SIMULATION APPROACH FOR AIRCRAFT SPARE ENGINES & ENGINE PARTS PLANNING

José A. Ramírez-Hernández[(a)], Steven Oakley [(b)], Mei Zhang[(c)], Alejandro Scalise[(d)]

[(a), (b), (c), (d)] Operations Research & Advanced Analytics, American Airlines, Inc.

[(a)]Jose.Ramirez-Hernandez@aa.com, [(b)]Steve.Oakley@aa.com, [(c)]Mei.Zhang@aa.com, [(d)]Alejandro.Scalise@aa.com

ABSTRACT

We present the application of a simulation approach utilized for planning the required ownership levels of aircraft spare engines and parts in American Airlines. Such planning is essential to efficiently support both flying and engine maintenance operations. As such, this problem is very important from a financial and operational perspective given the high cost and criticality of these assets. Our models can be utilized in single and multi-location settings and are utilized to estimate the minimum ownership requirements to satisfy a given service level. In addition, our models can provide other important operational information such as out of service related metrics. To illustrate the utilization and versatility of the models, three case studies with actual industry data are presented. Results from these case studies demonstrate the value of the information generated by our models which facilitates the ownership planning and can also be used to support other related decision processes.

Keywords: supply chain, inventory planning, simulation, spare engine and parts planning

1. INTRODUCTION

Airlines own or lease aircraft spare engines to cover the operation while engines are overhauled or repaired. Due to the high cost of these assets, accurate planning is necessary not only to support the airline's flying operations, but also to avoid unnecessary excess in engine ownership that can be expensive for the company. Moreover, accurate planning of the required levels of engine spare parts is required to support the engine overhaul process, which is essential to provide on-time engine repair completions and avoid a negative impact in the spare engines availability.

In this paper, we present the application of a simulation approach to determine the required ownership of both engine spares and engine spare parts. Due to the particular complexity of the maintenance programs conducted and the uncertainty of engine removal and repair processes, a simulation approach is utilized to determine the required spare engine and the engine spare parts inventory levels. Our team has also derived closed-form formulations that can be utilized for fast ownership calculations under specific operational conditions. The simulation approach is preferred when a more detailed modeling of the process is required.

The problem addressed by this application belongs to the field of repairable inventory systems (Guide and Srivastava 1997; Tysseland and Halskau 2007). This problem differs from the classical problem of consumable inventories in the sense that the items in inventory can be repaired. The main reason for the repair operation is the high cost of the items in the system. Usually, such items are cheaper to repair than to replace for new ones. Thus, these systems are characterized by a flow of items being consumed but that later are removed and sent back to repair shops for reconditioning before returning to the system. As mentioned in (Guide and Srivastava 1997), these systems can be modeled as either single or multi-echelon systems depending on having single or multiple inventory locations.

In the context of the application presented here, the problem to solve can be simply described by the question: How many spare engines and engine spare parts does the company need to own in order to support the completion of the promised flying schedule? Planning for spare engines and its parts is an important problem for the company not only from the financial point of view but also from the operational perspective. On the one hand, engine spares and its repairable components or assemblies are expensive. The cost of a commercial airline aircraft engines usually ranges in the millions of dollars. For example, a CFM56 engine used in Boeing 737 aircraft can cost more than $10 million apiece (CFM International 2013). In the same way, some engine parts can reach the hundreds of thousands of dollars per unit. Thus, from the financial point of view, the ability to avoid an incremental spare engine/part purchase can provide significant benefits. Moreover, among all the maintenance, repair and overhaul (MRO) processes conducted in the airline industry, the engine repair is one of the most expensive operations (Ackert 2011).

On the other hand, accurate calculations of the required ownership of spare engines/parts are required to support maintenance and flight operations when such type of assets are repaired periodically. For instance, during the engine repair process, the company is required to have enough spare parts that will be used to replace the components being repaired. In some cases, the parts are repaired before or by the time these are required, but in many cases the repair times are longer than the time when the parts are needed to build an engine. In those cases, having the right ownership of spare parts is essential to provide an on-time delivery of the engine. If there are not enough engine spares parts available, then the risk of not having available spare engines is increased. Moreover, acquiring more spare engines or spare parts than necessary will help the availability of the equipment but it could also have an unnecessary negative financial impact. Thus, the approach presented in this paper aims to facilitate the estimation of spare engines and parts ownership needed to efficiently support both repair and flying operations while avoiding the excess in spare engine/parts inventory.

There is an important body of research in the field of repairable inventory systems for both single and multi-echelon cases, such as it is discussed in (Guide and Srivastava 1997; Tysseland and Halskau 2007). Although, to best of our knowledge, there are no specific papers describing simulation applications for both engine spares and engine parts in the context of commercial aviation, research has been conducted in the field of aircraft components in general. Interestingly, research in the field of repairable inventory systems is rooted on military applications for aircraft components, such as it is described by the seminal paper by (Sherbrook 1968) which focuses on the mathematical formulation to optimally stocking repairable parts for aircraft in a multi-location setting. The formulation; however, requires to adhere to specific mathematical assumptions. A more recent paper in the context of aircraft components is presented by Simao and Powell (2008), which provides the application of an approximate dynamic programming approach for a multi-echelon problem, where a combination of simulation and optimization techniques are utilized to determining the inventory levels of aircraft components at different locations in the system while maximizing the financial benefits from such decisions. This approach is applied at the higher level considering warehousing, locations, suppliers, but without entering in specific details of the repair process of the assets.

The application presented in this paper considers both a single and two-echelon repairable inventory system, referred hereafter as single-location and multi-location models, respectively. In contrast with previous work, our application is focused on the repair processes of engines and engine parts in the context of commercial aviation, and considers specific modeling details proper of the repair processes of these types of assets. Such features include borrowing or cannibalization of parts, scrapping processes, modeling of capacity constraints for the repair shops in the system, and work conducted pre and post the introduction of engines into the repair shops. In addition, our models consider so-called harvesting processes in which available stored or retired engines with remaining useful life can be advantageously reconditioned at smaller costs and with shorter turn-times than other more expensive repair programs. By using simulation, our application also allows increased flexibility to model complex details proper of engine repair processes as well as additional flexibility to use any probabilistic distribution that could be found to properly fit, e.g., repair times, demand patterns, transportation times. Moreover, our models provide the possibility to assess the impact that the planning of engine parts or "shop pool" can have in the spare engine ownership, and also provide information of other important performance metrics such as Out of Service (OTS) events due to unavailability of spare engines. Currently, our models are being utilized by the company for decision support in the ownership planning process of four different types of engines.

In terms of previous practices in the company, many of the procedures employed in the past to plan for the ownership of spare engines and parts involved the use of manual calculations. These were prone to errors and did not consider the variability of the repair processes. Thus, our application also represents a step forward towards a more accurate and reliable approach to estimate the ownership requirements to support the company's engine maintenance operations.

It is important to mention that all the procedures, simulation, and algorithmic details of this application are part of an original invention property of American Airlines, Inc., which is patent pending according to the United States patent law, and based on corresponding provisional filing conducted before the United States Patent and Trademark Office.

The organization of this paper is as follows: section 2 provides an overview of the process modeling and simulation. A description of the implementation of the models and corresponding calculation tools is given in section 3. In section 4 we present three case studies to illustrate how our simulation application has been utilized in different analysis and for decision support purposes. Finally, in section 5 we provide some conclusions.

2. PROCESS MODELING & SIMULATION

In this section we provide details of the process modeling and simulation implementation. The application considers single-location models for both spare engines and engine spare parts. Also, a multi-location model is utilized for spare engines when there are several stations with an allocation of spare engines to support the flying operations. An overview of these models is presented next.

2.1. Single-Location Model for Spare Engines

2.1.1. Simulation Process Overview

Figure 1 below illustrates the flow diagram followed by the engine repair process for the single-location case which utilized in the simulation. As depicted in the figure, the process starts with the arrival of engines for repair, where each arriving engine corresponds to a removal from an aircraft due to either failure or a planned maintenance procedure. Actual data indicates that Bernoulli or Poisson process provide good approximation of the arrival or removal process. However, the model also allows the use of other arriving processes if necessary.

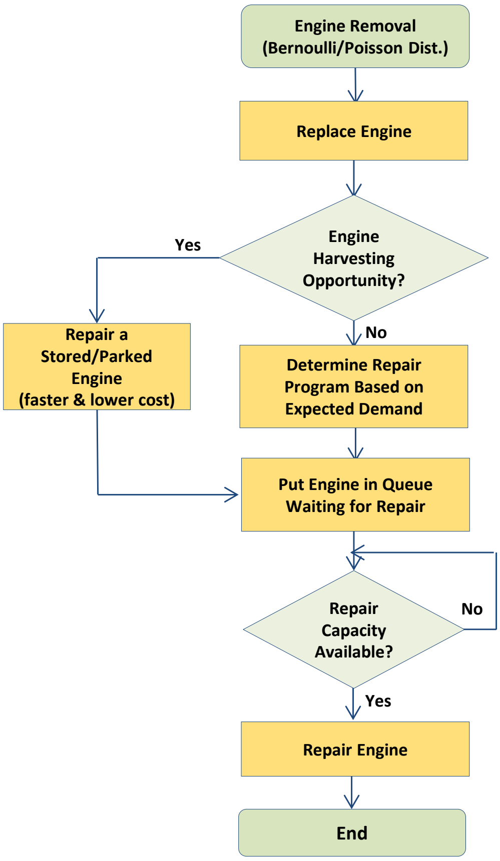

Figure 1: Flow Diagram for the Engine Repair Simulation Process, Single-Location Case.

Once an engine is removed and arrives for repair, the spare ownership level is updated given that each engine removed is replaced by an available spare. Then, a decision is made to determine if a harvesting process can be conducted. Harvesting is the process in which the company can defer the removed engine from being repaired, and instead, advantageously reconditioning a stored or previously retired/parked engine with remaining useful life but with lower repair costs and shorter repair turn-around-time than the removed engine. If harvesting is not an option, then a regular repair program is assigned. Such repair programs can be of a light type in which case the time needed to complete the repair is usually short. Similarly, there are also more extensive repairs usually referred as “heavy” which are more expensive and require longer turn-times than light repairs. The expected number of repairs to be conducted under each repair program is obtained from probabilistic engine removal forecast models developed by the company.

As indicated in Figure 1, the repair process could also have capacity constraints that may limit the number of engines that can be repaired concurrently. Thus, engines that arrive for repair are initially included in a queue where these wait until capacity is available to introduce the engine into the repair shops. While in queue, the engine can undergo a series of paper work and other procedures in preparation for the actual repair. Once an engine is sent into the shops, there is a repair time required to complete the reconditioning which is dependent on the repair program. In our application, repair times are random and are modeled using probabilistic distributions. For instance, we tested normal and gamma distributions by fitting them to various types of engine repair data using the method of moments. The gamma distributions produced more accurate fits overall and we concluded they were appropriate for the repair time distributions. At the end of the process, after completing the repair, the engine becomes available to the system and the engine spare ownership level is properly updated to reflect the addition of a new spare.

Some of the basic parameters utilized for modeling the single-location engine spare model include demand and repair time distributions, desired service level (percent of successfully fulfilled spare requests), and repair shop capacity constraints, e.g., maximum number of engines allowed under repair. Most of these parameters are obtained from historical data available from the company’s information systems and forecasting models. Moreover, there are specific parameters for the simulation runs that need to be defined: number of replications, simulation length, and warm-up period. In the case of this model, a total of 100 replications of 20 years each have been utilized with a warm-up period of one year. The warm-up period was determined by examining simulation outputs (replications) and determining the time required to reach the steady-state condition. Also, the number of replications and simulation length were selected after testing different combinations of these two parameters in order to achieve a desired accuracy in the estimated values. For instance, under the chosen values of simulation length and replications, we obtained a tight 95% confidence interval for the estimated ownership, with a lower and upper bounds that deviated from the average value by around 1%, but without sacrificing in simulation speed for practical purposes.

2.1.2. Model Output

The output generated by the simulation model corresponds to the variation of the engine spare level in time due to events such as engine removals and repair completions. Company proprietary statistical and optimization methods are then utilized to estimate the minimum required levels of total spare ownership such that a pre-specified service level is achieved.

In addition to the estimation of the total spare ownership requirements, the simulation model can be utilized to evaluate the performance of the system under the calculated ownership requirements. For instance, using specific ownership levels our models can be utilized to generate OTS-related metrics including the expected number of OTS events per year and statistics of the duration of such events. Similarly, the simulation output provides statistics of the WIP at the repair shops, average number of engines waiting in queue before being introduced for repair, and average spare count (ASC).

2.2. Single-Location Model for Engine Spare Parts: Shop Pool

In the case of engine spare parts, the modeling and simulation process follows a similar approach to that of the spare engines model. Again, the goal is to estimate the total required ownership of engine parts based on a pre-specified service level. However, in this case it is necessary to consider additional details proper of the engine parts repair process also called “piece-part repair” (PPR) process.

Figure 2 below depicts a general version of the simulation process for the repair of engine parts. Again, the process starts with the removal and arrival of an engine for repair. After arrival, a repair program is assigned and the engine parts are separated from the main engine assemblies. It is important to indicate, that an engine can have multiple parts of the same type, e.g., blades, vanes. Thus, at the time of separating the parts from the engine assemblies, there are components for which multiple units of the same type are sent for repair.

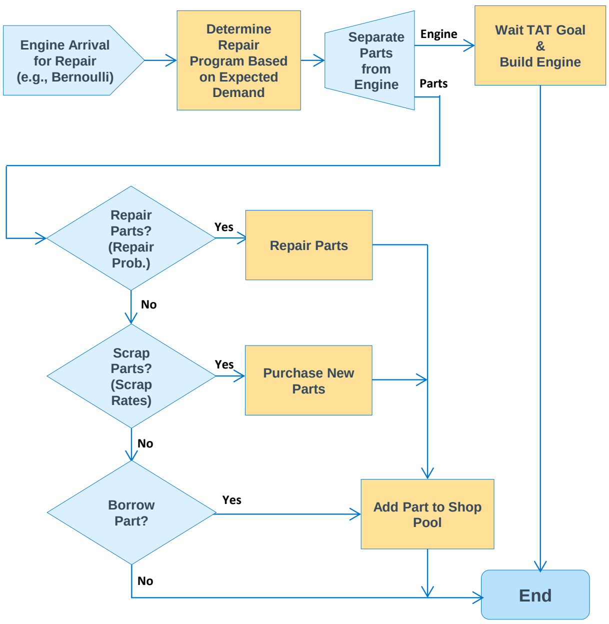

Figure 2: Flow Diagram for the Engine Spare Parts Repair Simulation process, Single-Location Case.

Once separated from the engine, the parts are evaluated to decide if repair is necessary, or there is need for scrapping the parts, or if the part is in good condition and can be borrowed to be used in the building of other engines downstream in the process. When the parts are sent for repair, the delay in the repair process is modeled using a probabilistic distribution, e.g., gamma. After completion of the repair process, the repaired parts are added to the available inventory or shop pool.

In the process there is also the possibility of having the parts scrapped. In that case, a purchasing order is generated and a random lead-time is used to model the delay after which the new purchased parts are back and available in the shop pool. The repair process of the parts and the engine is usually guided by turn-around-time (TAT) goals for the completion of the repair process. Thus, when parts are separated from the engine, the main assemblies wait for a specific amount of time or TAT goal before collecting the required spare parts from the shop pool and continue to complete the repair process. In some cases this goal is much shorter than the time required for completing the repair of parts. Thus, specific level of available inventory in the shop pool is necessary to support the repair process.

The parameters utilized by this model include, among others, demand and repair time distributions, repair probabilities, scrapping rates, capacity constraints, and desired service level (percent of successfully fulfilled spare part requests). As in the case of the spare engine model, most of the parameters are obtained from historical data available from the company’s information systems. Other parameters such as the number of expected engine removals (demand) are estimated using probabilistic forecasting models developed by the company. For practical purposes, the model has been designed to be run for individual engine parts. That is, repair time, repair probabilities, and scrapping rates, and other parameters are specified at the part level. Similarly, the model requires specifications for the simulation runs. For example, this model has been run using 500 replications of 15 years each with a warm-up period of one year which is appropriate to obtain enough simulation samples in steady state conditions. In this case; however, more simulation replications were utilized to obtain a tight 95% confidence interval for the estimated ownership (upper and lower bounds with around 1% deviation from average value). The reason for this is the increased variability of the simulation output given there are components with multiple units per engine.

2.3. Multi-Location Model for Spare Engines

The multi-location spare engine model can be seen as a two-echelon system with repairable components in which different locations or stations can have a spare allocation to support the operation. In addition to the stations, the system includes a location with engine repair shops and where engines are sent back and forth after removals and repair completions, respectively.

Figure 3 below illustrates the flow diagram of the simulation process for the multi-location model. As before, the process starts with the removal of engines due to failures or planned maintenance procedures. In this case, the removal process occurs at different stations. To model such removal processes, Bernoulli and Poisson distributions have been found to be a good fit based on actual data.

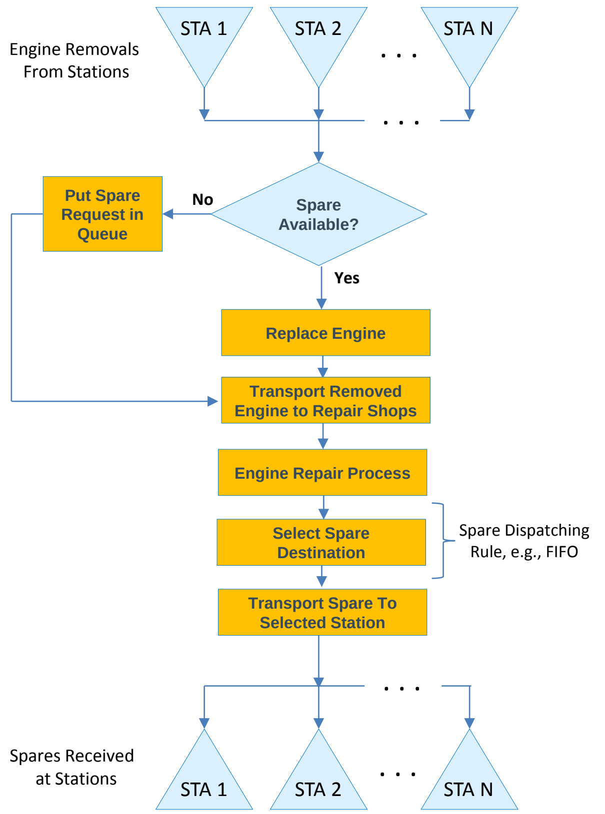

Figure 3: Flow Diagram for the Multi-Location Engine Spare Simulation Process.

As soon as a removal occurs at a station, then a new spare available from the shelf is utilized to replace the removed unit. However, if there are no available spares, then an OTS is generated and the spare request is put in queue until it is fulfilled by using a spare sent back from the repair shops, or in some cases, from another close by location that could allow the borrowing of a spare. All the removed engines are then transported to the repair shops and such transportation is modeled as a delay to arrive to the engine repair location. Once the engines arrived to the repair shops these follow the same engine repair process depicted in Figure 1. After completing the repair, a spare dispatching decision is conducted to select which station will receive the new spare. The selection is made based on the queue of outstanding spare requests from the different stations and by using specific dispatching rules. Some of the rules utilized include static type of rules such as First-In-First-Out (FIFO), and longest transportation time (LTT), as well as dynamic policies based on current inventory levels. After applying one of these rules and selecting the spare destination, then a transportation process is conducted before the engine arrives to the corresponding station.

In this model each of the stations in the system has a defined spare allocation, and each station can be measured in terms of service level and OTS-related metrics. Again, the output generated by the simulation is utilized to determine the total spare engine ownership per station that will meet a desired service level. Moreover, the output generated by the model allows measuring the performance at system-wide level, e.g., system-wide service level and expected number of OTS events per year.

The model utilizes different process parameters, including removal rates by station, desired service level by station, and transportation times between stations and the repair shops. Again, most of these parameters are obtained from actual data available from the company's information systems. Also, parameterization of the simulation runs is required. For instance, the model has been run using 50 replications of 30 years each with a warm-up period of one year.

3. IMPLEMENTATION OVERVIEW

The models described in the previous sections have been implemented within so-called calculation tools for the end-user. Figure 4 below depicts the general architecture followed in the implementation of these tools. As illustrated in the figure, the implementation includes both a user and an external server side. The user side has a Graphical User Interface (GUI) built in MS-Excel and using Visual Basic (VBA), for easy use and portability. The GUI is utilized to facilitate the configuration and specification of the parameters required by the model as well as to execute the calculations (run simulations). In the case of the engine spare model, the GUI offers the possibility of setting multiple scenarios, each with a different set of parameters, which allows the comparison of ownership requirements for different operational conditions.

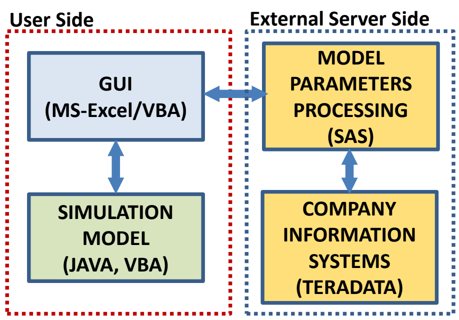

Figure 4: General Implementation Architecture for the Spare Engine and Spare Parts Simulation-based Calculation Tools.

Also within the user side of the implementation is the simulation model which uses the parameters specified by the user in the GUI. In this application the simulation models have been implemented using Java (as independent executable files) and VBA (integrated within MS-Excel). Prototyping of the models were also conducted in Arena (Kelton et al. 2006) before implementing the models in VBA and Java. In particular, Java allows fast simulation runs which is advantageously utilized for the engine spare parts model where the model is run for dozens of parts and speed is a factor. For developing purposes, for example in Java, in addition to the basic programming libraries (Oracle Corporation 2014) we also utilized the Commons Math library (The Apache Software Foundation 2014) for random number generation and statistical functions. Moreover, the discrete event simulation model was built utilizing a fixed time increment simulation clock with a 1 day interval. For each interval, the number of arrivals, processing times, repair completions and other events were simulated using the corresponding probabilistic distributions, and to update the state of the system. The approach taken to simulate the system was to build a non-terminating simulation that allowed us to study the steady state behavior. We utilized Java1.7 to code the model, including a calendar of events, queues counters, and the statistics engine.

On the external server side of the implementation, we have the company information systems, e.g., Teradata (Teradata Corporation 2014), which are used to extract historical information needed to obtain the parameters utilized by the models. The required processing of these data is conducted using specifically designed code in SAS (SAS Institute Inc. 2014). The resulting parameters after this processing are sent back to the GUI and are then made available for simulation and calculation purposes.

As part of the model development, a process of verification and validation (Sargen 2010) of the simulation models was conducted. First, the conceptual models were verified by the business units in charge of the process, e.g., engine repair production control and asset management teams, to verify that the assumptions and modeling details were correct. Then, a process of verification of the programming language code utilized to implement the models was conducted to ensure that the code was properly representing the conceptual model. In addition, verification of the models was also conducted by comparing the simulation output against closed-form formulas utilized to estimate the ownership under more relaxed conditions, e.g., infinite capacity. For validation purposes, historical data of spare requirements based on actual demand was utilized and compared with the estimations provided by our models. That is, under the same actual demand conditions, our models provided close calculations of the ownership requirements, with only small deviations from actual requirements.

Finally, and although a formal accreditation process has not been conducted, the models have been extensively and successfully evaluated by final users in the engine production control and asset management groups in American Airlines. Among the different groups involved in the development of these models, there has been consensus on the fact that simulation was the right approach to provide ownership estimations given the complexity of the process. The positive acceptance of implemented models in the business units has also facilitated the expansion on the application of our models to estimate ownership requirements for other type of aircraft assets such as Auxiliary Power Units (APU).

4. CASE STUDIES

The implemented models described in this application have been utilized for different analysis aimed to support the company's decision process regarding engine spares and parts planning. To illustrate the use of our models, we present next three case studies considering both single and the multi-location scenarios.

4.1.1. Impact of Engine TAT in Spare Ownership & Shop Pool Investment

The first case study is focused on the impact that the engine repair TAT has on both the engine spares ownership and the engine spare parts or shop pool investment. The shop pool investment is directly proportional to the required ownership, and in this case, it represents the aggregated amount from all parts considered in the calculations (more than 200). This case study was part of an analysis conducted with the objective of selecting a feasible and appropriate engine TAT goal for the engine repair shops and compatible with the company's flying operations objectives.

Figure 5 below shows some of the results obtained from this analysis. Please notice that the values of the scales in the chart of Figure 5 are omitted because these correspond to company proprietary information.

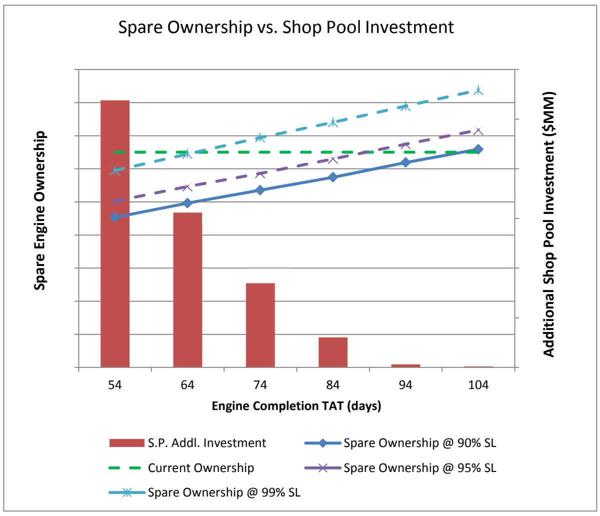

Figure 5: Impact of Engine Repair TAT in Spare Ownership and Shop Pool Additional Investment.

We conducted calculations of the spare engines and engine parts ownership for engine repair TAT's ranging from 54 to 104 days in average. As indicated in the chart, the additional investment in the shop pool (obtained at a 98% service level) is decreasing with the increment in the engine TAT. Conversely, as the engine TAT is increased, the engine spare ownership also increases. These results are expected given that increasing the engine TAT will allow additional time for completing the repair of engine parts, which leads to smaller shop pool requirements. However, increasing the engine TAT will also lead to increments in the engine spare ownership levels.

The chart in Figure 5 also shows the required spare ownership for different service levels (SL): 90%, 95%, and 99%. As indicated in the figure, higher service levels required higher engine spare ownerships. As a reference, the chart also includes the current ownership at the moment of the analysis (dashed line). For example, the results illustrated in the chart indicate that for the current ownership it is possible to obtain up to 99% service level when the engine repair TAT is about 64 days. However, such conditions also require an additional shop pool investment. Thus, the results generated by the models can provide valuable information to evaluate the trade-off between engine spare ownership, shop pool investment, and engine repair TAT.

4.1.2. Impact of Engine Spare Borrowing Between Stations on the Duration of OTS Events

The reduction or mitigation of OTS events is an important task in the planning and management of spare engines. Thus, using the simulation models presented here, the company is able to conduct analyses to estimate the impact that certain planning strategies may have in the occurrence of OTS events.

In this particular case study, the goal was to determine the impact that borrowing of engine spares between stations could have in the reduction of the average duration of OTS events. For this purpose, we considered the possibility of transporting available spare engines from a selected station to other stations in the system that could have outstanding spare requests. This task was conducted by setting a threshold point for the inventory of the station borrowing the spares. Thus, if the inventory level was equal or larger than such threshold, then it was allowed to transport a spare from the borrowing station to another station holding an outstanding spare request. Figure 6 below illustrates the results obtained across the different stations and by using borrowing threshold levels from 0 to 5 spares.

For the experiment conducted, the repair base and station was selected as the location from which spares can be borrowed. The reason for this was that such station not only had the largest spare inventory but it is also located at the same place as the repair location, which provides added flexibility in handling the inventory. Results from this case study indicate that decreasing the borrowing threshold has a positive impact in the duration of the OTS events, such as it is illustrated in the chart from Figure 6. That is, when more spares are allowed to be borrowed from the repair and base location, then the time that stations with outstanding spare requests have to wait is reduced. In particular, notice how stations 2, 3, and 4 have a significant reduction (more than 50%) in the duration of OTS events when borrowing is allowed compared to the case of no borrowing at all. Also notice that the repair base station was not negatively impacted in the duration of OTS events when borrowing of spares was allowed.

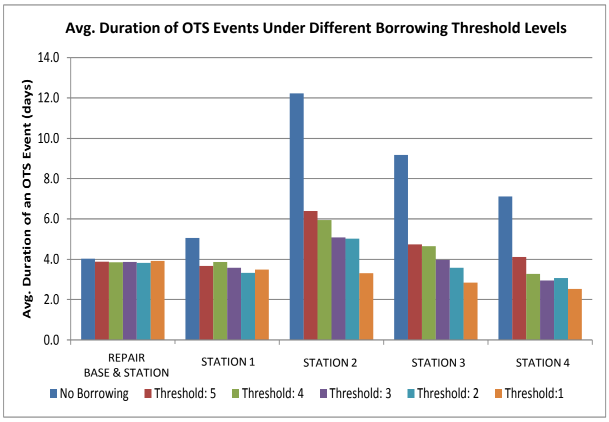

Figure 6: Impact of Engine Spare Borrowing Between Stations in the Duration of OTS Events.

In fact, there is also a slight reduction in the OTS duration observed at that station. Thus, in this case study our models are useful to demonstrate the positive impact that the borrowing of spares between stations has in reducing the average duration of OTS events.

4.1.3. Impact of Engine Harvesting Process on the Service Level

Another type of decisions that the implemented simulation model is able to support, are those of determining the level of engine harvesting required to achieve specific service levels in the system.

As mentioned earlier, the engine harvesting process consists in the reconditioning of stored or retired engines which still have useful working life, but that require less expensive repair procedures and shorter repair turn-around times than other more extensive repair programs. Thus, if the initial plan considers heavy type of repairs on several engines, but there is an opportunity to conduct a harvesting process, then the heavy repairs may be deferred and instead the harvesting candidate engines could be reconditioned and made available as spares.

In this case study we considered the impact that different levels of harvesting have in the service level observed across the different stations. The harvesting level is defined here as a percent of heavy repairs that are deferred to conduct harvesting of stored or retired engines instead. We tested harvesting levels from 0% to 30% of the planned heavy repairs while maintaining the total ownership constant. Figure 7 below illustrates the service level obtained at the stations under the different levels of harvesting. It also shows the impact in the system-wide service level.

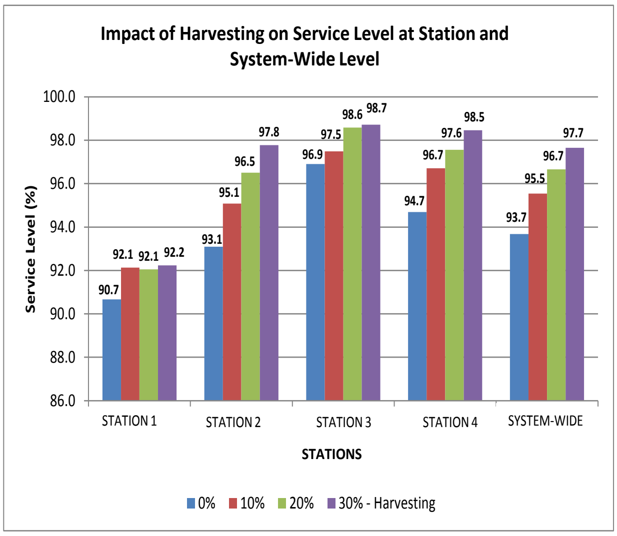

Figure 7: Impact of Engine Harvesting Process on Service Level at Station and System-Wide Level (multi-location case).

The results indicate that the harvesting process has a positive impact in the service level observed not only at station level but also system-wide. As expected, the service level is increasing with the level of harvesting. That is, as more engines can be harvested instead of conducting the heavy repairs that require long turn-times, then the new spares are made available faster to the stations, which in turn, leads to a reduction of number of OTS events and, consequently, to the increase of the service level. Moreover, notice that in this case the harvesting process is advantageously utilized to increase the service level without increasing the total ownership which is constant. Thus, important financial benefits can be obtained for the company by avoiding the need of additional investment in either new spares or expensive heavy repairs while achieving higher service levels. As an example, notice that station 2 increases its service level from 93% to about 98% by allowing a 30% of harvesting in the process.

5. CONCLUSIONS

In this paper we presented the application of a simulation approach for planning the required ownership levels of spare engines and parts. In maintenance operations for the airline industry this is an important problem not only because of its financial significance due to the high value of the engines, but also because of the need to effectively and reliably support the company's flying operations.

The implemented simulation models have demonstrated to be a useful tool not only for planning but also for decision support as illustrated in the different case studies presented here. Moreover, the models implemented have demonstrated its versatility in providing additional performance measures such as OTS-related metrics, which are useful for planning and selecting optimized strategies for engine maintenance operations.

In the context of aircraft maintenance operations, engines are not the only high value assets for which the models can be applied. There are an important number of other repairable aircraft components for which the models can be also applied. Currently, the company is in the process of extending the application of these models to other type of assets.

Finally, there is still opportunity to expand and refine the models described here by including additional and more sophisticated features. For instance, advanced simulation-based optimization approaches could also be considered in the case of the multi-location model in order to better tune the required inventory levels at different stations while meeting specific performance requirements.

ACKNOWLEDGMENTS

The authors thank Matt Pfeifer, Richard Czuchlewski, and Juan Leon from Operations Strategic Planning at American Airlines, as well as the Engine Production Control team at the American Airlines Tulsa Maintenance Base, for all their support and the valuable insight provided in the development and refinement of the models described in this paper. Special thanks to Jim Diamond, Managing Director of Operations Research & Advanced Analytics in American Airlines, for supporting the publication and presentation of this paper in the 13th International Conference on Modeling and Applied Simulation.

REFERENCES

Ackert, Shannon 2011. *Engine Maintenance Concepts for Financier*, Aircraft Monitor. Available from: http://www.aircraftmonitor.com/uploads/1/5/9/9/15993320/engine_mx_concepts_for_financiers___v2.pdf [accessed 20 March 2014]

The Apache Software Foundation 2014. *The Apache Commons Mathematics Library*. Available from: http://commons.apache.org/proper/commons-math/ [accessed 7 July 2014]

CFM International 2013. *2013 CFM56 and LEAP engine orders soar*, CFM Press Center. Available from: http://www.cfmaeroengines.com/press/2013-cfm56-and-leap-engine-orders-soar/722 [accessed 20 March 2014]

Guide, V. Daniel R., Srivastava, Rajesh, 1997. Repairable inventory theory: Models and applications. *European Journal of Operational Research*, 1 (102), 1–20.

Kelton, W. D., Sadowski, R. P., Sturrock, and D. T., 1997. *Simulation with Arena*, 4th ed. USA: McGraw-Hill, 2006.

Oracle Corporation 2014. *Java SE at a Glanse*. Available from: http://www.oracle.com/technetwork/java/javase/overview/index.html [accessed 7 July 2014]

Sargent, Robert G., Verification and Validation of Simulation Models, *Proceedings of the 2010*

Winter Simulation Conference, pp.166-183. December 5-8, Baltimore, MD (USA).

SAS Institute Inc. 2014. *SAS® 9.4*. Available from: http://www.sas.com/en_us/software/sas9.html [accessed 26 March 2014]

Sherbrooke, C. (1968): METRIC: *A multi-echelon technique for recoverable item control*, in Guide Jr., V. Daniel R., and Srivastava, Rasjesh (1997): *Repairable inventory theory: Models and applications. European Journal of Operational Research*, 1 (102), 1-20.

Simao, Hugo, Powell, Warren, 2008. Approximate Dynamic Programming for Management of High-Value Spares. *Journal of Manufacturing Technology*, 2 (20), 1–20.

Teradata Corporation 2014. *Teradata Database*. Available from: http://www.teradata.com/products-and-services/Teradata-Database [accessed 26 March 2014]

Tysseland, B.E., Halskau, H. 2007. Spare Parts Inventory – A Literature Review with Focus on Initial Provisioning and Obsolescence Management. *The 19th Annual NOFOMA Conference Proceedings*, pp. 1075-91. June 7-8, Reykjavic (Iceland).

JOSE A RAMIREZ-HERNANDEZ received the Ph.D. degree in Electrical Engineering from the University of Cincinnati, Cincinnati, OH, in 2010. He has B.S. (1995), Licenciatura (1996) and M.S. (1999) degrees (both with honors) in Electrical Engineering from the University of Costa Rica, San Jose, Costa Rica. From 1996 to 2001, he was an Adjunct Professor in the School of Electrical Engineering at the University of Costa Rica; from 2001 to 2007, he was member of the Systems Modeling and Information Technology Laboratory (SMITLab) at the University of Cincinnati; and since 2010 he works as Senior Operations Research Consultant within the Operations Research & Advanced Analytics and Operations Strategic Planning groups at American Airlines, Inc. His research interests include simulation, Markov decision processes, dynamic programming, and simulation-based optimization methodologies with applications to commercial aircraft maintenance operations. Dr. Ramirez-Hernandez is also member of IEEE, INFORMS, and the Omega Rho honor society.

STEVEN OAKLEY Received the Ph.D. in Systems and Industrial Engineering from University of Arizona (UA), Tucson, AZ, in 1990, his M.S. in Industrial Engineering (UA), 1986, and his B.S. Engineering Mathematics (UA, High Distinction), 1985. In 1990-91, he was Visiting Assistant Professor at the University of Tennessee in Knoxville, TN. Since 1991 he has been a technical consultant for American Airlines in the areas of revenue management and maintenance and engineering. He specializes in applied probability modeling and statistical science. Dr. Oakley's professional affiliations include INFORMS, American Statistical Association, and Tau Beta Pi.

MEI ZHANG received the Ph.D. degree in Industrial and Systems Engineering from Georgia Institute of Technology in 1997, M.S. in Geography from University of Tennessee, Knoxville (UTK) in 1993, B.S. in Mathematics from UTK in 1991. She has worked at Sabre and i2 Technology briefly before joining American Airlines in 2001. Currently, she is Senior Manager of Operations Research and Advanced Analytics at American Airlines supporting Technical Operations (Tech Ops) which includes the Maintenance and Engineering organization. Dr. Zhang and her team apply business process modeling, optimization, simulation and statistical analysis methodologies to help business units to improve efficiencies in their planning and operations. Dr. Zhang is a member of INFORMS and IIE.

ALEJANDRO SCALISE received the Ph.D. in Industrial Engineering from University of Pittsburgh (UP), Pittsburgh, PA, 2001. M.S. in Industrial Engineering (UP), 1999, B.S. in Industrial Engineering, Universidad Nacional de Cuyo (UNC) Mendoza, Argentina, 1986. Between 2001 to 2004 he was Airport Consultant at TransSolutions, Fort Worth, TX. Since 2004 he has been a consultant for American Airlines in the areas of airports and customer service, providing advanced analytics consulting for decision support. Dr. Scalise specializes in simulation modeling, operations research, and computer programming.