RAND PROJECT AIR FORCE

This PDF document was made available from www.rand.org as a public service of the RAND Corporation.

Jump down to document ▼

The RAND Corporation is a nonprofit research organization providing objective analysis and effective solutions that address the challenges facing the public and private sectors around the world.

Support RAND

Purchase this document

Browse Books & Publications

Make a charitable contribution

For More Information

Visit RAND at www.rand.org

Explore RAND Project AIR FORCE

View document details

Limited Electronic Distribution Rights

This product is part of the RAND Corporation monograph series. RAND monographs present major research findings that address the challenges facing the public and private sectors. All RAND monographs undergo rigorous peer review to ensure high standards for research quality and objectivity.

Programmed Depot Maintenance Capacity Assessment Tool

Workloads, Capacity, and Availability

Elvira N. Loredo, Raymond A. Pyles, Don Snyder

Prepared for the United States Air Force

RAND PROJECT AIR FORCE

The RAND Corporation is a nonprofit research organization providing objective analysis and effective solutions that address the challenges facing the public and private sectors around the world. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

RAND® is a registered trademark.

Summary

Aging Air Force fleets have accrued material deterioration problems that have resulted in increasing maintenance workloads, which have, in turn, led to reduced availability of the fleets for operations and training. Nowhere has this problem been more apparent and severe than during the periodic inspection and repair of aircraft structural elements of PDM (see pp. 5–8).

PDM is conducted in large organic or contractor facilities where aircraft can be partially disassembled, inspected, and repaired. A typical PDM visit may require between 2,000 and 50,000 labor hours (depending on the fleet) and substantial material. The total labor required to complete PDM is expected to increase as a function of the age of the fleet. However, there are different perspectives on the form that this increase may take. One analytic community (which we refer to as *the engineers*) relies on engineering judgment and current planned workloads to theorize that future workloads might stabilize over the near term; another group (*the statisticians*) rely on statistically based cost and workload trends to theorize that workloads and costs will grow and that availability will decrease.

Traditional Modeling Approaches Have Limited Applicability

While detailed resource and process simulation models can be constructed for a specific facility at a specific point in time, the workload, processes, and resource availabilities change constantly. More prob-

lematic, the specific workflows used by competing entities (organic or contractor) are seen as a proprietary matter that affects their ability to compete for future workloads. As a consequence, few facilities are willing to share detailed information on their specific work processes.

We developed PDMCAT to be able to estimate the number of aircraft in PDM status, future inductions, and production levels and to rely only minimally on detailed information from inside a facility (see pp. 9–12). We also sought to rely on easily observable features, such as the number of docks for performing maintenance and recent measures of actual performance, so that having "inside" information was not critical to forecasts of future inductions or numbers of aircraft in PDM status (i.e., not available for operations and training).

To that end, we extended and elaborated the BJB model (Zahorjan et al., 1982) to include multiple servers within each job stage. The original model was developed for the operational design of computing time-sharing systems. Appendix A discusses queuing theory related to this model. The BJB model required very little information in the first place, and we were able to simplify its data requirements further and apply it to the PDM process. Chapter Three describes application of the model and its development; Appendix B presents more detail on our extension.

Testing and Demonstrating PDMCAT: The KC-135 Case

To test and demonstrate the model's capabilities, we applied it to the KC-135 PDM process described in Chapter Four, first examining how well the model was able to forecast recent PDM performance, then comparing two alternative forecasts of the future workload and evaluating capacity and PDM process-improvement options to maintain acceptable availability levels. That fleet was chosen because there was an ample amount of information about its recent workloads, number of aircraft in PDM status, and changing capacity. More important, that fleet had experienced a substantial change in the number of aircraft in PDM status during the years 1998–2002, so we believed it would

constitute a good test of the PDMCAT model's forecasting capabilities (see pp. 13–22).

The alternative forecasts reflected engineers' versus statistical workload predictions. The fleet reduction program example demonstrates how changes in fleet size would reduce the number of aircraft in PDM status as the number of aircraft inducted each year diminishes. Figure S.1 shows the aircraft purpose possession history of the KC-135 tanker fleet from the second quarter of fiscal year 1995 to the first quarter of fiscal year 2004.[3] This chart shows the increase in the so-called *depot-possessed aircraft* and the consequential decrease in that

Figure S.1
Changes in Depot Capacity and Required Workload Created a Bubble in Depot-Possessed Aircraft
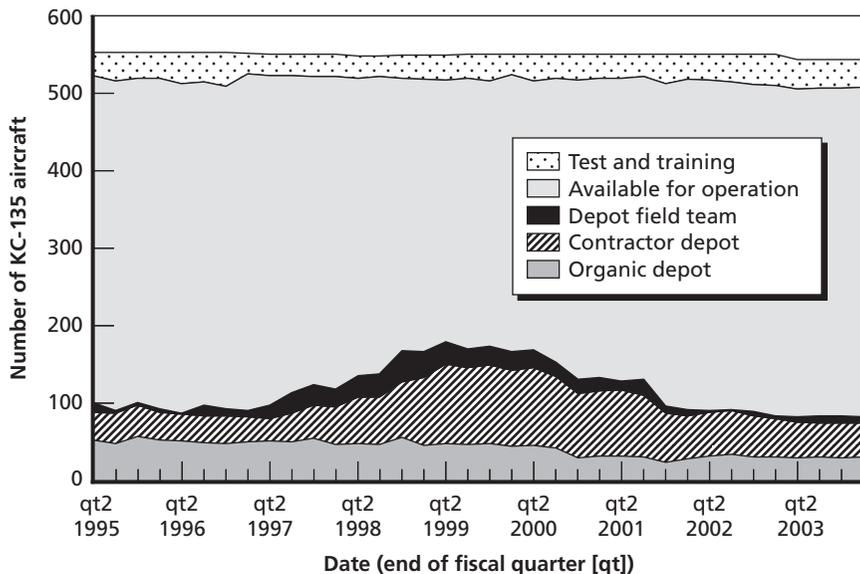
[3] The aircraft purpose possession history indicates how many Air Force aircraft are possessed for different purposes (e.g., test, training, modification, maintenance). It is constructed from detailed daily possession status change reports for each aircraft serial number. Most important for this study, it contains information from which one can compute the historical number of aircraft in PDM status and the number that entered PDM each year.

aircraft's availability for operations starting in the third quarter of 1997 and peaking in the second quarter of 1999—with almost 200 KC-135 tankers either in possession of depot field teams or at organic or contractor depot facilities. Our initial analyses addressed the PDMCAT model's ability to replicate that experience.

Initial Analysis of the PDMCAT Model

We used historical workload data to compare the model's forecasts to actual aircraft in PDM status during a critical transition period—from 1997 through 2003. During this time, the number of aircraft in PDM status increased by more than 50 percent, then returned to levels below the initial 1997 level. Figure S.2 shows that the PDMCAT model accurately reflected the increase and subsequent decrease in aircraft in PDM status.

Figure S.2
PDMCAT Forecasts Using Actual Workloads Match Actual In-Work Forecasts Using the PDMCAT Model
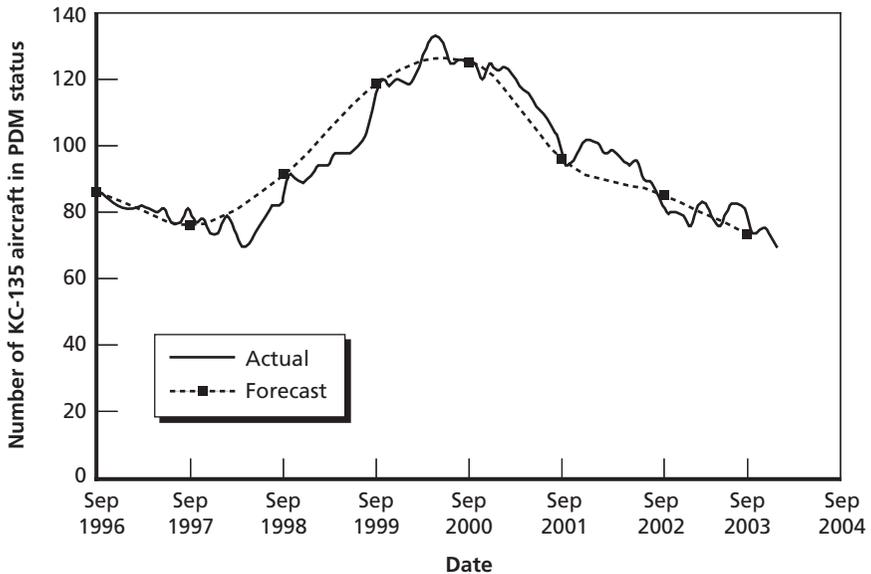
Using the PDMCAT Model to Assess Assumptions About Future Operations

Finding the historical match acceptable, we applied the model to test how assumptions about workload plans, induction schedules, labor application rates (often called *burn rates*), depot capacity, and fleet size would affect the forecast of near- and long-term inductions, production quantities, and aircraft in PDM status. A sample of how we used PDMCAT to test various assumptions is shown below.

Forecast of Future Workloads

Two forecasts of future PDM workloads were used in the Chapter Four analyses. The first, developed by the *KC-135 Economic Service Life Study* (ESLS) (Sperry et al., 2001), uses both statistical analysis and expert engineering judgment to estimate the effect of fatigue cracking and corrosion growth on future PDM workloads. The second is a purely statistical equation drawn from a PAF study that sought to discover and characterize maintenance life-cycle workload patterns that were common across all Air Force fleets, rather than a pattern that may reflect some idiosyncratic temporary behavior in a single fleet's history (Pyles, 2003). (See pp. 18–21.)

We used the model to examine both near-term (one to five years) and long-term PDM performance. In the near-term cases, we assumed there was only limited opportunity to increase PDM capacity, but that the PDM induction policy (i.e., the interval between subsequent PDMs) could be used to manage the workflow and aircraft availability. In the long-term cases, we assumed that it would be possible to add physical capacity (docks where aircraft could receive PDM maintenance) and to introduce process improvements that could increase the labor application rate (the number of labor hours that can be usefully applied to a single aircraft in a single day). (See pp. 21–24.)

Using PDMCAT to Moderate the Effects of Changes in Aircraft Induction Intervals on Near-Term Work in Process

The KC-135 fleet PDM process has experienced a turbulent period during which previously stable flow times and production rates were disrupted by a period of low production outputs followed by a period of higher-than-usual production outputs. If the KC-135 PDM managers were to follow Air Force Technical Order (AFTO) 00-25-4 (U.S. Air Force, 2003) interval prescriptions exactly, those production fluctuations would reappear as induction fluctuations, creating an imbalance between depot capacity and incoming workload requirements (see pp. 26–29). PDM managers have some leeway in adjusting aircraft induction intervals. This was the case in 2002 and 2003, when we found that the depot inducted five more (in 2002) and 28 more (in 2003) aircraft than required by AFTO 00-25-4 (see U.S. Air Force, 2003).

Figure S.3 shows how we used the PDMCAT model, along with the PAF workload forecast, to demonstrate the effect of those early inductions on aircraft in PDM status in subsequent years. Over the near term, the model projects a temporary reduction in the number of aircraft in PDM status, followed by an equally temporary increase in that number that would begin to approach the peak number of aircraft in PDM status from 1997 through 2003. The later increase was caused by a forecast increase in PDM workload coinciding with the scheduled return to PDM of the additional aircraft produced in 2003–2004. By adjusting the annual induction rates during these periods, we were able to use the model to identify an alternative induction plan that would reduce the peak number of aircraft in PDM status to acceptable levels through 2009.

Using PDMCAT to Test Assumptions About Long-Term Workload Growth, Increases in Capacity, and Burn Rates

Looking to the long term, which is depicted in Figure S.4, we found that the more pessimistic PAF workload projection would cause the depot flow times to increase until the "aircraft in PDM" status would

Figure S.3
PDMCAT Near-Term Forecasts Modulated by Changing Inductions
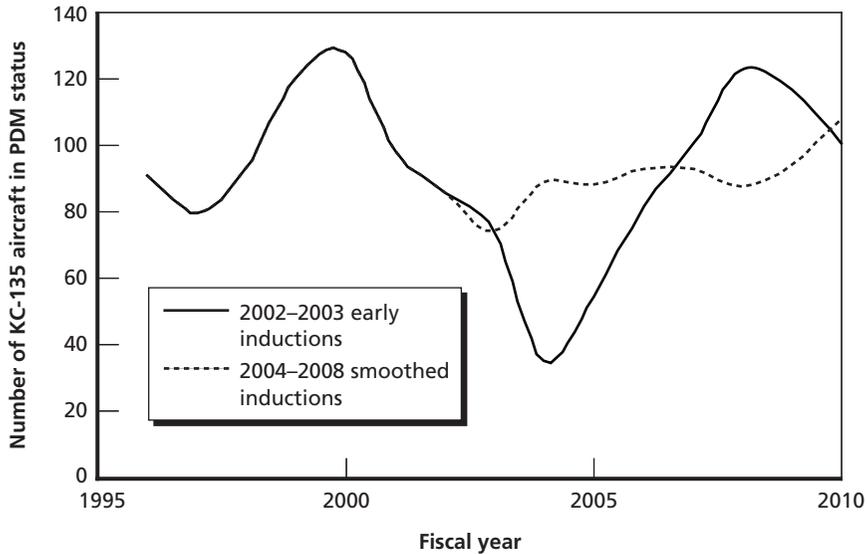
reach the 1997 to 2003 peak by 2013. We then increased either the physical capacity (number of docks where maintenance can be performed) or the labor application rate (a composite factor reflecting both labor available across all shifts and the degree of parallel operations in the PDM process) by 50 percent in 2010 to evaluate how those capacity increases might change the availability forecast. We learned that the increases both reduced the number of aircraft in PDM status and prolonged the time until the 1997 to 2003 surge peak was reached. The labor application rate option performed better, not reaching the 1997 to 2003 peak until 2024, compared to 2020 for the capacity increase case. We next examined the implications of the ESLS engineering-based workload forecast, which yielded a much more optimistic long-term outcome, never quite reaching the 1997 to 2003 peak (see pp. 32–43).

Figure S.4
Adding Capacity and Increasing the Labor Burn Rate Delay Impact of PAF Workload Forecast
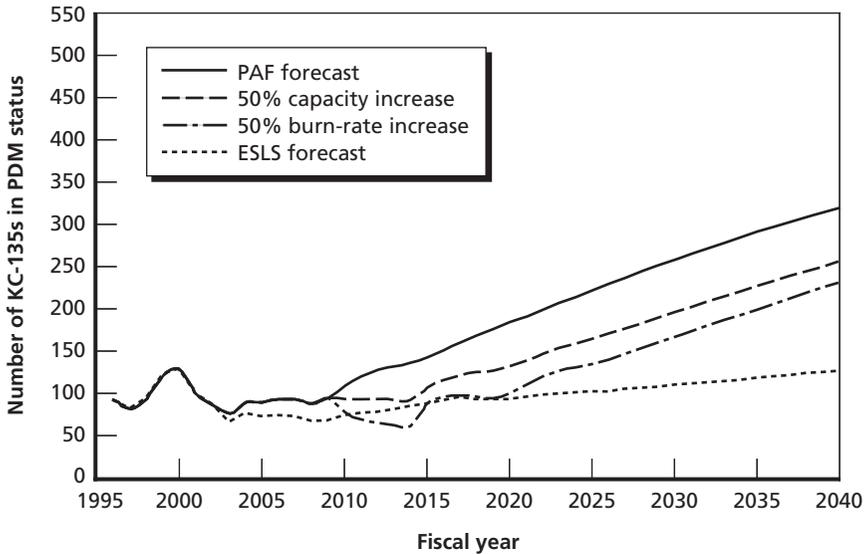
Using PDMCAT to Forecast the Effect of Changes in Fleet Size

In Chapter Five, we compared the PAF forecast against the KC-135 system program director's (SPD's) engineering-based forecast (see p. 20), assuming that the Air Mobility Command (AMC) plan to retire KC-135Es would have been implemented until only 490 aircraft remain in the fleet: 417 KC-135R/Ts and 73 KC-135Es.[4] We further assumed that the capacity would change in proportion to changes in the projected workloads (see pp. 47–49). With the KC-135 Tanker Sustainment Group's moderate forecast of PDM workloads, the PDMCAT

---

[4]   This plan was not implemented, but the analysis sheds light on how it would have affected KC-135 aircraft availability.

model projects that the aircraft in PDM will not reach the 100-aircraft level until after 2050.[5] Under the less optimistic PAF forecast, the PDMCAT model projects that the number of aircraft in PDM status will reach 100 as early as 2013, even if the fleet size is reduced as planned. This projection is contrasted with the results shown in Figure S.5 (PAF forecast 1). The conjunction of reducing the KC-135 inventory and increasing capacity significantly reduces the effect of increased workloads on aircraft availability.

Figure S.5
Reducing KC-135 Inventory and Increasing Capacity Dampen Surge in Aircraft in Work

---
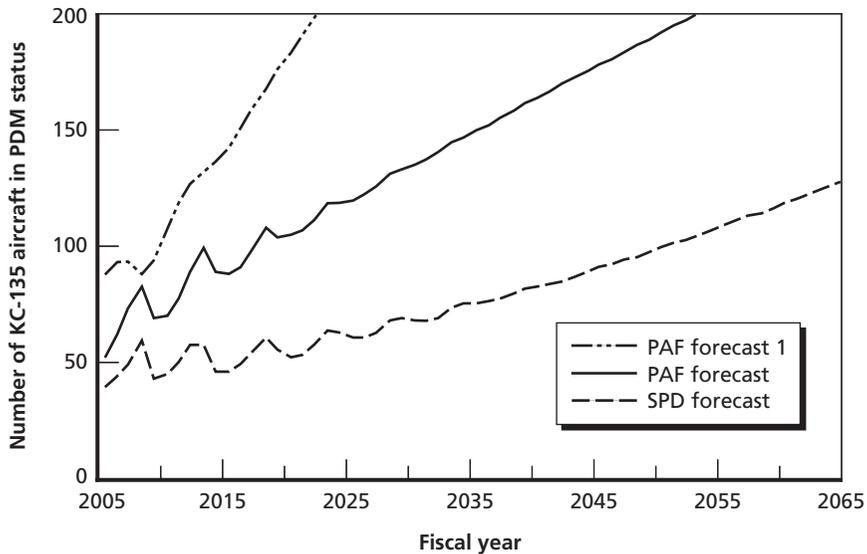
[5] The office's formal designation has recently been changed from the KC-135 SPD office to the 437th Tanker Sustainment Group (437 TSG). The forecast was very similar to that for the KC-135 ESLS but was based on more-recent decisions that eliminated some near-term tasks and postponed others.

Limitations of the PDMCAT Model

PDMCAT is a macro-level forecasting model. As with all forecasting models, it is sensitive to the accuracy of the factors used to generate the forecast. PDMCAT requires three critical factors: a forecast of future workloads, an estimate of the maximum labor application rate, and an estimate of the depot capacity.

Future PDM workloads are the subject of some debate. Pyles (2003) found a general cross-fleet pattern for PDM growth as fleets age and a significant second-order term related to age. An analysis focused solely on the KC-135—the KC-135 ESLS (Sperry et al., 2001)—also projected continued growth on KC-135 PDM workloads, although at a less pronounced rate than that found by Pyles. The 437 TSG workload forecast closely mirrors the ESLS forecast in terms of rate of growth, but projects fewer hours per PDM. The PDMCAT forecast of aircraft in work will vary depending on the workload forecast used. While workloads have grown in recent years, this is hardly conclusive evidence that the trend will continue into the future. Some argue that the workload growth will necessarily taper off as all or most of the key components on the KC-135 are repaired or replaced. Therefore, users of the PDMCAT model to forecast long-term trends in aircraft availability (20 or 30 years into the future) should periodically review and refine the available workload forecasts to reflect more-recent information that may reduce those differences in workload forecasts (see pp. 46–49).

An estimate of the maximum labor application rate (sometimes called the maximum hands-on burn rate), the rate at which labor can be applied to PDM workload, may change over time as process improvements, learning, and technology allow depot personnel to work more efficiently. As it becomes possible to apply more labor simultaneously to each aircraft, the PDM flow times will diminish. However, some changes in the underlying processes, such as subcontracting some tasks to outside entities, may reduce both the measured workload and the measured maximum labor application rate without necessarily reducing the flow time as the PDM process waits for the completion of subcontracted work. When workloads are contracted out or otherwise

moved from the formal PDM package, it is important to reestimate the labor application rate. As an estimate of depot capacity, the PDMCAT model measures depot capacity in docks—the number of aircraft that can receive work simultaneously at the maximum labor rate. The modeler has the option of entering a constant number of docks or of increasing the number of available docks over time. However, the PDMCAT model does not assess how the addition of docks may change the labor skill mix and affect the burn rate, nor does it consider how additional docks are added. That is, PDMCAT does not differentiate the addition of docks within an existing facility (by freeing up space currently occupied by other workloads) from the addition of docks by hiring contractors or by otherwise increasing physical capacity.

Although the underlying mathematics of the PDMCAT model support both lower- and upper-bound calculations on PDM throughput, the model produces only an estimate of the upper bound. Estimating the lower bound requires additional information about the imbalances across various stages of the PDM processes (i.e., the times and resources devoted to different PDM tasks) that would seldom be available to an external observer because of the competitive value of that information. In addition, we assume that PDM process managers will allocate their resources across those tasks to maintain a balanced production process, in which the average throughput rates at each production stage are roughly equal.

Conclusions

We were able to use the model to examine some important near- and long-term issues associated with the KC-135 fleet. While we were impressed with the model's existing capabilities, we have already begun to extend it to deal with multiple fleets using shared facilities, fleets with induction periods of less than a year, and modification workloads (see pp. 56–57).

With regard to the KC-135, we found that the shapes of the availability and cost forecasts did not grow in proportion to workloads, as

assumed in many studies.[6] Future studies forecasting PDM costs and aircraft availability may need to consider using PDMCAT or equivalent calculations to estimate how changing PDM workloads will affect fleets' budgets and availability (see pp. 53–55).

---

[6] The KC-135 Analysis of Alternatives study (Kennedy et al., 2006) is an exception. A version of PDMCAT was used to estimate the number of aircraft in PDM and modification status, and the PDM costs associated with several different workload forecasts.