# Analytical Study and Versatile Simulation of Tool Crib Operations in an Aircraft Maintenance Hangar

Edward Gingell and Michael A. Saliba*
Faculty of Engineering, University of Malta, Msida, Malta
Email: gingelledward@gmail.com, michael.saliba@um.edu.mt

***Abstract*—Large commercial aircraft maintenance facilities are characterized by the simultaneous layovers of multiple aircraft undergoing a wide variety of maintenance checks and modifications. The organizational layout of a maintenance hangar normally includes a tool dispensing store, or tool crib, from which aircraft technicians sign out special tools to carry out specific tasks. Efficient operation of the tool crib ensures that the non-value-adding time spent by technical personnel queuing to be served, and while being served, is kept as low as possible, therefore increasing productivity and minimizing the risk of delay. In this work, a detailed study of the operation of a tool crib in an aircraft maintenance facility is carried out. A detailed and versatile simulation of the operations is developed, based on real data collected at the facility with the aid of a new software tool to facilitate event synchronization. The empirical analysis is augmented by and compared to a theoretical study of the queuing behavior of the environment. The simulation is then used to investigate quantitatively the effects on the overall waiting time of varying the stock levels of critical tools, of reorganizing the tool storage locations, and of changing the number of human servers. In the context of the case study, it is found that relatively minor changes in tool crib organization can result in substantial savings both in the overall waiting time and in unsatisfied tool requests.**

***Index Terms*—aircraft maintenance, simulation, queuing theory, optimization**

## I. Introduction

The aircraft maintenance industry is one in which work precision and efficiency are two crucial characteristics. The Maintenance, Repair and Overhaul (MRO) industry for aircraft conducts a large number of checks to certify these aircraft as airworthy. This operation is extremely delicate, requires a large number of skilled workers and costs huge amounts of money to run.

The work reported here involves the detailed study and optimization of one of the activities within aircraft maintenance, namely the operation of the tool storehouse in a narrow-body aircraft hangar environment. The basic operation is to issue and receive back tools required on the shop floor within this hangar. The tool crib stores all the tools which may be needed during an aircraft maintenance check, and therefore any aircraft technician working on the shop floor needs to contact this tool crib directly to retrieve or return a tool at the beginning or end of some particular job.

The main objective function to be addressed in this project is the number of man-hours in queue away from value-added work, i.e. "wasted" by the production workers while waiting to be served by the tool crib personnel. Lowering the value of this function within the existing constraints (e.g. cost and available storage space) means that the tool crib is working in a more efficient manner and that the workers are not taken off production for a long time while requesting a tool to be issued or while returning a tool.

The approach taken in this work involves the collecting of data on the prevailing processes within the tool crib of an aircraft maintenance hangar; carrying out preliminary analysis on these data to identify potential areas of improvement; building a simulation model of the operation of the tool crib; and running the simulation on the real data while varying critical operating parameters as identified in the previous step, to quantify the effects on the objective function. The three main parameters within the model are (i) the number of servers, (ii) the stock level of each item, and (iii) the storage location of each item. The model can be run on different values of these parameters, while observing the effect on the average number of man-hours spent by the workers in queue.

The data collection was carried out within the Narrow-Body Hangar of an aircraft maintenance facility which specializes in C checks and comprehensive aircraft modifications. All the processes, discussions, parameters, simulations, results, and comments in this study relate directly to this environment.

## II. Theory and Literature

### *A. EASA Regulations*

The European Aircraft Safety Agency (EASA) is the European body which governs the aircraft maintenance industry and sets all the regulations required. Among others, the three most important sections are Part-66, Part-145 and Part-147 [1].

Manuscript received March 5, 2018; revised April 10, 2019.



doi: 10.18178/ijmerr.8.4.523-530

The EASA Part-66 regulations state what qualifications a person must have to receive an official license to work in aircraft maintenance, together with any limitations he/she might have. Part-145 is the section of the EASA Regulations which governs the places where aircraft maintenance is carried out, such as the site in the present case study. These rules specify a very wide range of conditions, from the need to have everything organized on the shop floor on clearly-recognizable racks, and the need for good housekeeping throughout the company, to the strict need of keeping the company in line with all the latest Health and Safety regulations. Part-147 sets the guidelines for an institution to be officially recognized as a learning centre for persons wanting to further their studies by registering for standard modules for the B License Type Courses, Fuel Tank Safety Courses, and Human Factors Training among others.

### B. MRO & Warehouse Management

Maintenance, Repair and Overhaul (MRO) is the term used to describe the operations of the companies in the aircraft maintenance industry [2]. This is mainly due to the various types of checks and maintenance routines they can carry out on an aircraft. For example, an A or C check only requires a number of inspections and some routine repairs, while a D check is commonly known as an overhaul since the maintenance program is very exhaustive.

An important phase in this whole maintenance procedure concerns the operations taking place within the store which keeps the necessary tools and equipment required for the workers to carry out their daily jobs. This place is a large space with organized storage spaces to segregate all the different tools and pieces of equipment, in order to make it as easy as possible to access them. It is well known in the literature that improving the placement of products in a warehouse can lead to gains in productivity in the picking operation, and studies have aimed to optimize the warehouse layouts and aisle designs both in existing warehouses and in the planning stages of new ones (e.g. [3]). This is done since inefficiency within this working site can have detrimental effects on the overall production efficiency. This is highlighted in the present work since every maintenance worker in the facility needs to access the same tool crib to issue or return a piece of equipment, irrespective of which of the (three) production lines he/she is working on. This can give rise to a production bottleneck if not handled correctly and thus limit the efficiency of the rest of the line.

The main objectives of a warehouse are to: [4]

- maximize the use of space;
- maximize the use of equipment;
- maximize the use of labour;
- maximize accessibility to all items; and to
- maximize protection to all items.

These points should be put into practice in the best way possible to achieve an efficient working environment and form the general basis of the present work.

### C. Queuing Theory

In general, waiting time is a variable to be reduced or minimized, and systems normally strive to make sure that the services required are completed within the shortest possible time. This applies to humans waiting in line to receive a service and means that customers in queues demand quick services with minimum waiting times. This challenge extends to a vast array of situations and environments. In a commercial setting, reducing waiting time has implications that go beyond a reduction in frustration, and can result in substantial financial gains through the corresponding increase in operational efficiency.

Queuing models can be applied to real-life situations by studying the real-life tasks thoroughly and making sure that any data collected reflect the particular situation in the most realistic way possible. These data will eventually help the observer decide which model to implement to carry out the analysis and to achieve the desired results. The potential of these models is in their ability to allow modification of specific variables, such as for example the number of servers, to find that value which gives the required maximum, or average, waiting time for customers.

Not every possible queuing system can be successfully modelled. There are numerous distributions which define the particular queue being studied, as are for instance classified according to Kendall’s notation (see [5]), but not all of them can be used for computational models. Case studies such as the one at hand require a good understanding of the existing system to make sure that the appropriate model is implemented. It should also be noted that the Poisson Process [6] [7], as modelled using, for example, an M/M/1 or M/M/c/K model, are quite idealistic and assume an exponential distribution of both arrival and service rates. Actual situations are not ideal and therefore will differ slightly from the M/M/c models, but before conducting experiments, one will firstly have to make sure that they are within the boundaries of the simulation and that they can thus return meaningful results. As Saaty states [8], queuing theory “*attempts to formulate, interpret, and predict for the purpose of better understanding of queues and for the sake of introducing remedies*”. Therefore, it is very important that realistic data are fed to these models, so that the estimations calculated are applicable to the actual situation being investigated.

## III. Case Study and Problem Formulation

As stated above, the main target to be reached in this work was the investigation and improvement of a specific objective function within real and existing constraints. This was addressed by first acquiring a solid grasp of the nature of the job at hand, the environment in which all the tasks are executed, the workforce and work patterns, and other relevant facts related to this work site. This was all carried out during a residency spent by one of the authors inside this work environment, both inside and outside the tool crib and in the rest of the production areas. Detailed observations were noted down continuously, and reports were drawn up concerning both aircraft layovers and the functioning of the tool crib. A general report at the end summarized the general considerations and important points to be addressed.

The approach taken involved building a simulation model based on specifications that were as precise as possible, in order stay faithful to real-life operations. Therefore, one major consideration was the way in which the data and information could be transferred from the real world to the simulated model. This problem was approached by collecting extensive data concerning the actual requirements, parameters, and operations of the tool crib within its working environment, and by constructing a simulation based on these data. The identified critical parameters could then be modified within the simulation environment while running the real data sets through the program and observing and evaluating the outcomes of the simulations.

The general structure of the simulation program is shown in Fig. 1. The input parameter file shown in the figure consists of the tool crib parameters to be addressed. The input use data file consists of real use requirements data collected within the actual environment during the data collection stage of the research. The tool crib operations are simulated on the use data using the set parameters, and the critical output is the average time spent by personnel in the tool crib queue.

## IV. Data Collection

One factor in this project which proved to be pivotal early on was data collection and filtering. This was most evident where the data had to serve as an input to the simulation model. As seen in Fig. 1, the use data input file required a certain structure comprising the three characteristics that described the relevant inputs of the outside environment to the tool crib: the worker identifications; their corresponding times of entry into the system; and the tool(s) they needed issued or needed to return. A sample of these inputs was collected during actual tool crib operations under the existing layout.

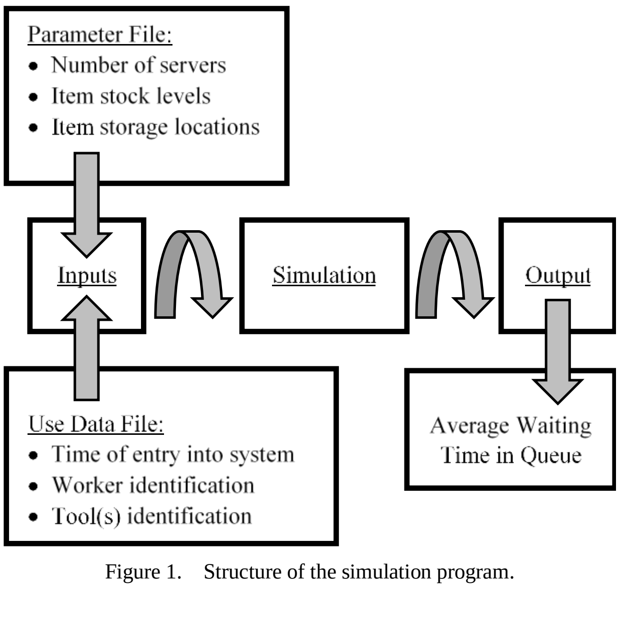


Figure 1. Structure of the simulation program.

This set of data needs to replicate the real-life situation in as faithful a way as possible. It was very important to define clearly the three use characteristics, to formally get the source of information from where these pieces of data should be extracted or collected, and how they would finally be put together to make sense in the data input file. "Entering the system" was defined in a general sense as the moment when a production technician shows up at the tool crib to be served. It is either the time that he/she arrives at the tool crib, finds an idle server, and starts to be served; or, in the case of a queue, it is the time that the technician joins the queue to wait to be served. From these inputs, the model can access internal data (also generated from data collection) to calculate the processing times required to serve the required technician. In this time, the server concerned is shown as busy and any new tasks coming into the system will have to be organised in a queue.

There were two sources of input use data: the first being the in-house electronic system that maintains records of the identity of the technician being served, the time that the technician started to be served by the server, the time that the technician's transaction with the server ended, and the tool(s) being issued or returned during the transaction. All these data are stored in electronic spreadsheets which can be accessed easily by company personnel. For the purposes of this work, this was a convenient way of getting the list of transactions that took place sequentially inside the tool crib. Crucially missing from these data, however, is the time that each technician joined the queue and the cases where a requested tool was not available (since the latter case results in no tool being issued or returned). Thus a second source of data was required to fill in these gaps, and this was acquired through manual (albeit computer-aided) collecting of data as described below.

The manual data collection had to be set up in such a way so as to eventually be easy to synchronize with the electronic data, i.e. the technician who joined the queue at time $t_1$ would be the same one who started to be served by server $i$ at time $t_2$, collected (or returned) tool $x$, and finished the transaction with the server at time $t_3$. The variables $i$, $t_2$, $x$, and $t_3$ for every event are associated with each other through the electronic data spreadsheets. While collecting the complementary manual data, it was important to record $t_1$ every time a technician joined the queue, as well as at least one of variables $t_2$ and $t_3$ in order to have a common time stamp with the electronic data. Recording the server $i$ manually as well proved useful as a double-check that synchronicity was not lost when combining the two data sources. This proved to be quite complex to carry out using solely manual means, due to the multiple inputs that had to be tracked by the person collecting the data.

This issue was resolved by writing a "timekeeper" program in C# to present the data collector with a user-friendly interface which made use of comfortably placed clickable buttons whenever a request entered the system (either at the end of the queue or straight to an idle server) and whenever one of the servers started to service a new request. This user interface is shown in Fig. 2. Whenever a technician entered the tool crib, the data collector clicked on the appropriate server "Start" button (if the technician started to be served immediately), or on the "Add" button of the queue (if the technician joined the queue). In either case, the technician was assigned a provisional identity by the program that consists of his time of entry into the system. If there were persons in the queue, then pressing the appropriate server Start button would transfer the person next in line (still identified by his original arrival time) onto the server, while also recording the time that the person started to actually be served. The timekeeper displays for a non-empty queue and for a busy server are shown in Fig. 3. Whenever a tool crib server ended a transaction with a technician, the corresponding Stop button was pressed by the data collector. In the background, this program stored the data in a text file, where each transaction consists of the time when the request entered the system, the time it started being served, the time it finished being served, and whether the request was for an issue or return of goods. The data from the timekeeper program could later be correlated to the data on the tool crib records system through the common time stamps of start and end of service by the (appropriate) server, giving a full record for every event. The number of unsatisfied requests could also be extracted from the combined data, since each of these would appear as a transaction in the manually collected data without a corresponding time stamped entry in the electronic data.

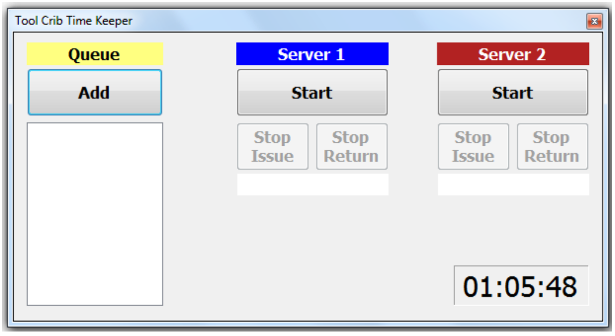


Figure 2. User interface of the timekeeper data collecting aid.

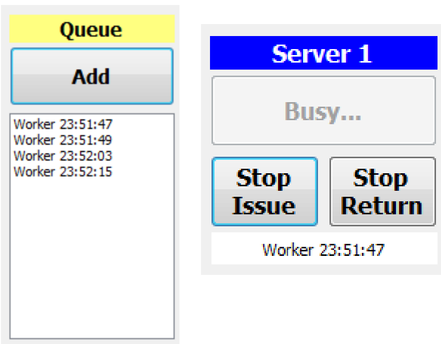


Figure 3. Timekeeper displays for a non-empty queue and a busy server.

This method was tried and tested on several occasions and after some initial problems, the desired result was achieved. The input data files could finally be set up properly with real inter-arrival times and real corresponding pieces of equipment. The final use data set was taken on a normal, busy working day with three aircraft layovers running simultaneously at different phases of completion over a four hour period, between 07:00 and 11:00.

## V. Analysis and Simulation

### A. *Overview*

The next phase of the project involved the detailed development of the simulation model and setting up the input parameters to give flexibility to the simulation. In parallel with this process, a separate software tool was also used to analyse the queuing behaviour of the environment being studied from a theoretical point of view. This latter task was performed to understand better the queue behaviour of the case study, and to obtain a partially independent value of the expected average waiting time which could later be compared with the output of the simulator.

The simulation model needed to be built up in such a way so as to cater for the three input parameters identified in section I and shown in Fig. 1. These three parameters needed to be analysed separately to get a good understanding of their individual effects on tool crib performance.

### B. *Queue Analysis*

The queuing analysis was addressed using a standard off-the-shelf software package: *Queuing Theory Simulator* [9]. This software package offers a variety of queuing models to test for, and for each model it offers the possibility to enter the basic characteristics of the queue. Based on these, it can calculate the probability density function and other important quantitative features to describe the queue behaviour.

The distributions of the inter-arrival times and processing times were extracted from the real data that had been collected. Some changes from the ideal queuing models specified by Kendell were evident. *EasyFit* software from *Mathwave* [10] was used to obtain these distributions. A compromise was found between the best-fit distribution offered by the software and the queuing models available on the queuing theory simulator. Finally, it was concluded that the situation approximates an M/E(k)/c queuing model rather than the ideal M/M/c/K model. The arrival times were exponential but the server processing times were ruled by an Erlang distribution of shape factor 3. The best fit server processing time distribution is shown in Fig. 4.

The values pertaining to the inter-arrival rate and service rate were calculated from the manually collected data and inputted into the system to give the following results:

- there was a general server utilization of 89%;
- the server utilization over the busy period during the four hours of data collection was 100%;
- the approximate system size is 6.

These early results gave an indication that there was a definite potential for improvement in the operations. This encouraged the next phase of building the model and working towards the optimization of the performance parameters.

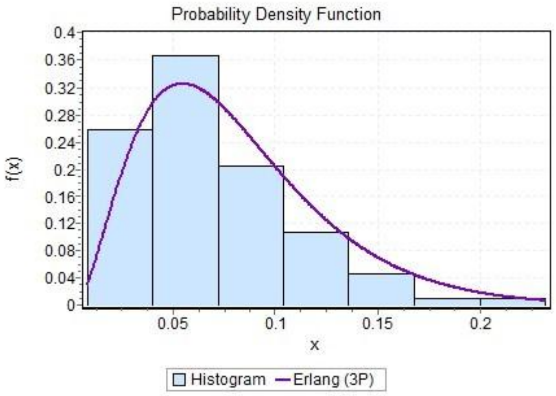


Figure 4. Best fit for server processing times: Erlang distribution with shape factor of 3.

## C. Equipment Location Analysis

This analysis was very important to get a clear picture of the effect of equipment location on tool crib performance and consequently to identify improvements that could be made. The way this was carried out was via a basic analysis of all the transactions of tools and equipment issued over a specific period of time. Firstly, the tool crib was systematically defined into nine areas to facilitate the analysis. The locations were defined according to the distance from the servers (Area 1 being the closest and Area 9 the furthest). The time phase chosen for collecting data for this preliminary analysis was an eleven-day period of normal business operations with varying workloads throughout the time span, and the data could be extracted from the electronic records of the company. All the tools requested by personnel over this time span were then analysed, and frequencies of tool demand per location area were calculated and drawn up. These are shown in Fig. 5.

These results gave an indication of the alterations that could be investigated using the simulation when it was fully functional. For example, it is seen that while Area 6 is relatively far from the server station, the frequency of tool demand from that area is considerably high. Thus, the relocation of a number of high demand items from Area 6 to a closer area would be expected to result in a decrease in average travelling time by the server, thereby reducing server processing time and increasing the overall system performance.

## D. Equipment Availability Analysis

The motive behind this analysis is to investigate whether the stock level of certain important pieces of tooling is sufficient. For this exercise, data were collected by preparing a form that would be filled manually by the servers to record all instances when they were unable to supply a technician with a requested tool, indicating the date, the tool, and giving a reason for why it was unavailable (i.e. “already issued to production”, “loaned to another hangar”, “undergoing calibration”, or “other”). This exercise was carried out over a period of 15 days. The results of this exercise helped to identify those items where an increase in stock level was most likely to improve tool crib performance and to later test this improvement quantitatively in the simulation. For example, through this exercise, it was specifically noted that three particular items accounted for more than 50% of unsatisfied requests. A histogram illustrating the percentage frequencies of the four reasons for unavailability of requested tools is shown in Fig. 6.

## E. Simulation of Operations

The full simulation model was developed in C# concurrently with the above analysis. It was programmed to use the approach shown earlier in Fig. 1, where it would receive a use data input file with activities, distribute these activities appropriately to the queue or to an idle server, and calculate all of the output variables associated with the operation of the system. These calculations involved the use of standardized durations for communication, travelling and data entry during server transactions, derived from the empirical data collected during the data collection stage of the work (and also using Monte Carlo random number generation where appropriate). During its evolution, the program would be displaying and recording statistics of queue lengths and average waiting times, with the aim being to minimize the latter as much as possible once all the three performance parameters were tested. Fig. 7 features a typical dynamic display of the simulation, showing the status of the servers; the length and contents of the queue; the cumulative average waiting, service and queuing times; and the cumulative number of unsatisfied requests.


Figure 5. Tool demand per location area.

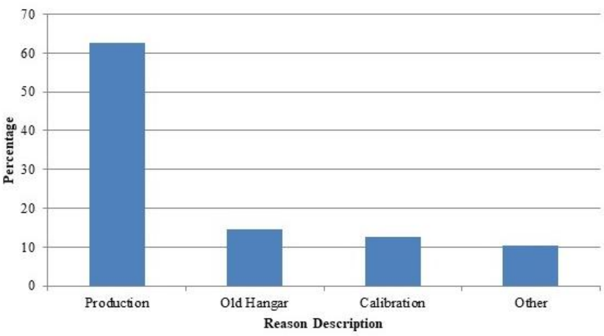


Figure 6. Reasons for unsatisfied tool requests.

## VI. Simulation Results

### A. Prevailing Situation

The first set of simulation runs were executed using input parameters that were equivalent to the prevailing situation in the tool crib, i.e. two human servers, and the tool stock levels and tool storage layout at the time. It was assumed that the servers were working 100% efficiently. The differences between the properties of these two servers are that (i) their distances to each and every storage location are different due to the different location of the server stations; and (ii) Server 1 is assumed to be more experienced than Server 2 (as is the real situation at the tool crib) and is also equipped with a bar code reader; while Server 2 is less experienced and is not equipped with a bar code reader. All these assumptions were reflected in the input files and in the unique settings files for the separate servers used in the simulator.

Ten different simulation runs were made using the same input use data file, with slightly different results being obtained due to the random number generation for communication, travelling and data entry. Each run generated an average waiting time (WT, i.e. the average total duration of a person in the system) and an average service time (ST, i.e. the average duration of the time spent with a server). The average of averages of the waiting time (AWT) and of the service time (AST) were found to be 485.6 s (8.1 minutes) and 58.2 s respectively. The average number of persons in the queue at any time (AQL: average queue length), was also extracted from the ten runs and was equal to 5.93. The resulting graph of queue length against time for this simulation is shown in Fig. 8(a).

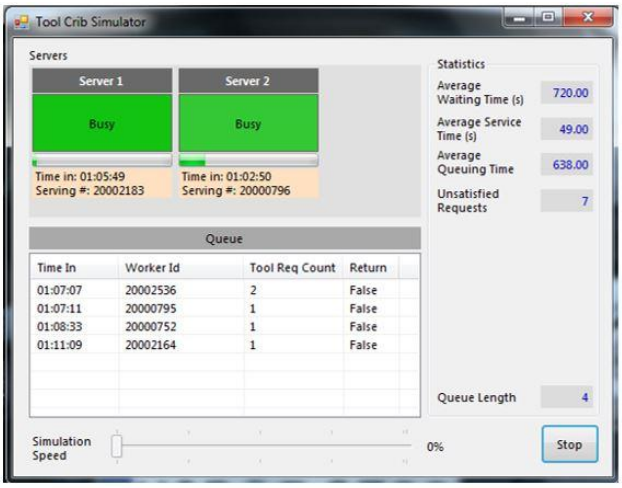


Figure 7. Dynamic display of the simulation program.

### B. Increasing the Tool Stock Levels

The second test involved increasing the stock levels of those items that had been identified as critical in the equipment availability analysis (section V. *D*). The main result from this test was that the number of unsatisfied requests decreased from ten to five. As expected, the waiting and queue parameters did not change significantly (from ten runs: AWT = 472.8 s (7.9 minutes); AST = 58.0 s; AQL = 5.85). The graph of queue length against time for this simulation set is shown in Fig. 8(b).

### C. Improving the Tool Storage Locations

The third test involved changing the tool storage locations according to the recommendations from the equipment location analysis (section V. *C*). It is noted that the changes that were made (i.e. simulated) adhered to the space constraints of the actual tool crib, i.e. all the changes made were physically feasible in the real system. The results from ten runs were AWT = 442.5 s (7.4 minutes); AST = 53.7 s; and AQL = 5.43. There was, therefore, a reduction of 9% in the average waiting time and a reduction of 8% in the average queue length. The graph of queue length against time for this simulation set is shown in Fig. 8(c).

### D. Increasing the Number of Servers

The fourth and final test involved the addition of a third server. The additional server was given the characteristics of Server 1 (an experienced server with a bar code reader) and ten simulation runs were executed as before. There were very large reductions in the average waiting time and in the average queue length, with AWT = 81.4 s (1.4 minutes) and AQL = 0.46. The graph of queue length against time for this simulation set is shown in Fig. 8(d).

## VII. Discussion and Conclusion

This work involves the development and use of a versatile simulation program to investigate the effects of several parameters on tool crib performance. The simulation was run on real data, and the consistent employment of the same input use data set enabled a direct assessment of the effects of system parameter changes on system performance. Separate offline analyses were carried out to better determine the specific inputs to be focused on in the simulation (e.g. which particular tools may need higher stock levels). The results confirm that increasing the stock levels of certain highly used tools will have a strong effect (reduction) on the number of unsatisfied demands. They also demonstrate that a simple reorganization of tool storage locations would be expected to result in significant reduction in the average waiting time, and that this reduction can be quantified. Finally, the results show that the addition of one server would result in a very significant improvement in system performance, with average waiting time reductions in the order of 80% and a reduction of peak queue lengths in the order of 75% (from 27 persons to 7 persons in the simulations that were run).

With respect to the alterations to storage locations, it is noted that the financial expense to transfer the identified items to a new storage location would be expected to be minimal, since the items do not require specialized or cumbersome storage facilities. Thus, the 9% reduction in average waiting time can be obtained with minimal cost and effort, and would be expected to result in significant annual financial gains. On the other hand, the alteration of stock levels would depend on the capital investment that the company would be prepared to allocate to this end. Prior to such an investment it would be prudent for the company to investigate the results summarized in Fig. 6 (reasons for tool unavailability), since it may be possible, for instance, that tools are not being returned promptly to the tool crib after use.

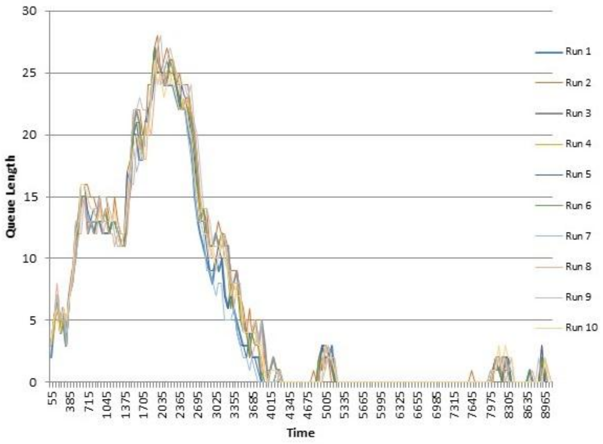


(a)

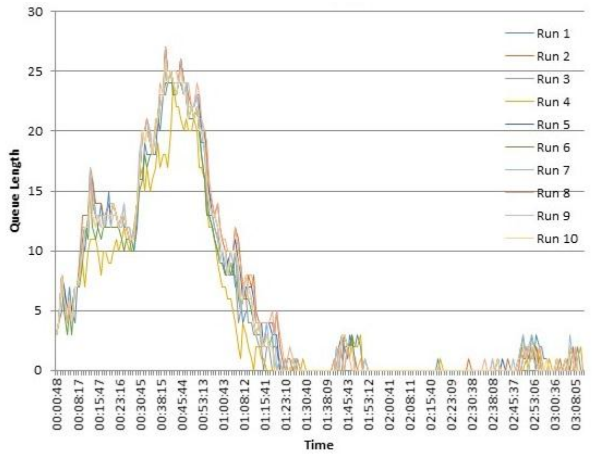


(b)

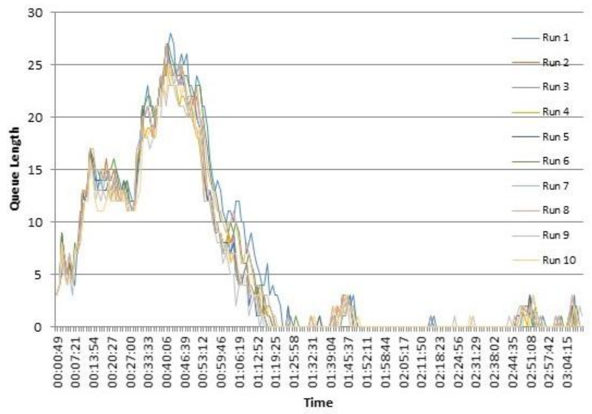


(c)

Figure 8. Queue length against time for the simulated time period: (a) for the current situation, (b) with increased tool stock levels, (c) with improved tool storage locations.

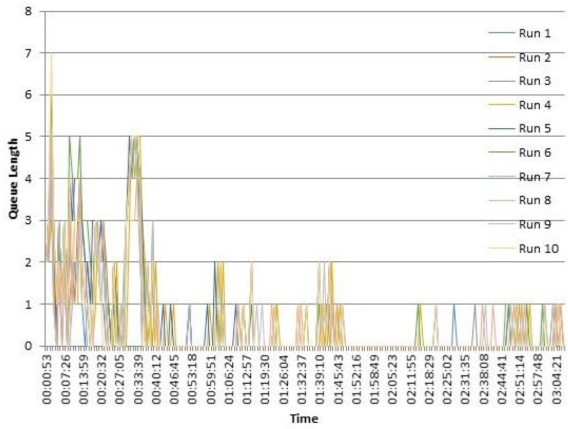


(d)

Figure 9. (continued) (d) with an additional server.

Finally, in regard to the additional server, it is noted that the use data input file employed in the simulation represented a very busy morning (somewhat of a worst-case scenario), and that therefore the tool crib may not be subjected to the same high workload every day. Consequently, the gains indicated by the simulation results might only apply for busy periods, and it may result that the third server would be redundant during lighter periods. It might be advantageous to find a balance by altering the shift patterns of the tool crib technicians to have three workers manning the tool crib during peak hours, such as the ones simulated in this work (for example, by calling in an additional tool crib server on overtime). Conversely, the number of servers could be scaled back to two (or even to one, as appropriate) at other times. In general, the company will need to carry out a detailed financial analysis to determine whether the savings in waiting time outweigh the cost of employing an extra server in the tool crib during busy periods. In future work, the simulation that has been developed can be used to analyse further data sets collected during periods of different work intensity, in order to obtain further results that can inform the company's decisions in this regard.

## REFERENCES

[1] European Aviation Safety Agency (EASA), Regulations https://www.easa.europa.eu/regulations [Accessed 28 February 2018].

[2] H. A. Kinnison, *Aviation Maintenance Management*, NY: McGraw-Hill, 2004, p. 35.

[3] U. Venkatadri and S. Kubasad, "Optimizing product placement for manual picking in warehouses," in *Proc. the 19th International conference on Flexible Automation and Intelligent Manufacturing* (FAIM), July 2009, pp. 581–589.

[4] The Warehouse Management Handbook, J. A. Tompkins and J. D. Smith Editors-in-Chief, McGraw-Hill, USA, 1998.

[5] H. C. Tijms, "Algorithmic analysis of queues," in *A First Course in Stochastic Models*; Chichester, U.K.: Wiley, 2003.

[6] I. Adan and J. Resing, *Queuing Theory*, Eindhoven University of Technology, The Netherlands, 2002.

[7] A. B. Clarke, "Maximum likelihood estimates in a simple queue," *Ann. Math. Statist.*, vol. 28, pp. 1036-1040, 1957.

[8] T. L. Saaty, *Elements of Queuing Theory: With Applications*, NY: McGraw-Hill, 1961, pp. 3-53.

[9] D. Gross, J. F. Shortle, J. M. Thompson, and C. M. Harris, *Fundamentals of Queuing Theory*, 4th ed. Hoboken, NY: Wiley, 2008. *QTSPlus* queueing theory software available at [Online]: ftp://ftp.wiley.com/public/sci_tech_med/queueing_theory/ [Accessed 3 March 2018].

[10] Mathwave, data analysis and simulation, *Easyfit* software [Online], http://www.mathwave.com/products/easyfit_desc.html [Accessed 3 March 2018].

**Edward Gingell** was born in Pieta, Malta in 1988. He obtained a B.Eng. (Hons) degree in mechanical engineering from the University of Malta, Msida, in 2010, and an M.Sc. in project management from the University of Northumbria at Newcastle in 2014.

He has been working in the manufacturing industry since graduating, and is currently a Launch and Change Manager with Methode Electronics Malta Ltd. Prior to this position he was working as a Qualification Engineer within the pharmaceutical industry at Actavis Ltd. in Malta.

Ing. Gingell is a member of the APM, PRINCE2, CIBSE and ASHRAE. He is a licensed Engineer in Malta.

**Michael A. Saliba** was born in Paola, Malta in 1963. He obtained a B.Mech.Eng.(Hons) degree from the University of Malta, Msida, Malta in 1986, a M.A.Sc. in mechanical engineering from the University of British Columbia, Vancouver, Canada in 1991, and a Ph.D. in physics, also from UBC, in 1998.

He is an Associate Professor in the Department of Mechanical Engineering at the University of Malta. Prior to his postgraduate studies in Canada, he was a Planning and Development Engineer with Air Malta Co. Ltd. in Luqa, Malta. His current research interests include service robotics and reconfigurable manufacturing.

Prof. Saliba is a Senior Member of the IEEE, a Member of the ASME, SME, and APS, and Corporate Member of the Maltese Chamber of Engineers. He is a licensed Engineer in Malta.