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

LIFE-CYCLE ENGINE FLEET SIMULATION FOR SPARE PART INVENTORY MANAGEMENT WITH ADVANCED CONDITION INFORMATION

Ana Muriel
Michael Prokle

Dept. of Mechanical & Industrial Engineering
University of Massachusetts Amherst
160 Governor's Drive
Amherst, MA 01003, USA

Robert Tomastik

Global Materials & Logistics
Pratt & Whitney
400 Main Street
East Hartford, CT 06118, USA

ABSTRACT

The cost efficient management of spare parts for low-volume high-tech equipment is inherently difficult. In this on-going study, we seek to improve the OEM's spare parts inventory management by incorporating the condition information from a large number of distributed working units in the field. For that purpose, the condition information relayed by sensors is put in context with usage parameters, preventive replacement policies, customer plans, and current economic indicators to create an aggregate forecast and inventory ordering policy. This requires a synthesis of the state of the art knowledge from multiple research streams. In this paper, we outline a simulation environment of the maintenance management of a jet engine program over its life cycle, and provide preliminary results highlighting several modules for future research to improve the performance of spare part inventory policies and assess the value of health monitoring.

1 INTRODUCTION

Stochastic part deterioration makes the prediction of equipment maintenance and associated spare part demand very difficult. Yet, accurate forecasting and spare part inventory management is crucial to keeping expensive equipment running and maintenance costs manageable. Condition monitoring involves collecting real-time sensor information from a functioning device to make predictions regarding the health condition and lifetime of each unit, and is thus posed to improve maintenance decisions. By aggregating over the condition of an entire fleet, this information not only promises improved maintenance scheduling but also better management of the resources needed - in particular spare parts.

Our work is part of a multipronged and interdisciplinary study that develops the methodologies necessary to utilize sensor readings from a large number of distributed working units in the forecasting and inventory control of the spare parts necessary for maintaining those units. The research consists of four key milestones (as outlined in the NSF Abstract #1301188):

1. "Advancing sensing methods and the interpretation of signals to diagnose equipment condition".
2. "Developing procedures for transforming these data into predictions of time-to-overhaul and resource-requirements".
3. "Building part forecasting methods and inventory policies that aggregate this information across equipment, under consideration of field usage and economic conditions".
4. "Creating a simulation tool for the monitoring and maintenance of a large fleet to validate the methodology".

This paper focuses on the last milestone, building a simulation environment to validate, compare, and further optimize the study's proposed methodologies and inventory policies. The ultimate objective is to highlight the economic value of advanced sensing techniques. We focus on a particularly complex and high-stakes application: jet engine fleet management. The associated spare part inventory management for maintenance, repair, and overhaul (MRO) services is critical to the availability of the fleet but is inherently challenging for the following reasons: (1) A high fraction of parts are very capital intensive since jet engines not only push the limits on what is technically possible in terms of operational performance, but also provide a safety critical service that requires the best material and design to ensure reliability. (2) Tier one suppliers need to account for a large number of distinct spare parts for which safety stock and inventory easily adds up to a large operating capital. For example, Mabert et al. (2006) report that Pratt & Whitney stocks more than 22,000 distinct parts. (3) The demand rates are very low and often lumpy and intermittent which can lead to long and costly inventory holding. (4) Demand for these spare parts may result from actual part failure, operators' economic decision-making to procure and stock these parts, or decisions to perform maintenance out-of-schedule to accommodate lease agreements, optimize overall fleet operations, or take advantage of maintenance contract terms. (5) Tier one suppliers experience long and highly variable procurement lead-times for critical parts. In order to counteract the uncertainty on the demand and supply side, the implementation of highly expensive safety stocks is the common practice to prevent stock-outs.

These challenges can be further illustrated in the light of the most recent 2008/2009 economic downturn. The available seat kilometers followed closely world GDP growth. Aircraft usage declined significantly and airlines were facing a significant loss of revenue. Accounting for the high cost of ownership of an airplane, deferring MROs helped airlines reduce fleet costs. However, it led to significant inventory build-up of replacement parts impacting tier one and tier two suppliers significantly. In fact, the value of inventory of parts stocked at tier one, tier two, and small part suppliers has been estimated at \$40 billion in 2010 (Clearwater 2011), a number that corresponds to the entire MRO market size in 2010 (Reals 2010). In an industry that is already inherently difficult, this triggered companies to re-examine their operations and spare part inventory management. In this capital intensive environment, small improvements have significant financial impact.

2 LITERATURE REVIEW

Our study requires a synthesis of multiple traditional management science research streams, including demand forecasting, inventory management, and maintenance scheduling combined with insights from aerospace part degradation modeling and the interpretation of on-board condition monitoring sensor signals. There is a remarkable number of studies that take on each complex individual task, or a subset of these tasks; for a full literature review, we refer the reader to Prokle (2017). However, previous research that takes on the holistic approach of translating condition-based sensor readings into maintenance scheduling decisions and spare part inventory ordering decisions is very sparse (see Park and Ryan 2015; Park 2015). In this section, we focus on previous fleet management simulation studies.

General research in this area has been rooted in industry or defense environments, driven primarily by pressing problems and practical incentives to solve them. These case studies provide excellent insights but are often limited in revealing the full scope of the technical details involved. In collaboration with Bombardier, Gharbi et al. (1997) seek to optimize a yearlong maintenance program for the Canadian fighter CF-18. The authors attribute poor schedule performance to the variability caused by changes in work scope authorizations, deterministically estimated work-step durations, and unplanned maintenance activities which delayed and shifted resources between different projects. Their simulation model uses the initial production plan (status quo) as input and delivers an improved final production schedule by incorporating elements, such as resource constraints, unpredicted failures, or variability in work durations, and allowing for specific user input. Using historical data, the model is found to not only produce superior production plans, but also provide high value in performing 'what-if' analysis, helping to negotiate maintenance aspects with customers. Gatland et al. (1997) approach a similar capacity and facility loading problem for a fleet of engines and provide insights from Delta Airlines. The authors build an Arena simulation model which analyzes the impact of varying (1) engine removal times, (2) engine disassembly start times, (3) disassembly work schedules, and (4) engine workscope mix. KPIs considered in the study include engine turnaround time, engine throughput, engine service level, machine utilization, personal utilization, part turnaround time, part throughput, average throughput, and work-in-progress. The industry funded study of Stranjak et al. (2008) uses a multi-agent approach to the problem of overhaul prediction and scheduling to navigate the competing objectives of minimizing operational maintenance costs and decreasing waiting times. The reliability of the engine is being approximated by the Weibull function, aggregating the part specific probability distributions. Using the function's scale and shape parameters, the authors distinguish between different life stages of the engine and types of disruptions (i.e., infantile, random, and wear-out) and schedule overhauls as close as possible to their predicted optimal overhaul date under capacity restrictions. Besides utilization, turnaround times, and aircraft-on-ground occurrences, the authors also capture the impact of the number of spare engines available. Painter et al. (2006) use Arena to simulate fleets of engines using military mission profiles and estimate the long-term cost effects (i.e., life-cycle costs) of maintenance policy decisions influencing key performance indicators like expected time-on-wing, cost-per-engine flight hour, and the operational fleet availability. The authors argue that historic data is a bad estimator for future costs. Instead, they develop a simulation coupled with data mining techniques to, first, generate a data set of maintenance history and cost statistics, and, then, build a life-cycle cost model that uses appropriate static and non-static cost estimation parameters (e.g., mission profiles or operating environments). Statistical sampling determines scope, timing, and location of failure. Data-mining, regression, classification, and clustering techniques are used to identify key life-cycle limit cost drivers. Mattila et al. (2008) also simulate flight missions and model the maintenance of fighter aircrafts for the Finish Air Force under normal and conflict situation conditions with the objective to improve decision-making in fleet maintenance operations. In their Arena model, the authors consider three types of maintenance needs: (I) periodic maintenance, where the model criteria are cumulative flight hours and predetermined service intervals, (II) failure repair, modeled as time between failures, and (III) battle damage, which is modeled as pass-fail probabilities. The configuration of fleet and maintenance operations constitutes the simulation input, and aircraft availability, maintenance, and flight performance statistics are generated as output. Maintenance network locations are considered and incorporated into their model. The authors highlight the challenge of the scarcity of data, its confidentiality, and the important insights of subject matter experts to overcome some of the unknown factors of the model.

In summary, the research areas this study incorporates are expansive and multifaceted. No previous work has been found integrating all, to model the entire process from condition monitoring to spare part inventory management and assess the value of sensor information over the life-cycle of an engine program.

3 SIMULATION FRAMEWORK

This section provides a high-level view on the simulation framework that will allow us to assess the value of incorporating fleet sensor information into spare part inventory management. Details on the various modules will be given in the next section.

The main input parameters can be classified into four major categories: (1) economic indicators, (2) engine profile, (3) cost parameters, and (4) sensor information. Figure 1 depicts relevant parameters within each category and their dependence. The engine profile details unique characteristics of the engine and its main parts, their usage history, condition information, and maintenance information. Part usage history (flight cycles flown) is necessary to capture the repair and reuse of parts, while keeping track of the parts' life limit. We use flight cycles as the key parameter affecting the engine condition since airplane starts put most stress on the engine and are the major driver of degradation. Economic conditions impact both engine usage and maintenance decisions (airlines strapped for cash may delay expensive overhaul by assigning the engine to a lighter flying schedule or using it as a spare). We assign a status to each engine reflecting its current activity in the fleet: in service, currently in overhaul, spare engine, or retired. Generally, sensor data must be translated into probabilistic information regarding 1) when the engine will be overhauled, 2) what modules will be included in the workscope, and 3) what parts will need replacement.

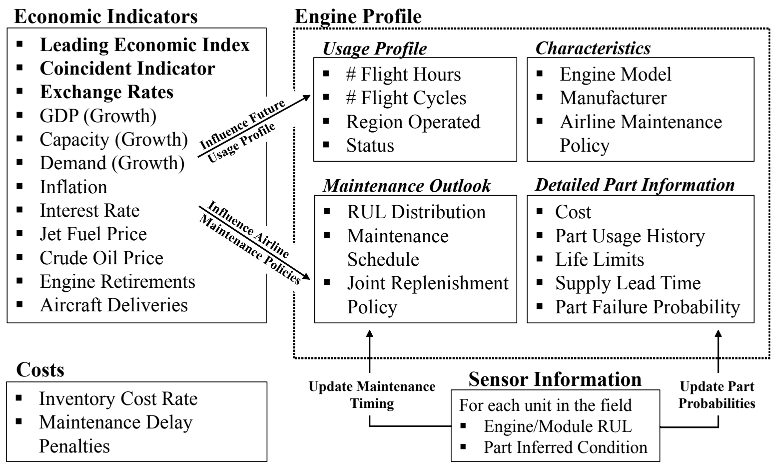

Figure 1: Major inputs to the simulator.

We develop a discrete-event simulation with weekly time steps. Figure 2 (left) provides a high level outline of the simulation framework to compare a traditional spare part inventory management approach to our proposed condition-based inventory control and fleet management process. Figure 2 (right) provides further detail on the engine program life-cycle simulation. Engines are added to the fleet according to a given production schedule, and assigned a degradation and initial usage profile. Every week (time step), we simulate the fleet operation, that is, cycles flown for each engine during the week. The sensor readings after that number of flights are provided by the NASA C-MAPSS simulator, and the particle filtering approach in Wang et al. (2015) is used to generate a remaining useful life (RUL, in flight cycles) distribution for the engine, and potentially for its modules and parts. Economic conditions are used to predict engine usage over time and translate the life cycles into calendar time. For each modeled spare part, the distribution of spare part demand over its lead time is computed by aggregating the probability distribution of demand during the part lead time over all the engines in the fleet using the condition-based RUL distributions. [A traditional inventory policy ignoring condition information can also be used, and its performance compared to assess the value of condition information].

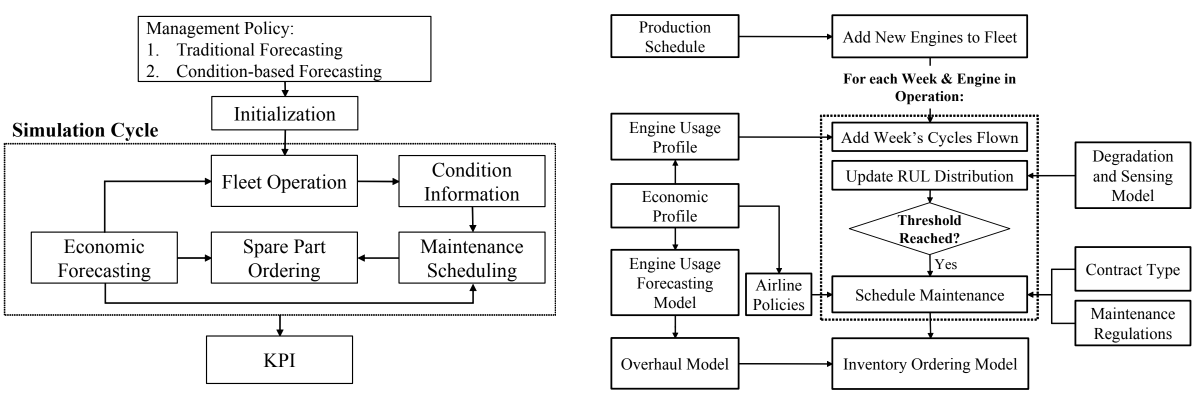

Figure 2: Overviews of the simulation framework.

When the engine's condition reaches a certain threshold, an overhaul is scheduled according to individual airline policies, contract types, and maintenance regulations.

While the parameters in the engine profile and sensor information are updated weekly, economic indicators follow quarterly updates, as they change slowly over time and have a longer-term influence on flight cycles flown and airline behavior.

The following steps illustrate how a particular engine transitions during the simulation over the following stages:

1. The engine is being produced and enters into service. We add the engine to the existing fleet of engines and assign random weekly flight cycle usage. The engine stays in service until an overhaul is scheduled due to abrupt fault, specified part life limit, or regular wear and tear.
2. The engine has reached the overhaul criteria. This is a prespecified threshold defined as either a fixed number of flight cycles remaining with a specified probability or an individual part life limit. This triggers a spare engine into service and sends the engine into overhaul.
3. The engine has reached the end of overhaul. The maintenance duration is randomly assigned and might further be increased in weekly increments when parts or resources are not available. The engine is added to the pool of spare engines after maintenance is complete.
4. The engine switches its status from spare to service and replaces an incoming engine to be overhauled.
5. Steps 1-4 will be repeated until the engine is retired. However, duration and timing might change significantly based on evolving engine and system conditions and the stochasticity in each simulated period.
6. The engine reached a specified threshold in age measured as number of lifetime flight cycles and is retired.

4 SIMULATION MODULES

The simulation is developed in Matlab 2017a environment to interface with the previous work of our collaborators (Milestone 1-3 in introduction). The following describes the key modules necessary to capture fleet degradation and maintenance operations over time.

4.1 Engine Usage Forecasting Model

The economic forecasting model seeks to forecast engine usage, measured in flight cycles, given the available economic data introduced in Figure 1. This is done in two steps. First, we select the most relevant subset of the economic indicators (features) on-hand using Matlab's R2017a Statistics and Machines Learning Toolbox and its sequential feature selection function *sequentialfs* which builds pools of subsets and sequentially adds and test features. Random partitioning of training and test sets as well as tenfold cross-validations assure statistical validity. Second, once the most reliable subset of features has been selected, a regression model with ARIMA time series errors is built using Matlab's Econometrics Toolbox (*regARIMA* class). This allows the estimation of regression coefficients, forecast future flight cycles, and confidence intervals. Furthermore, it accounts for typical seasonality in flight cycles flown while also testing time lags of the features used for the forecast.

In forecasting cycles flown four quarters ahead in a preliminary study, the feature selection algorithm identified four predictor features: Worldwide GDP, Number of Installed Engines, Worldwide Rate of Inflation, and OECD Composite Leading Indicator (MEI). The ARIMA model applied to these features results in a Mean Absolute Percentage Error (MAPE) of 2.79%.

The engine usage forecasting model is used to predict the number of flight cycles flown by each engine over a particular time period. That is, it is used to translate calendar time into cycles flown by the engine, since engine degradation is modeled over cycles flown, but overhaul operations and part supplies are scheduled over time. In the simulation, the flight cycle predictions will be updated only on a quarterly basis since economic conditions represent a high level view and change slowly over time.

4.2 Degradation and Sensing Model

Jet engine sensor data allows improved estimation of time-to-overhaul, overhaul scopes, and associated part requirements. Various methodologies and techniques have been proposed for reliable remaining useful life (RUL) predictions. Most recent work of our collaborators includes using deep convolutional neural networks for health monitoring and fault classification (Wang et al. 2017), automated performance tracking (Wang and Gao 2017), as well as Bayesian approaches and particle filtering techniques for wear predictions and lifetime estimation (Wang and Gao 2016; Wang and Gao 2015; Wang et al. 2015). In this section, we provide a high level description to provide an understanding for RUL distribution and its derivation, as inputs to the simulator.

Generally, gas path analysis aims to detect physical faults in a part (e.g., fan, compressor, or turbine) which caused changes in performance (i.e., efficiency or flow capacity) producing changes in measurable parameters (i.e., pressures, temperatures, or speeds). The analysis can estimate the state (i.e., efficiency) of a given part based on the sensed parameters. The posterior distribution for the state is estimated using particle filter Bayesian approaches on the weekly updated observable parameters.

Figure 3 illustrates how sensor data is continuously measured and used to estimate the RUL of the engine. The figure illustrates the measurement of sensor data over 100 flight cycles. Although measurements naturally are very noisy, a trend can be observed. Using the previous measurement, a future path can be predicted using the particle filtering method.

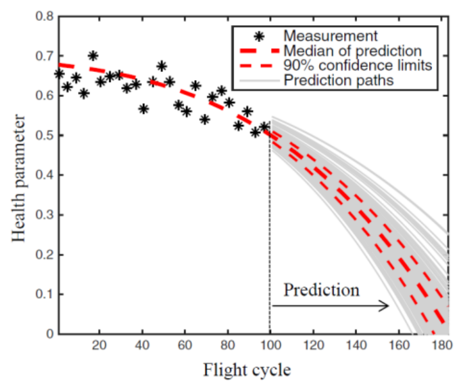

Figure 3: Estimation of remaining useful life distribution using particle filtering (Wang and Gao 2016).

The noise and uncertainty in the system provides multiple paths and results in a range of possible remaining flight cycles until the engine reaches an exogenously specified failure threshold that would prompt overhaul. These values can then be used to compute the probability distribution for the remaining useful life of parts or the overall engine:

$RUL_e(t)$: Remaining useful life for engine e as estimated at time t.

$RUL_{ie}(t)$: Remaining useful life for part i on engine e as estimated at time t.

Between overhauls, each simulated engine is randomly assigned a specific degradation curve from a library of possible degradation curves. Some engines may experience abrupt flight disruptions (e.g., bird strikes), which are captured by a sharp step decline in the degradation curve. Sensor readings and the subsequent RUL predictions are updated weekly according to the number of flight cycles flown in the corresponding week. Each simulation cycle, the health status of the engine is checked and the engine is scheduled for overhaul when (i) the probability of a remaining useful life of the engine is below 150 flight cycles with a probability of $P \geq 0.95$ or (ii) the life-limit of a part has been reached.

As a starting point, our case study uses sensor information to determine RUL distributions at engine level. RUL is defined as the number of cycles until the engine gas temperature (EGT) reading reaches a certain threshold, as is common in the modeling of the degradation of engine condition (e.g., Saxena et. al 2008). Future developments in degradation modeling of the various modules and major parts will be incorporated to generate RUL distributions at the module and part levels. This is a key driver to future progress in our research context since it has the potential to refine the individual forecast of workscopes and associated spare parts needed during an overhaul process. Our simulation framework can be extended and used to assess the value of additional condition information through new sensors.

4.3 Overhaul Model

Figure 4 provides insight into the modeling of a maintenance event within the simulation. We focus on maintenance events that require engine removal and an available spare engine as replacement. The swapped engine and its defective parts are then either being repaired and returned into inventory for reuse, or scrapped. Each maintenance step may have limited resources (e.g., spare parts or mechanics) and may infer delay. We assign a random duration to each overhaul and introduce further delay when required spares are not available at the point of engine reassembly. The engine will only move to reassembly when all required parts are assigned to it. We initially consider the following two policies to determine which parts will be required during overhaul:

1. *Deterministic Policy*: Exchanges a predefined set of parts in every overhaul.
2. *Random Policy*: Uses a random number generator and part specific probabilities of failing to determine which parts are being exchanged.

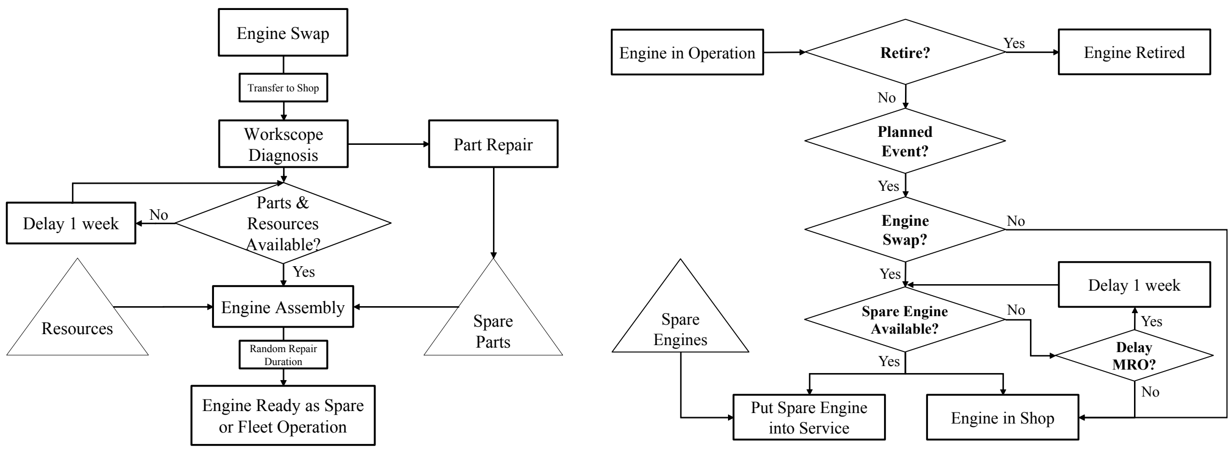

Figure 4: Overhaul process (left) and engine state transitions (right).

4.4 Inventory Ordering Model

The spare parts inventory ordering model uses the condition of the individual units in the field to provide an aggregate view of the distribution of part demand over the uncertainty period (lead time plus reorder interval) and generate a base-stock level as follows. For each particular engine e and time t:

- The sensing module provides a distribution of the number of cycles remaining useful life of the engine.
- This distribution is transformed into a distribution of remaining useful life $RUL_e(t)$ in calendar time, using the predictions on cycles flown per week for that particular engine.
- For a part i with lead time L_i, the probability of engine e requiring part i over the part lead time plus the reorder interval (1 week) can then be determined as $P_{ei}(t) = P[RUL_e(t) < L_i + 1] * p_{ei}$, where p_{ei} is the probability of part i being required for the overhaul.
- The aggregate demand for part i over the lead time plus reorder interval, that is over the relevant uncertainty period $[t, t + L_i + 1]$ can then be approximated by a normal distribution with mean $\mu = \sum_e P_{ei}(t)$ and standard deviation $\sigma^2 = \sum_e (P_{ei}(t) * (1 - P_{ei}(t)))$.
- The base-stock for part i at time t required to achieve a desired service level α is given by

  $S_i(t) = \sum_e P_{ei}(t) + z_\alpha * \sqrt{\sum_e P_{ei}(t) * \left(1 - P_{ei}(t)\right)}$, where z_α is the standard normal safety factor.

This provides a basic inventory model built upon the aggregation of the condition information of units in the field. A major thrust in the future work is to improve upon this model. In particular, the current inventory ordering model considers each part in isolation and thus ignores the assemble-to-order nature of overhaul operations, where all required spare parts in that specific overhaul need to be there to proceed with the reassembly of the product. If the goal is to support overhaul operations and ensure with high probability that engine maintenance is not delayed by the unavailability of spare parts, the inventory ordering problem should be formulated as a multi-product assemble-to-order (ATO) problem subject to order fill-rate constraints with non-stationary demands. The order fill-rate (in our case overhaul fill-rate) only considers the fraction of orders for which all needed parts can be provided within the specified time window without delay. In the ATO literature, the demand process is generally assumed to be stationary. Our setting, however, involves a demand process that is continuously changing based on the condition of the fleet of engines as they age and accumulate flying hours. In addition, the number of products that need to be considered to represent all potential overhaul part requirements is 2^n, where n is the number of parts being modeled; any combination of new parts may be needed, especially since some parts can be repaired or taken from a used part pool instead of purchasing a new one.

4.5 Key Performance Indicators

To assess the performance of the proposed inventory policy, we use the following key performance indicators:

- *Inventory Cost*: The overall cost incurred for parts held in inventory over the simulation period.
- *Part Fill-Rate*: Percent demand for individual spare parts, as needed in maintenance operations, satisfied directly from stock.
- *Part Induced Delay*: The average delay in engine overhaul that an individual part caused during maintenance operations.
- *Maintenance Fill-Rate*: Percent of engine overhauls for which all spare parts are directly available from stock.
- *Average Engine Maintenance Delay*: Average engine delay during maintenance operations.

In addition, to evaluate how spare part inventory policies impact and interact with the overall fleet management system, we also calculate:

- *Aircraft-on-Ground* (AOG): Number of aircraft grounded as their engines require overhaul and no spare engine is available.
- *Number of Spare Engines*
- *Fill-Rate of Spare Engines*: Percent demand for spare engines satisfied directly from stock, to replace engines in the field grounded for maintenance operations.
- *Average Maintenance Turnaround Time*: Average time from induction to shipment after maintenance operations, which is affected by the availability of parts.

Using these metrics of system performance, we can characterize the value of condition monitoring by comparing (i) traditional demand-based stock levels and (ii) condition-based stock levels as described in Section 4.4. Furthermore, we can iteratively refine and test both inventory policies as we better understand their impact on overall fleet management performance. Finally, we would like to extend the simulation framework to assess and characterize the value of placing additional sensors and virtual sensing methods.

5 CASE STUDY AND RESULTS

In this section, we describe a simple case study carried out to highlight the power and capabilities of the current version of the simulation. Increased detail of the complex industry environment and improved decision-making models are still in the works. As an example, we simulate four parts and a moderate size engine program. The simulation platform can easily scale up.

The simulation runs over a period of 20 years (1040 weeks) and simulates weekly engine production of two engines for the first ten years, i.e., weeks 1-520. In addition, spare engines are introduced to the fleet on a continuous basis to fulfill a prespecified 10% requirement of spare engines in the fleet. Each engine is assigned to a fixed degradation profile, and associated noisy sensor readings after each cycle flown; 10% of the engines are assigned a fixed degradation profile that includes an abrupt fault (see Figure 5 for an example of regular (left) vs. abrupt (right) degradation profile). The number of flight cycles flown are randomly assigned to each engine and assumed to be constant in the short-term, and only newly assigned after either (i) new economic conditions are incorporated each quarter of a year, or (ii) the engine finishes maintenance operations and is installed in a new aircraft. The assigned flight cycles reflect the U.S. flight cycle numbers for passenger airplanes in the years 2003 to the end of 2012 available from the United States Department of Transportation. Engine retirement age is set to 40,000 flight cycles.

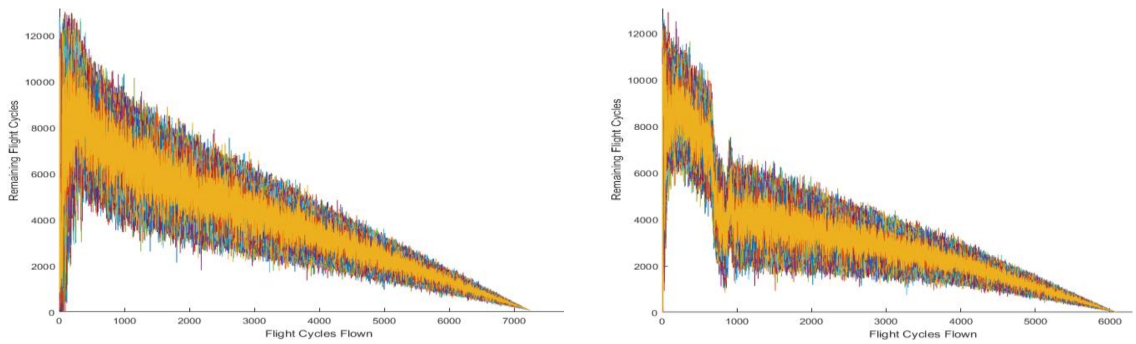

Figure 5: Remaining useful life distributions (# of remaining flight cycles) as a function of cycles flown.

We use the following part Lead Time (LT in weeks), Cost (C), and Starting Inventory (SI) characteristics:

- Part 1: 5 (LT) | 2000 (C) | 10 (SI)
- Part 2: 13 (LT) | 2000 (C) | 15 (SI)
- Part 3: 26 (LT) | 1000 (C) | 20 (SI)
- Part 4: 36 (LT) | 1000 (C) | 30 (SI)

Engines scheduled for overhaul are assigned a random maintenance duration according to a distribution with range 6-27 weeks and highest probability between 9-11 weeks. Parts are required to be in physical inventory three weeks before the scheduled end of the random maintenance durations. This reflects the time that is needed to reassemble the engine after the required replacement parts are available. In this initial case study, we consider a policy that requires all parts modeled to be exchanged at each maintenance visit. A missing part delays the maintenance by one additional week until all parts are available. Spare engines are installed in the fleet once an engine requires an overhaul.

Figure 7 (left) shows the total number of overhauls in each period over the simulation length. New engines are introduced to the fleet at a rate of two per weekly period, starting period 1 (until period 520). The figure shows how overhaul operations only start around period 150. Different RUL distributions and flight cycles flown influence the random time of the overhaul. A sudden spike of overhauls can be seen at period 580. Older engines requiring their second major overhaul start to overlap with newer engines requiring their first overhaul. Observe that we are not modeling the end of the engine program's life cycle, since most engines do not reach their life-span limits within the 20 years simulated. As a result, the number of overhauls is still close to its peak in the final weeks of the simulation. The inventory of spare engines over the 20 years is displayed in Figure 6 (right). Initially, the number of spare engines grows according to the specified 10% spare engine production requirement, as very few engines require overhaul. The figure clearly shows that the 10% spare engine production requirement overestimates the actual need for spare engines, as well as the fact that much of their production should be postponed to later in the program when overhauls will be prevalent. The sudden spike in overhauls translates into a sudden decline in spare engines, since every incoming engine requiring maintenance is replaced by a spare engine. Engines finished with maintenance operations are added to the spare engine pool, explaining the growth in spare engines even after engine production stops in week 520.

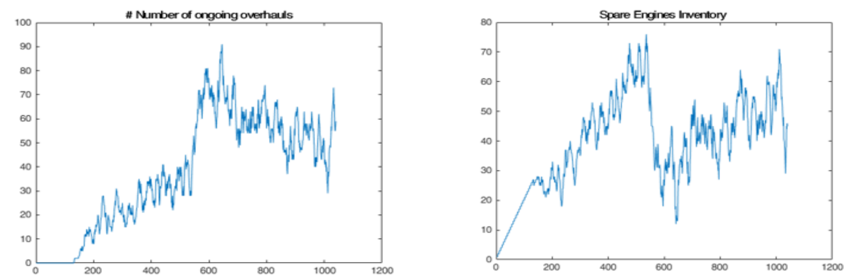

Figure 6: Engine overhauls (left) and inventory of spare engines (right) over time (weeks).

Figure 7 shows that the condition-based inventory order-up-to methodology (see Section 4.4) supports the trend in number of overhauls over time. The order-up-to level at a given point in time also reflects the lead time of the part. Parts 3 and 4 with the highest lead time have a higher inventory buffer than parts 1 and 2 with shorter lead time.

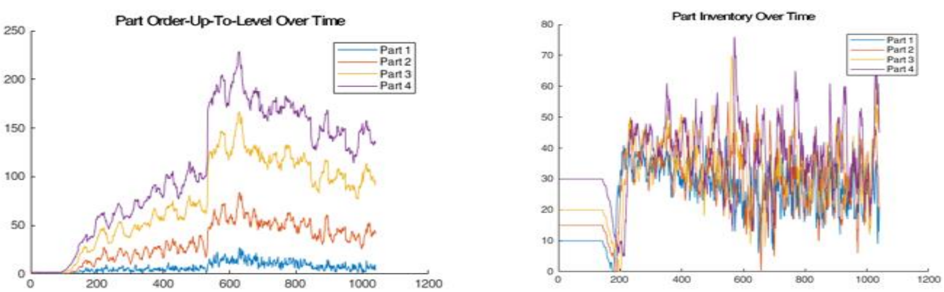

Figure 7: Order-up-to levels (left) and on-hand inventory (right) for each part over time (weeks).

The physical on-hand inventory shows high fluctuations. Orders arrive after their deterministic lead time and inventory is depleted according to the number of engines that require maintenance operation at a given point in time. We start with an arbitrary (rather large) initial inventory. The figure shows how inventory is depleted when demand for maintenance operations starts in period 150. Inventory only reaches zero twice in this simulation run in approximately periods 190 and 650.

6 CONCLUSIONS

This is one of the first studies to integrate multiple research streams (sensing, degradation modeling, RUL predictions, economic conditions, part demand forecasting, and inventory management) into one simulation framework. Various assumptions are essential to narrow down the most important building blocks for the study and reduce the complexity. Similarly, many extensions are possible in continuously refining the assumptions to better capture reality and understand the impact of various parameters and policy decisions. This will require continued close collaboration with our industry partners and much data gathering and analysis.

We see the broad impact of our work in the following. First, this study will help demonstrate the economic value of condition monitoring for improved spare part demand forecasting. Second, our condition-based inventory management approach should contribute to the reduction of inventory costs as well as increase fleet availability for commercial and military aircrafts. Third, the simulation framework can be used to evaluate the strategic implications for future development of sensor technologies by identifying the operational value of adding different sensors to the engine. Fourth, the study supports better decision-making for MRO service contracts over the full life of the program. The manufacturer can better assess costs and risks over the life of the engine and is, therefore, able to better assess their service offering (e.g., guaranteed availability of an engine) and set the pricing of the service contracts accordingly. Fifth, further research on condition sensing, data gathering, and analytics evaluating engines' health status contributes to keeping airplanes safe and reliable.

The simulation framework developed allows for consideration of different policies and models within each of its modules. In the current simulation tool, we have chosen a simple engine introduction and retirement schedule, a particular method of forecasting cycles flown given current economic conditions, a basic method for the inventory control of spare engines, a simple spare part inventory management policy given condition information, etc. Further research is needed for the careful selection of other models and policies to use within each module. Extensive experiments need to be run to evaluate the performance of different inventory policies, and understand the impact of the many parameters at our disposal. There are further relevant research directions and extensions we would like to highlight. First, our research should have significant impact on MRO service contracts, such as 'power-by-the-hour' agreements, which are designed to guarantee asset availability to customers under a predefined fixed cost model and help airlines build stable financial plans by lowering the risk of unexpected operational expenses. A higher confidence in spare part forecast will have significant effects on promised contract conditions (e.g., service levels) and pricing of the service (Justin and Mavris 2015), which would require further evaluation. Second, there are multiple streams of demand for a part, such as military spares, commercial spares, military production assembly, and commercial production assembly. Each demand stream has different requirements such as demand lead time or service level. Modeling the interaction of all streams appropriately is complex but is of practical relevance (Kocaga and Sen 2007). Third, there are multiple agents that are involved or influence MRO decisions. A future study could further model agents' behavior and evaluate their individual competing objectives.

ACKNOWLEDGEMENTS

This study has been supported by NSF CMMI-1301188 and is based on collaborative work with Professor Jennifer K. Ryan, Professor Robert X. Gao, and their PhD students. We thank UMass REU students Jessica Peterson and Nicole Lynch for their contribution, and Dr. Raj Subbu (PW) for his comments and insights.

REFERENCES

Clearwater. 2011. Aerospace Global Report - A Clearwater Industrials Team Report. http://clearwaterinternational.com/aerospace-global-report-2011/, accessed 19.06.2016.

Gatland, R., E. Yang, and K. Buxton. 1997. “Solving Engine Maintenance Capacity Problems with Simulation”. In *Proceedings of the 29th Conference on Winter Simulation 1997*, 892-899, edited by S.Andradottir et al. , Piscataway, New Jersey: IEEE.

Gharbi, A., J. Girard, R. Pellerin, and L. Villeneuve. 1997. “Bombardier Turned to Simulation to Validate the CF-18 Maintenance Program”. *Interfaces* 27(6): 22-34.

Justin, C. Y., and D. N. Mavris. 2015. “Aircraft and Engine Economic Evaluation for Fleet Renewal Decision-Making and Maintenance Contract Valuation”. In *Proceedings of the Institution of Mechanical Engineers, Part G*: *Journal of Aerospace Engineering:* 1–14.

Kocaga, Y. L., and A. Sen. 2007. “Spare Parts Inventory Management with Demand Lead Times and Rationing”. *IIE Transactions* 39(9): 879–898.

Mabert, V. A., A. Soni, and G. Campbell. 2006. “A Case Study of Pratt and Whitney Aircraft's Commercial Spares Planning”. *Communications of the IIMA*. 6(3): 2.

Mattila, V., K. Virtanen, and T. Raivio. 2008. “Improving Maintenance Decision Making in the Finnish Air Force through Simulation”. *Interfaces*. 38(3): 187-201.

Painter, M. K., M. Erraguntla, G. L. Hogg Jr, and B. Beachkofski. 2006. "Using Simulation, Data Mining, and Knowledge Discovery Techniques for Optimized Aircraft Engine Fleet Management". In *Proceedings of the 38th Winter Simulation Conference,* 1253-1260, edited by L.F.Perronne et al., Piscataway, New Jersey:IEEE.

Park, H. 2015. *Decision-making in Advanced Manufacturing*. Doctoral Dissertation, Department of Industrial and Systems Engineering, Rensselaer Polytechnic Institute. digitool.rpi.edu:8881/dtl_publish/11/176738.html.

Park, H., and J. K. Ryan. 2015. "Condition-based Inventory Management with Planned and Unplanned Maintenance". In *Proceedings of the 2015 Industrial and Systems Engineering Research Conference*, May 30 - June 2, 2015, Nashville, TN.

Prokle, M. 2017. *Theory and Practice of Supply Chain Synchronization*. Doctoral Dissertations, Department of Mechanical & Industrial Engineering, University of Massachusetts Amherst. https://scholarworks.umass.edu/dissertations_2/1120.

Reals, K. 2010. Forecasts Detail 10-year Outlook for Global MRO Market. https://www.flightglobal.com/news/articles/forecasts-detail-10-year-outlook-for-global-mro-market-347933/, accessed 30.01.2016.

Saxena, A., K. Goebel, D. Simon, and N. Eklund. 2008. "Damage Propagation Modeling for Aircraft Engine Run-to-failure Simulation". In *International Conference on Prognostics and Health Management 2008*, 1–9.

Stranjak, A., P. S. Dutta, M. Ebden, A. Rogers, and P. Vytelingum. 2008. "A Multi-agent Simulation System for Prediction and Scheduling of Aero Engine Overhauled". In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multi-agent Systems 2008*, May 12th – 16th, 2008, Estoril, Portugal, 81-88.

Wang, P., and R. X. Gao. 2015. "Adaptive Resampling-based Particle Filtering for Tool Life Prediction". *Journal of Manufacturing Systems* 37 (2): 528–534.

Wang, J., P. Wang, and R. X. Gao. 2015. "Enhanced Particle Filter for Tool Wear Prediction". *Journal of Manufacturing Systems* 36: 35–45.

Wang, P., R. Yan, and R. X. Gao. 2017. "Virtualization and Deep Recognition for System Fault Classification". *Journal of Manufacturing Systems* 44: 310-316.

Wang, P., and R. X. Gao. 2017. "Automated Performance Tracking for Heat Exchangers in hvac". *IEEE Transactions on Automation Science and Engineering* 14 (2): 634–645.

Wang, P., and R. X. Gao. 2016. "Markov Nonlinear System Estimation for Engine Performance Tracking". *Journal of Engineering for Gas Turbines and Power* 138 (9): 091201.

AUTHOR BIOGRAPHIES

ANA MURIEL is an Associate Professor and Associate Department Head in the Department of Mechanical and Industrial Engineering at University of Massachusetts Amherst. She received her M.S. and Ph.D. in Industrial Engineering and Management Sciences from Northwestern University. Her research focuses on various aspects of logistics and supply chain management. Her email address is muriel@ecs.umass.edu.

MICHAEL PROKLE is a Senior Data Scientist at Fortune Brands and Lecturer at Northeastern University. He received his Dipl. Wi.-Ing. from the Karlsruhe Institute of Technology, Germany and M.S. and Ph.D. in Industrial Engineering & Operations Research from the University of Massachusetts Amherst. His email address is michael.prokle@gmail.com .

ROBERT N. TOMASTIK is Associate Director of the Materials Management Center of Excellence at Pratt & Whitney, a division of United Technologies. In this role, he is responsible for decision support and analytics supporting PW's manufacturing operations, supply chain, and aftermarket. He has a Ph.D. degree from the University of Connecticut in Electrical & Systems Engineering. His email address is robert.tomastik@pw.utc.com.