementation describing technicians flow to attend the tasks.

The right part of Figure 5 depicts an agent-based flowchart for the task lifecycle from a technician's perspective. Technicians start by transitioning to the **waitForPlan** state. When a new **DailyPlan** is received, it is assigned to the parameter **plan** and the agent moves (1) to the **assignment** state. Here each task is loaded in sequence into the process modeling, and executed according to the plan received (2) from the optimizer, until either there are no more tasks, and the flow moves to the **goHome** state and subsequently to **waitForPlan**, or *endTime* is reached. If a callback is allocated to the technician while in the **assignments** state, the task being executed is preempted and the emergency callback is attended to (3). Further, to model the behavior of technicians that work exclusively on unplanned tasks, a transition (4) from **waitForPlan** to **callbackNoPlan** is included. If during the execution of a callback a scheduled action is assigned to a technician (5), the agent moves back to the **callback** state. Finally, in case multiple callbacks are assigned to the same technician, these are preempted depending on their priority, and modelled as the nested transitions (6) in two states **callbackNoPlan** and **callback**.

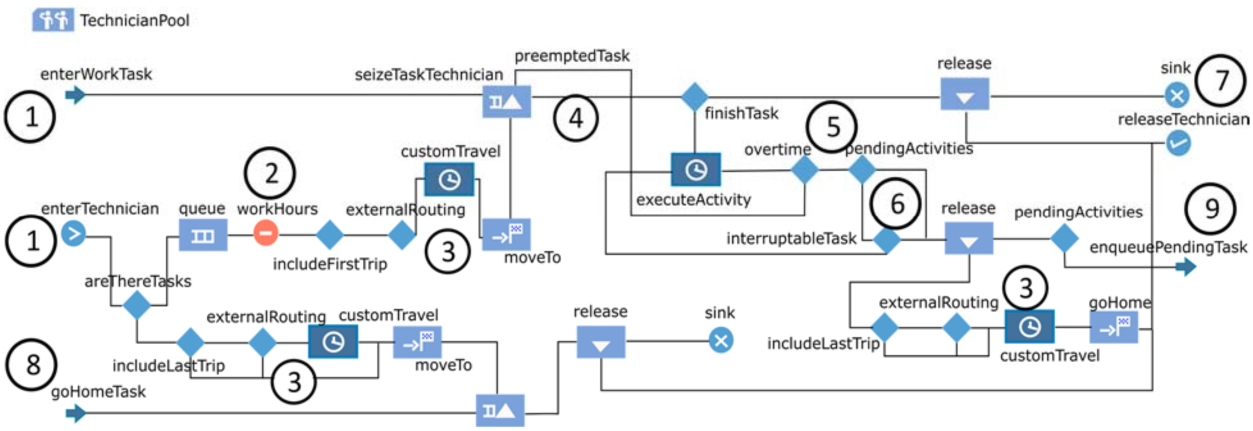

Figure 6: Process modelling implementation for tasks and technicians flow.

The process modelling in Figure 6 shows the behavioral aspects of a technician. The main aspects are:

1. Technicians of the same technician unit are the resources of the pool accessed through **enterTechnician**. Each task is independently loaded by the **enterWorkTask** input from agent-based **assignments** described in Figure 5.
2. The component **workHours** controls the number of hours that each technician works. This parameter is given by the optimizer and blocks travel for the technician if it cannot be completed between *startTime* and *endTime*.
3. Travel times between locations in AnyLogic are deterministic and simulated at constant speed. To add stochasticity to the travel times, to simulate delays due to traffic, adverse weather conditions or other events, the block **customTravel** enables future inclusion of these. We plan to incorporate weather and traffic historic data to analyse the impact of these factors on the overall schedule. This is currently planned as future research.
4. Once technician units are assigned to a task (seize block), a task can only be preempted if an unexpected task (*callback*) with higher priority is assigned to the same technician.
5. Activities are executed in the delay block **executeActivity.** Their duration is parametrized by type -- maintenance, repair or test -- and is specified as an interval of minimum to maximum duration in minutes. The skill level of the technician will then determine the duration within that interval. In the current model, these values are sampled from a uniform distribution. When a technician continues working beyond *endTime*, the additional time is counted as overtime. The **overtime** conditional block models this behavior, and also controls the execution flow redirecting the technician either to home or to execute the next activity.
6. For pending activities when the technician reaches *endTime*, the conditional blocks **pendingActivities** and **interruptableTask** decide if the technician can leave those activities to the next day, or if all activities within the task must be finished before ending the working day.
7. Technician and tasks are released at the same time.
8. goHomeTask is an dummy task used to model the time required for the technician to travel to their final destination (e.g. home, hotel or depot).
9. Finally, for tasks that can be interrupted, the pending activities are enqueued to be re-scheduled.

3.2.3 Communication Interfaces

As depicted in Figure 1, the optimizer developed for the workforce scheduling which generates the monthly plans, and the simulation model developed in the AnyLogic engine, are exchange information using a *communicator* component. This component receives the model description from the experiment to be simulated and, in order to mimic reality, weather predictions are incorporated for the day to be simulated – if available – to add a deviation to the transport time between locations.

For the communications with the optimizer, a REST server is used, allowing the optimizer and simulator to be decoupled and run on different servers, and allows the simulator to be called on demand by the optimization engine in a Simulation-as-a-Service fashion. JSON has been used as data standard for the communications.

4 EXPERIMENTS

This section describes the experiments conducted to test the framework. The stochastic components considered in the simulator are service time and callback occurrence. The service time is modeled by a uniform distribution with input lower and upper bounds. The callback occurrence is modeled by a Poisson distribution with an input parameter λ representing the average number of occurrences in a given time interval. Each unit has its unique λ value that is computed based on its historical data. In this set of experiments, only one type of callback is considered, and there are no CBM tasks. When a callback occurs, the assigned technician preempts his current task to work on the callback. Once the callback is resolved, he resumes his work on the preempted task.

The framework is executed in order to evaluate different design decisions. Each design i is evaluated with 50 independent replications. Each replication consists of one month of work, including 25 working days, and using the set of random streams $\{s_i\}_{i=1,\ldots,50}$.
Each replication is performed as follows:

1. Day 1 is simulated according to the monthly schedule.
2. When simulating the activities on a certain day, the tasks actually performed may be different from the tasks listed in the monthly schedule. This is because new callback tasks may be generated and some of the scheduled tasks may not be completed.
3. The daily re-scheduler computes a new schedule for the next day based on the actual completed activities of the current day in the simulator.
4. The next day is simulated and we jump to step 2 unless the end of the month is achieved.

After performing all the replications for different designs, we consider a paired-t test for testing the hypothesis of whether there is a statistically significant difference in the performances of the designs. The described computational setting is used to perform some real world experiments, in an attempt to use statistical evidence to provide answers to given business needs. Section 4.1 includes more details on the use cases. Section 4.2 describes and analyzes the results.

4.1 Description of Use Case

The framework is designed to evaluate the performance of the field service operations under different conditions. The final goal is to guide business process decisions with respect to some conflicting key performance indicators (KPIs). These KPIs belong to two different categories: (a) operational cost and (b) customer satisfaction. The first category includes travel time and idle time. The goal is to minimize these measures such that technicians have a higher productivity and operational costs are reduced. The second category includes the percentage of on-time task completion by task priority. In this context, the company's goal is to maintain a high customer satisfaction level while keeping the operational cost low.

The company is interested in analyzing the impact of different use cases including: (1) using overtime to finish tasks vs no overtime; (2) completing out-of-town work over multiple days vs driving back home each day; (3) impact of having a different number of technicians; (4) rescheduling policy after a callback.

These use cases explore different business process decisions related to technicians work rules, staffing and scheduling policy. This paper focuses only on use case 3. The aim is to provide decision makers with the analysis of the impact on KPIs of having different number of technicians.

The geographical area considered is a rural/urban area in North America, including thousands of units to be serviced.

4.2 Execution of the Framework

The developed framework is tested with a use case covering several regional offices and thousands of units. Due to page limitations, we will provide results on the impact of the number of technicians on several performance indicators such as percentage of tasks completed and total travel time only for one representative regional office. The baseline number of technicians is 8 and different levels of this design variable is varied from 9 (plus1) to 4 (minus4). Since each route is assigned to a dedicated technician, in the following figures a route refers to a single technician.

Since the number of callbacks received from units is independent of the number of technicians, we observed an increase in the average number of callbacks handled per route during the month as depicted in Figure 7. There are slight statistically significant differences between the different number of technicians, and Figure 7 shows practical differences between two groups of technicians as (9, 8 and 7) and (6, 5 and 4). Due to the increasing workload in designs with fewer technicians, the percentage of tasks completed is reduced with decreasing number of technicians, as shown in Figure 8. This results from more interruptions to the planned work due to the increasing number of callbacks per technician. The results observed in Figure 7 and Figure 8 complement each other. Note that Figure 8 shows no statistically significant difference between 9 and 8 technicians, and no practically significant difference between 7 and 6 technicians. We observe a similar result for the travel time per technician in Figure 9 where there is no statistical or practical difference in having fewer than 6 technicians. Since the optimization model (Antunes et al. 2018) tries to maximize the number of completed tasks, it assigns tasks in closer locations with shorter expected duration in cases with fewer technicians. This is the reason we do not observe significant difference in these designs. Similarly, total time per technician increases with the reduced number of technicians, since the workload increases as observed in Figure 10. It is interesting to note that in Figure 10, the variance in each design is quite low. It is observed that although the variance between individual days can be large, over the whole month time period the workload balances out fairly well.

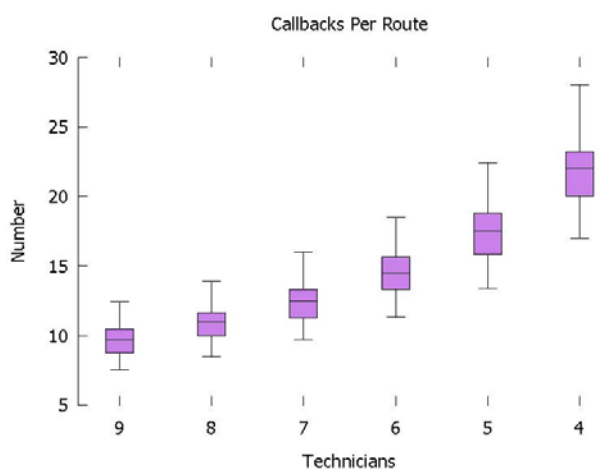

Figure 7: Callbacks per route (technician).

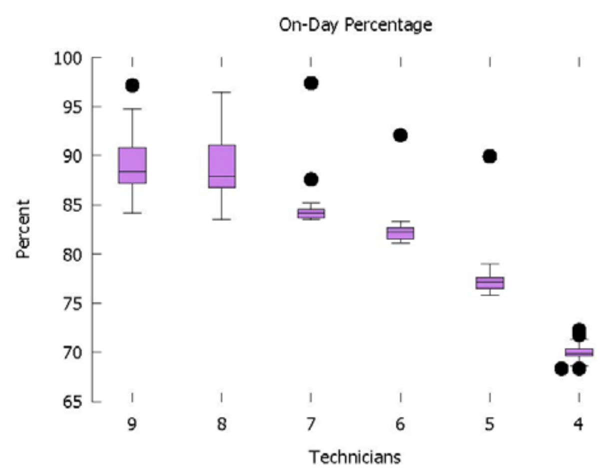

Figure 8: Percentage of tasks performed

Analysis of the difference between the performed (i.e., observed by the simulation) vs. actual work (i.e., planned by the optimizer) is presented for the baseline design in Figure 11 and for each of the 50 replications in Figure 12. The difference between these two indicators in each replication is due to the additional workload of unplanned callbacks as well as the stochastic factors considered in the simulation model (i.e., variation in service time). Furthermore, Figure 11 and Figure 12 show the independence of each replication within each design. Indeed, Figure 12 is a combined scatter plot representation of the same analysis. This figure gives a clear indication that observations fall apart for the minus 3 (i.e., 5 technicians) and minus 4 (i.e., 4 technicians) designs where additional travel due to callbacks significantly reduces the productivity on the scheduled work performed (y-axis).

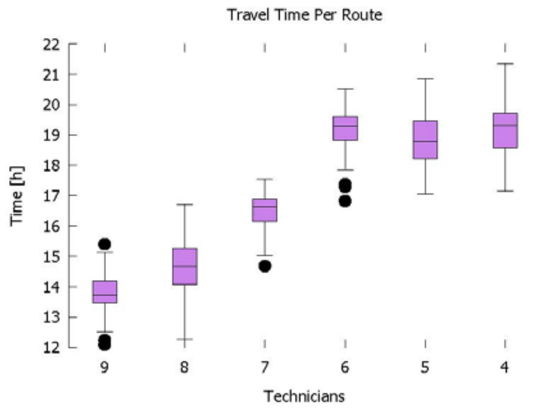

Figure 9: Travel time per technician.

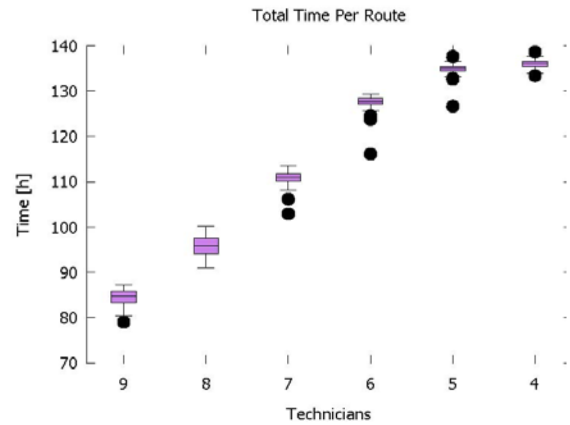

Figure 10: Total time per technician.

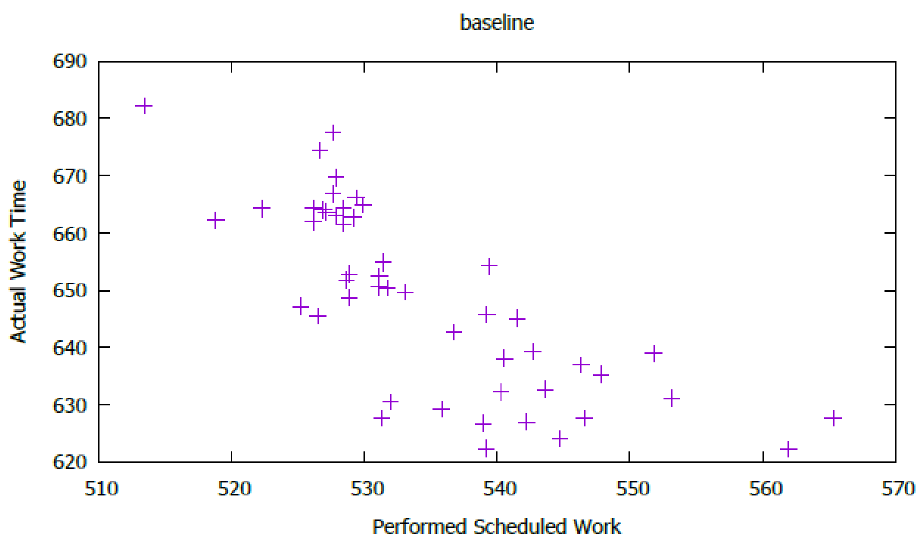

Figure 11: Performed scheduled vs. actual work for the baseline design.

Overall results indicate that the regional office considered in this analysis can provide the same level of service with 7 technicians. Reducing the number of technicians below this level will bring a statistically significant cost of compromising the number of completed tasks and travel time per technician.

5 CONCLUSIONS

This work presented a holistic framework for supporting business decisions in the context of field service with maintenance activities. The framework developed can be extended to different application domains due to its modularity (i.e., separate optimization and simulation components) and orchestration of these modules via the implemented interface.

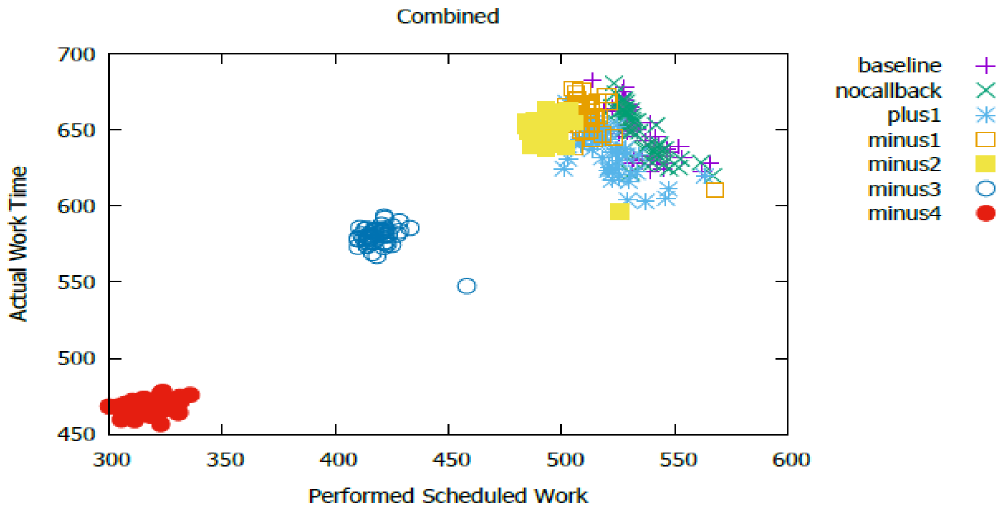

Figure 12: Performed scheduled vs. actual work over all replication

In this paper, as the primary purpose was to present the integrated framework and to demonstrate its effectiveness through a staffing use case, the detailed description of other use cases was omitted. The next steps for this study include enhancing the framework with real-time data stream from the field (Bader et al. 2017) as well as including recommender and incentive schemes from the technician's point of view. In addition, the integration of the proposed framework into existing field service management systems in the market will be investigated.

ACKNOWLEDGEMENTS

This material is based upon work supported by United Technologies Corporation under a UCC Collaboration Project and by Science Foundation Ireland under Grant No. 12/RC/2289 which is co-funded under the European Regional Development Fund.

REFERENCES

Alwadood, Z., I. Kassim, and R. M. Rani. 2010. "Maintenance Workforce Scheduling Using Arena Simulation". *Second International Conference on Computer Research and Development*, 7th–10th May, Kuala Lumpur, Malaysia, 517–521.

Antunes, M., A. Vincent, K. N. Brown, D. Desmond, G. Escamocher, A. M. George, D. Grimes, M. O'Keeffe, Y. Lin, B. O'Sullivan, C. Ozturk, L. Quesada, M. Siala, H. Simonis, and N. Wilson. 2018. "Assigning and Scheduling Service Visits in a Mixed Urban/Rural Setting". *2018 IEEE 30th International Conference on Tools with Artificial Intelligence (ICTAI),* 5th–7th November, Volos, Greece, 114–121.

APICS (Ed.). 2010. *APICS, Dictionary*. 13th ed. Chicago: The Association for Operations Management.

Bader, S. R., M. Vössing, C. Wolff, J. Walk, and M. Maleshkova. 2017. "Supporting the Dispatching Process for Maintenance Technicians in Industry 4.0". In *Tagungsband der 9. Konferenz Professionelles Wissensmanagement*, 5th–7th April, Aachen, Germany, 131–136.

Borshchev, A. 2013. *The Big Book of Simulation Modeling*. Oakbrook Terrace: AnyLogic North America.

Bouajaja, S. and N. Dridi. 2017. "A Survey on Human Resource Allocation Problem and Its Applications". *Operational Research* 17(2): 339–369.

Cheong, M. L., P. Koo, and B. C. Babu. 2015. "Ad-hoc Automated Teller Machine Failure Forecast and Field Service Optimization". In *2015 IEEE International Conference on Automation Science and Engineering (CASE),* 24th–28th August, Gothenburg, Sweden, 1427–1433.

Cortés, C., M. Gendreau, D. Leng, and A. Weintraub. 2011. "A Simulation-based Approach for Fleet Design in a Technician Dispatch Problem with Stochastic Demand". *Journal of the Operational Research Society* 62(8):1510–1523.

Dantzig, G. B. and J. H. Ramser. 1959. "The Truck Dispatching Problem". *Management Science* 6(1):80–91.

Dear R. G. and J. S. Sherif. "Using Simulation to Evaluate Resource Utilization Strategies". *Simulation* 74(2):75–83.

Fuchs, B. 2016. *Global Field Service Network Design: A Simulation-Based Decision Support System for Industrial SME's*. PhD Thesis. Department of Management, Technology, and Economics, ETH Zurich, Zurich, Switzerland. https://www.research-collection.ethz.ch/handle/20.500.11850/117597.

Hertz, P., S. Cavalieri, G.R. Finke, A. Duchi, and P. Schönsleben. 2014. "A Simulation-based Decision Support System for Industrial Field Service Network Planning". *Simulation* 90(1):69–84.

Holland, C., J. Levis, R. Nuggehalli, B. Santilli, and J. Winters. 2017. "UPS Optimizes Delivery Routes". *Interfaces* 47(1):8–23.

Kuo, Y.-H., J. M. Leung, and C. A. Yano. 2014. "Scheduling of Multi-Skilled Staff Across Multiple Locations". *Production and Operations Management* 23(4):626–644.

Lin, Y., A. Hsu, and R. Rajamani. 2002. "A Simulation Model for Field Service with Condition-based Maintenance". In *Proceedings of the 2002 Winter Simulation Conference*, edited by E. Yücesan, C. H. Chen, J. L. Snowdon, and J. M. Charnes, 1885–1890. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Olson, K. A., C. M. Overstreet. 2015. "Enhancing Understanding of Discrete Event Simulation Models Through Analysis". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 472–483. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc. .

Petrakis, I., C. Hass, and M. Bichler. 2012. "On the Impact of Real-Time Information on Field Service Scheduling". *Decision Support Systems* 53(2):282–293.

Souyris, S., C. E. Cortés, F. Ordóñez, and A. Weintraub. 2013. "A Robust Optimization Approach to Dispatching Technicians Under Stochastic Service Times". *Optimization Letter* 7(7):1549–1568.

Visser J. K. and G. Howes. 2007. "A Simulation Technique for Optimizing Maintenance Teams for a Service Company". *South African Journal of Industrial Engineering* 18(2):169–185.

Vössing, M., 2017. "Towards Managing Complexity and Uncertainty in Field Service Technician Planning". In 2017 *IEEE 19th Conference on Business Informatics (CBI),* 24th-26th July, Thessaloniki, Greece, 312–319.

AUTHOR BIOGRAPHIES

Gabriel G. Castañé is a Senior Researcher at the Insight Centre for Data Analytics in Cork. He holds a PhD in Computer science from University Carlos III of Madrid on the topic of energy aware cloud simulations. His email address is gabriel.gonzalezcastane@ucc.ie.

Helmut Simonis is a Senior Research Fellow at the Insight Centre of Data Analytics. His research interests are focused on developing decision support systems in different industries. His email address is helmut.simonis@insight-centre.org.

Kenneth N. Brown is a Professor at University College of Cork. His research areas are decision support and optimisation using constraint programming, distributed reasoning and other AI techniques. His email address is ken.brown@insight-centre.org.

Yiqing Lin is a Research Scientist at the United Technologies Research Center. She has extensive experience in the area of Operations Research including continuous and combinatorial optimization, and discrete-event simulation for supply chain management, manufacturing, as well as field service operations. Her email address is LinY@utrc.utc.com.

Cemalettin Ozturk is a Senior Research Scientist at the United Technologies Research Center and his expertise is in applications of combinatorial optimization and simulation methods in manufacturing, supply chain, service, and telecommunication networks. His email address is OzturkC@utrc.utc.com.

Michele Garraffa is a Senior Research Scientist at the United Technologies Research Centre. His research interest is on combinatorial optimization methods with a special focus on the development of hybrid exact/metaheuristic approaches for real-world industrial problems. His email address is GarrafM@utrc.utc.com.

Mark Antunes is Senior Regional Field Operations Manager at the United Technologies Corporation. His email address is mantunes@gmail.com.