Proceedings of the 2007 Winter Simulation Conference
S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, eds.

SIMULATION-BASED, ONTOLOGY DRIVEN RESOURCE PLAN DEVELOPMENT

Michael Graul
Perakath Benjamin
Arthur Keen

Knowledge Based Systems, Inc.
1408 University Drive East
College Station, TX 77840, U.S.A.

Frank Boydstun

3001 Staff Drive, Ste 2Y31
Oklahoma City Air Logistics Center
Tinker Air Force Base, OK, 73145-3025, U.S.A.

ABSTRACT

This paper describes the use of a planning ontology of the domain of Aircraft Maintenance, Repair and Overhaul [MRO] at a USAF depot to produce a discrete event simulation model of the aircraft ramp. The ramp is a very flexible, critical, shared resource for aircraft production. The "found work" variability inherent in the MRO process forces the ramp allocations and routing to be modified frequently. Unplanned aircraft moves are expensive and propagate queuing congestion and more moves. Therefore, it is necessary to develop a Rapid Ramp Reconfiguration Plan (RRRP). Ontology models allow the language of the planner to be harmonized with the language of the simulation analyst to speed model development. Simulation based planning is useful to mitigate the impact of "discovered work" by enabling the evaluation of and best selection from thousands of potential ramp resource allocation scenarios.

1 INTRODUCTION

This paper describes ongoing research and the application of that research into the use of ontology models to enhance simulation model development and configuration for use in MRO resource planning. The application domain, Aircraft MRO [Maintenance, Repair, and Overhaul] at a USAF maintenance depot is distinguished by a highly variable and constantly changing workload, making the planning and scheduling of work difficult. Further, the MRO domain has been hampered by the misapplication of manufacturing system terminology, philosophy, management practices, and information technologies to an inherently diagnostic type system. The gap between the manufacturing and MRO domains is very obvious in the technologies used to support resource planning. While the ontology of the manufacturing planner has received considerable attention since 1974, the MRO domain has suffered from the lack of a formal definition of the objects and relations—and therefore a distinctive set of terminology—that would allow planners to characterize and define the true nature of the system. Hence, technologies developed for planning in the MRO domain suffer from inconsistent and misleading terminology. This, coupled with the highly variable nature of MRO, has led the depot culture to develop decades of work arounds and tweaks to enable the business systems to support the shop floor. Most of these deltas between the "business school definition" and the actual need of the production shops result from the historic tension between production performance of the shop floor and business methods above the shop floor. The production organizations have developed practices to embellish and diminish the formal, impoverished business data to ensure success. So the depot systems become twisted artifacts of the culture. While traditional system modernization decisions have three sides, one that says "replace it because the technology is obsolete," a second that says "replace it because it will never work right," and the third that says "keep it because it has never been used right," the depot is unique in that all sides have an element of truth in their argument.

Ongoing research has shown that simulation models can be effectively leveraged by the MRO planner to analyze the impact of the inherent system variability—and provide robust resource configurations that mitigate the impact of the workload variation on schedule and cost. A simulation-based planning capability supported by an MRO domain ontology is a step toward establishing a correct fit for an MRO information technology capability and providing the planner with precise and coherent analyses rather than accounting system artifacts using a misapplied terminology set.

The first part of this paper describes the nature of the MRO domain and describes the core facets of MRO that make modeling its behavior essential. The second part of this paper discusses mapping these essential MRO facets into a formal definition, from the MRO planners perspective, of what exists in the MRO domain. Next we define a set of needs and requirements for system simulation constructs that, if developed, would enable computer modeling of the MRO environment and assist the

MRO production and resource planner by enabling the study of the system performance of various MRO processing scenarios over time. Finally, we describe an MRO ontology based simulation model that has been developed and is currently in use at the Oklahoma City Air Logistics Center (OC-ALC) for Ramp configuration planning. The subject matter of this paper is developed from the experience of performing simulation modeling at the OC-ALC at Tinker Air Force Base (AFB).

2 NATURE OF THE MRO DOMAIN

USAF maintenance depots perform the most radical and invasive maintenance, repair, overhaul, and modification tasks on Air Force weapon systems. Tinker is one of three Air Force maintenance depots. Tinker performs depot level maintenance on large body aircraft, jet engines, common aircraft commodities, and aircraft software.

The primary objective of the depot is to reverse the wear-out trend of the weapon system. This is normally achieved through MRO of the entire system and its subsystems, down to the individual nuts, bolts, and washers, as needed and as is economically reasonable. In addition, because of the expertise available in taking apart and putting together the weapon system, the depot is often tasked to modify the system, changing its configuration either to add to the functions to meet some need of the warfighter or to improve the reliability and/or maintainability of the system and thereby reduce the total cost of ownership.

This paper uses the depot workload of large aircraft in its discussion of the depot domain. Tinker performs Programmed Depot Maintenance (PDM) on large multi-engine airframes. The PDM consists of basically three phases, called Pre-Dock, In-Dock and Post-Dock. The Pre-Dock and Post-Dock phases are done mostly outside on the ramp. The exception is the stripping of the paint during Pre-Dock in the strip hangar, and the painting of the airframe during Post-Dock in the paint hangar. The In-Dock Phase is done inside a hangar. There are situations where multiple hangars are needed for the In-Dock work (Graul et. al 2002).

2.1 Depot Found Work

The single most significant difference between depot workload and manufacturing, and a difference likely true of *diagnostic systems* in general, is the nature of what we might call "found work." After the induction occurs, the system is opened up, investigated, and evaluated, and the detailed list of work really needed for that particular unit is determined. The work content continues to be "found" or identified as the system goes through the phases of depot maintenance. Analysis of history in the depot indicates that anywhere from 20% to 60% of the man-hours on a system will be identified and defined after the induction of the system into the production shop. How will this level of variance play havoc with the budgets and schedules planned so carefully and negotiated in prior years to support the ongoing needs of the warfighter?

2.1.1 Nature of Found Work

There are three characteristics that describe the nature of found work. The first characteristic is that a task of found work is not like any other task of found work. Each instance is unique in terms of its resource requirements. The easiest example is corrosion removal. While the corrosion may be on the same part in the same location, each is corroded more or less severely. The removal of each instance will take more or less time, and maybe more or less skill, because of the particular instance of that corrosion. While this may seem somewhat trivial to the reader, corrosion removal on an aging weapon system can be significant. Consider a planned level of 3,000 hours of corrosion removal for each of 50 aircraft, which is actually very low for a large body aircraft, especially the aging airframes of the tanker fleet. This year plan requires 150,000 hours, or corrosion work for about 75 people per year at a yield of 2,000 work hours per year per person. The historical variance of 20% to 60% gives a probability low of 180,000 hours, up to a high of 240,000 hours. The 75 employees for corrosion work will be working between 400 and 1200 hours of overtime in one year! This unique aspect of found work by instance is consistent for hydraulic, sheet metal, electrical, landing gear, etc. Obviously, the first approximation of improvement is to calculate for at least the low percentage variability, in this case plan on 180,000 hours per year, with 150,000 known precisely and another 30,000 to be determined as each airframe is evaluated. There is still a probability of up to another 60,000 hours, or 30% growth.

Another type of found work is the "leaker." After the airframe has been completely overhauled, the process of putting the airframe back in service after performing a functional check flight begins. This is the post-dock phase of the PDM. The fuel system of the aircraft is fully restored and must be tested. The aircraft is moved to one of the fuel/de-fuel pits on the ramp, loaded with fuel, and moved to another ramp spot. The fuel system is pressurized for twenty-four hours and then inspected for leaks. Since the fuel systems have been dry several months and the airframe has moved many times, finding leaks where no work has been done and where no leak was noticed on arrival is normal. Fixing the leaks usually requires removing the fuel by another trip to the fuel pit. Then the leaks are repaired and another pressure test is performed. It is quite normal to find additional leaks and repeat the cycle. In a recent data mining analysis of ramp movement, aircraft were observed moving to the fuel pit repeatedly. One aircraft's data indicated over thirty cycles of this movement. Expertise was sought in the engineering office of a former post-dock production chief who stated that it was not only accurate but also remembered the aircraft's tail number and the disruption the aircraft was to their schedule. He verified that what seems like very unusual numbers of cycles were normal PDM situations because those aircraft are "leakers." This was important to the analysis because a "leaker" begins a spike effect of congestion and rippled movement on the ramp. Leakers become like pin-balls that never leave the ramp, bouncing around and around. This found work causes havoc in post-dock, which is the most constrained portion of the PDM Schedule because all the impacts of all the other found work have consumed all the contingency options for recovery. The only option left in Post-Dock is overtime, working around the clock to meet the production dates. Each instance of restoring the fuel system on an aircraft is unique including the drastic impact to all other aircraft caused by "leakers."

The first characteristic, the unique nature of each instance of found work, drives the second characteristic of found work: the organizational level of experience for that task. The experience level can be qualified in three grades: mature, developing, or new. Found work can be a mature task that has been performed several times, is very well known, and is easily quantified, once discovered. Found work can be a developing task that has been performed a few times but that has not been performed enough to be completely and definitely known. Found work can also be a first time new failure. A new failure is not covered by the technical orders and repair procedures and, therefore, requires engineering approval from the aircraft management office.

The third characteristic describing the nature of found work is the principle that *work will always be found*, including new failures. It is logically impossible to determine everything that is actually wrong with a system until that system is "completely" evaluated. In many remanufacturing situations, this unknown has been bypassed by simply making a 100% disassembly and 100% replacement of parts and only reusing the core itself, which is either accepted or rejected at the collection point. This method of 100% invasive disassembly is impractical for large body jets. The nature of found work involves unique tasks even when the tasks are defined similarly. Each task has a level of experience associated with it, or a learning curve context. This means that in depot MRO, each unit of the workload is entirely unique from every other unit inducted and worked.

2.1.2 Impact of Found Work

From an industrial engineering perspective, the primary objects impacted are the planned resource requirement and schedule. The previously mentioned statistic—the historical variance of 20% to 60% of the man-hours being found work—is only for the impact on the direct labor resource. This level of task variance causes a corresponding ripple (or convulsion) variance in the other resources. It includes variability in the indirect material, direct material, back-shop workload, engineering support, and indirect hours attributed to supply, scheduling, and planning. The facilities, the equipment, and even the people, all to a certain level of duration, can simply be put on overtime and run 24 hours a day, 7 days a week. However, parts cannot be put on overtime. In depot MRO, the critical resource is always parts. This is especially true for the impact of new failures. In some of the aircraft fleets worked at the OC-ALC, aircraft parts are routinely replaced that, according to the original airframe useful life design, did not ever need to be inspected! The ripple effect of found work variability on the resource requirement, when compared to new manufacture variability, is like a tsunami from the ocean compared to a pebble on a pond.

2.2 Salvage of Parts

One of the other unique characteristics of the depot is the availability of parts from inducted systems. Often referred to as cannibalization or "canning" and sometimes referred to as "rob-back," it is the acquisition of parts from an inducted work-piece instead of from prescribed supply or back-shop sources. The depot culture has various degrees of overlap and distinction of these words from place to place, and they serve as another example of the unique characteristic of the depot workload. This practice has often meant the difference between success and failure in delivering a reliable combat ready system to the war fighter. Canning negatively affects accurate supply requirements: what is canned today will not be bought tomorrow if there is no record of its demand in the supply system. Theoretically, then, it is possible to have fewer of a certain critical part in the inventory than there are weapon systems that require that part. This is because there is a work-in-process (WIP) of non-available weapon systems at the depot and some of them have been canned. The root cause for canning has two components. One is that the part requirement was not known until the work was found. The second component is the lead-time for weapon system parts. They are not readily available at the local hardware store, and the procurement process is significantly more involved if the part is not already covered by an existing contract.

One of the strategies for dealing with the issue of "lead time to acquire parts" is to keep condemned parts. While the part may originally be condemned for cost of repair versus buy new, when production is at a work stoppage and the field needs critical components, fixing one part quickly and expensively is preferred to buying a new one slowly and cheaply. Engineering may be required to again look closely at tolerance variances or approve new repair procedures, but the product goes out the gate.

2.3 Depot Transformation

The large complex array of processes needed to perform depot MRO began during WWII, which also gave birth to many other types of large-scale operations and their management tools and methods. Over the decades, the depot has been extremely successful in providing the frontline defenders of our freedoms, those devoted souls who live at the tip of the spear, with reliable weapon systems. However, while the rest of the world has been driven by global competition to improve their processes, the depots have been driven by other forces to largely preserve their system of weapon systems MRO. In recent years, there has been an increasing effort and determination—from the Federal Government down to the shop floor—to apply the cost cutting and time saving methods developed in the private industry to the depots. The objective is to reduce the cost of depot MRO and harvest that benefit for the sake of the war fighter as well as the taxpayer. It is a very large undertaking in at least three aspects. One is the magnitude of the organization being changed, some 7 million man-hours per year. The second aspect is the magnitude of change attempted: the entire organization and every shop will be drastically affected by the complete rearrangement of the shop infrastructure, layout, and processes. The third aspect is described by the analogy "We want to ride the bicycle while we fix it!" Transformation of the depot MRO will not include shutdown for any duration that causes a loss in the bottom line production! The Oklahoma City Air Logistics Center's 76th Maintenance Wing will meet all its production commitments while transforming its production processes to globally competitive methods and practices.

3 TOWARDS AN MRO DOMAIN ONTOLOGY

What are some of the characteristics of a Maintenance, Repair and Overhaul Planning Ontology (MROPO)?

We all know that some of the key pieces in MROPO include Organizations, Assets, States, Processes, Resources, and Constraints. Organizations are important to model because they are the key actors in the MRO that control budgets, schedules, resources, assets, and the constraints. In Figure 1 is a beginning meta-ontology,

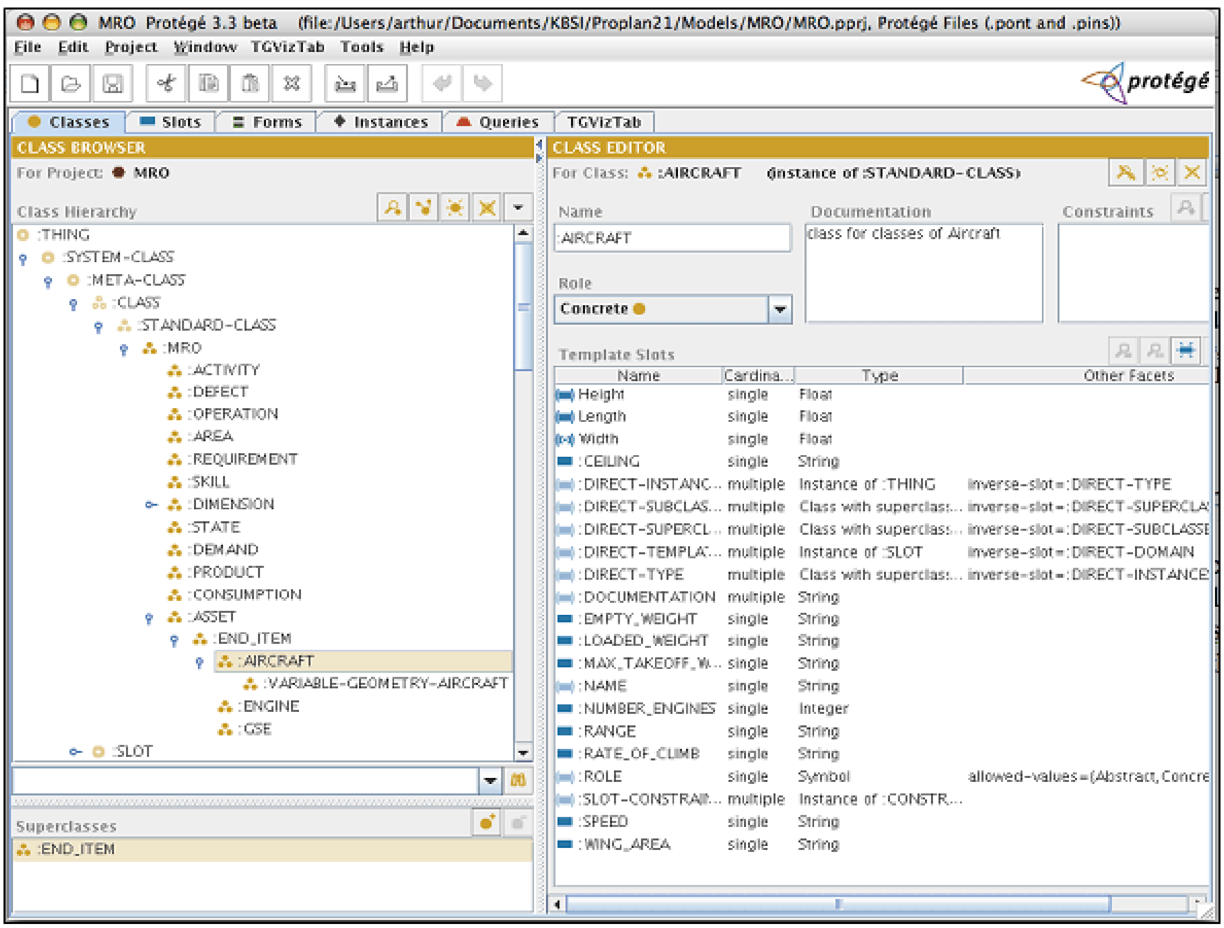

Figure 1: Partial MRO Meta-ontology Model

using the Protégé tool [http://protege.stanford.edu], of concepts that we have captured from domain experts, which include Activity, Defect, Operation, Area, etc. These models, when fully developed, can be leveraged by diverse projects that are well beyond the scope of MROPO.

The current analytical semantic mismatch between the state-of-the-art enterprise manufacturing systems and the real world of MROPO causes a flow of "near misses" to ripple across the enterprise. The compounded ripple error is analogous to impedance mismatch between electrical components. The electrical result is excessive power consumption, heat, and failure of components. The enterprise result is excessive resource consumption, schedule trauma, and materiel shortages. It also is analogous to the butterfly effect of chaos theory. Formalizing this MROPO knowledge in this ontology will not only have immediate benefits for ramp configuration simulation, but also form the foundation for addressing the following issues in today's MRO with follow-on next generation planning capabilities:

- Aging and retiring workforce with no formal long-term knowledge management;
- "Data Scrubbing" replaced with accurate data, information, and knowledge interchange between applications and individuals ;
- Formal and informal training deficiency replaced with realistic immersive training for planners;
- "Start from scratch" approach to resource planning replaced by automated "re-use" of MRO planning knowledge to speed up planning and re-planning;
- Collaboration tools across the enterprise for individuals involved in MRO;
- Current imprecise impact analysis causing excessive budgeting and schedule allotment replaced with accurate, quick decision support tools; and
- Imprecise knowledge of the informal organization / emergent behavior of the MRO organization is eliminated and predictable behavior established.

The advantage of the ontology modeling approach is the ability to re-use models to drive multiple analytical techniques, capture data in a common representation, and to share results. A major aspect of the MRO ontology is a representation for classes of assets as well as the assets themselves.

For example, the class B-52H is a subclass of B-52, which in turn is a subclass of bomber, which is a subclass of aircraft. The class B-52H has properties, for example the class has 8 engines. These properties are known as class properties because they hold for all instances of the class, so rather than having to specify them redundantly on each instance; it is far more efficient to specify them on the class. Figure 2 illustrates a characterization of the Asset type using the Protégé tool [http://protege.stanford.edu].

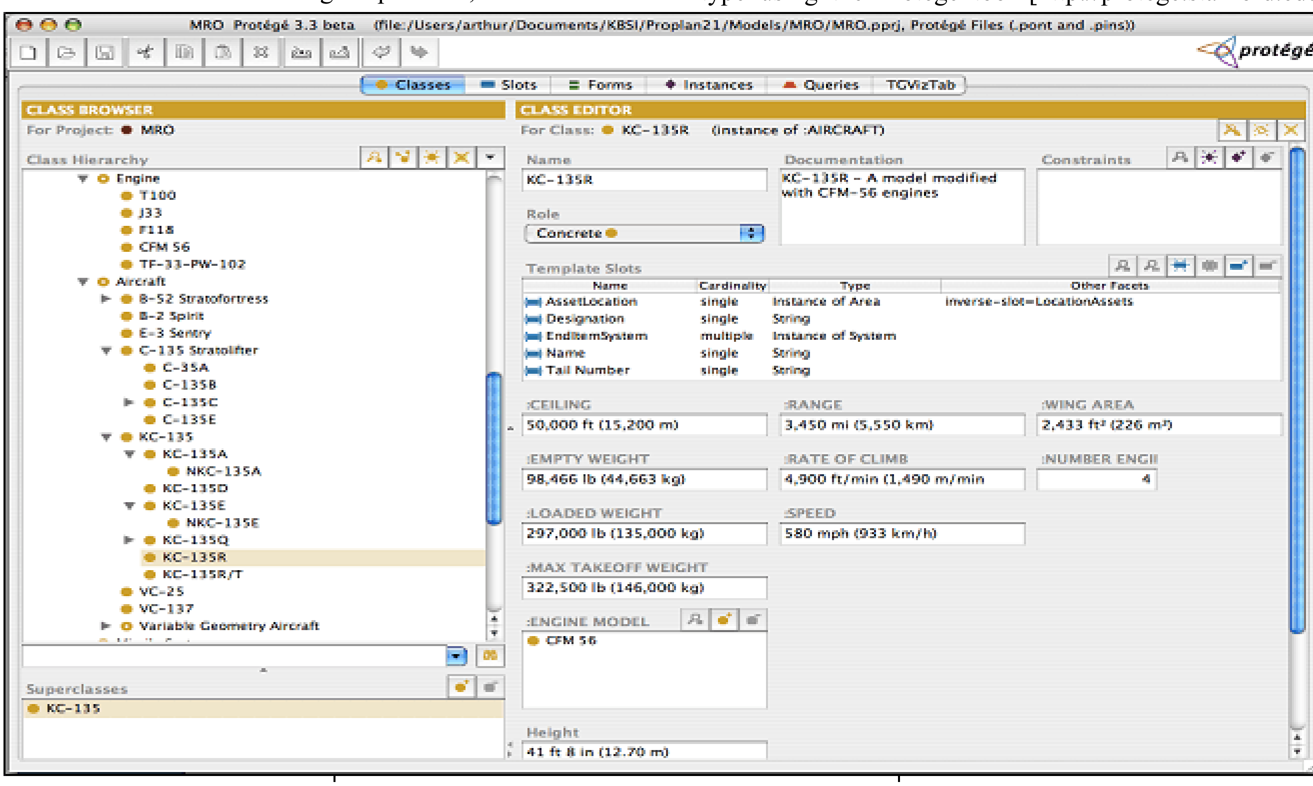

Figure 2: Partial MRO Asset Ontology Model

The resulting MRO ontology is large, but captures a definition of MRO that has never formally been documented. The ontology is a living document, and its representation can be used to drive not only simulation models but a variety of MRO planning and analysis tools and related business applications. The ontology model enables a formal language within a consistent framework and, most importantly, a distinctive definition of the system of systems that is MRO.

Another characteristic modeled in the ontology is a constraint associated with a resource. One of the four classical industrial engineering resources is facility, and there are facility space constraints. In the MRO world of large body jet aircraft, the airframe is put in a hangar dock as a facility. In addition, because work is done inside the airframe, the airframe becomes the facility in which the work is done. Nested facility constraints exist. So if the work is to be done in the cockpit on the avionics, there is only so much space for technicians, equipment, and parts. In one airframe organization, the planners have defined the term loading factor for this characteristic. The entire airframe, inside and outside, is broken into areas that are numbered and each area has a loading factor which is an integer of the number of technicians that can work in that area.

The largest shared resource in Tinker aircraft depot production is the facility resource of the aircraft ramp, which at Tinker is in excess of 80 acres of concrete. The aircraft PDM phases were mentioned earlier. These Phases and associated facility requirements are the result of decades of practical experience using mostly hangars that were built in the 1940's through 1960's. For the ramp, each area of the ramp has characteristics. Some were mentioned at the outset of this paper, such as aircraft weight versus concrete thickness. A fully fueled B-52H bomber cannot be parked just anywhere in the 80 acres. Another characteristic that is readily known in production is the constraint of the number and type of fuel pits where aircraft are fueled and de-fueled. Some fuel pits are new and efficient and some pits are old and unreliable. Fuel trucks can fuel and de-fuel aircraft in locations other than the fuel pits, but this method is more costly in time and money. The earlier discussion of the "leakers" as an unpredictable "found work" drives this resource constraint. "Leakers" can be modeled as a type of fuel system restoration work. The airframe either gets the "normal" fuel system restoration or the "leaker" fuel system restoration. Fuel system restoration happens in the Post-Dock phase, which is in the PDM work package, which is in the aircraft workload along with the other two workloads of modifications and Unscheduled Depot Level Maintenance (UDLM). A third constraint of the ramp, the final one mentioned in this paper, is the "run spots," which are locations where jet engines can be run for testing and tuning after re-installation in preparation for the final check flight. There are "initial engine runs" which simply involve getting the engine to fire up and test the rudimentary controls, and there are "full power engine runs" where the full spectrum of engine performance and control is validated according to specifications. Initial engine runs can be done with portable "blast fences" which deflect the propulsion exhaust of the jet engine upward in a safe manner. Full power engine runs require permanent blast fences that are built into the ground, especially the B-1 Bomber with its afterburner jet engines and the B-52H with its eight jet engines.

One of the objectives of this project is to advise the Maintenance Operations Control Center (MOCC) on how to configure the ramp functions according to the next few days/weeks of work to be done on the ramp due to the impact of found work on the work schedules of the aircraft in work. As aircraft processes have been leaned and production times reduced, the Air Force has realized a significant benefit in reduced aircraft in possession at the depots. Keeping the aircraft on schedule during the latter portions of flow, after all the impacts of found work, is one of the critical features and challenges of maintaining that benefit. These latter portions of the MRO flow are when the fuel is put back on the aircraft and the engine runs are performed. In other words, there is no slack time for waiting while ramp spots become available. When the jet needs a ramp spot, it is needed now!

4 IMPLICATIONS FOR SIMULATION MODELING

There are two specific challenges to be discussed that are faced by planners in their attempt to model and analyze the MRO domain for operational improvements (as discussed in the previous section) and relate these concepts to simulation model-specific needs. The first is a permanent characteristic of the variability inherent to MRO. The second is the current bow-wave of MRO transformation from traditional weapon system MRO through transformation to world-class lean MRO. So the technology will not only simulate the impact of variability, but also provide simulation capability that enables a transformation team to study the impacts of various transition strategies prior to initiating facility modifications and possibly compromising production opportunities. Normally, a system plagued by extreme variability in its key performance indicators is indicative of decision policies that are ineffective in controlling the influences of the transition step changes during the transformation. Hence, a useful way of modeling the MRO system is to isolate the production control logic from the production process logic that it supports. In doing so, the modeler can more easily specify the decision logic necessary to deal with cases of found work, for example, without worrying about how the decision logic will be executed by the simulation engine. This in fact largely simulates the transformation effect, changing more the way workload is managed and executed than the actual workload and resources themselves. The question then becomes, how must the underlying simulation engine operate in order to integrate the two components back together?

The following section outlines the simulation technology requirements to addressing these system simulation needs and dealing specifically with the challenges described previously for an MRO system undergoing transition.

5 OVERVIEW OF ONTOLOGY DRIVEN RESOURCE PLAN DEVELOPMENT USING RAMPPLAN

The ability to accurately represent OC-ALC depot processes in the MRO ontology will lead to the rapid deployment of *cost-effective* and *high-performance depot processes*. Transformation is inevitable, but the transition path itself must be selected to allow for the minimal disruption of resources. Moreover, the transient nature of MRO necessitates the study of variability and its impact throughout the transformation as the planned transitions proceed. The simulation-based ramp resource assessment tool requirements highlighted previously were used to develop a technology termed "RampPlan." The key innovations provided through RampPlan are described in this section.

5.1 Transformation Modeling

The following paragraphs describe some of the challenges of the transformation of the depot that drive simulation modeling characteristics. As the transformation steps forward, there are numerous transition events of resources and control policies to be evaluated (Benjamin et. al 2005).

5.1.1 Future State Modeling

RampPlan allows the modeler to define and project forward in time critical element states that characterize when resources will be available for use. For example, during the transformation process, particular facility resources are taken off-line [unavailable] for repair or modification, and when the facility modification is completed, the facility is brought back on-line and made available. It is essential that the planner be able to model these resource states, but also to run the projection to a particular state in time—for instance, just as a resource is taken off-line. At that point the model can be saved—at that particular world state—and the planner can modify the model and restart from that "world" state and perform another forward projection. This style of modeling and experimentation allows the planner to create a comprehensive set of projections that cover the transitions for key resources—each projection is a world state that can be used to start a new "thread" of experimentation for generation of possible future states.

5.1.2 Transition Process Modeling

RampPlan allows the planner to model and simulate different transition activity types. As discussed above, the tool must simulate different types of dock modification activities. Once a dock modification is complete, it becomes active and is available for use in the MRO process; and further, after dock modifications are completed, the dock may open with a different "configuration" (i.e., different physical characteristics of the dock and the type of work that can be performed in that dock space).

5.1.3 Flexible, Multi-Site Work Policy Modeling

RampPlan is capable of allowing the planner to simultaneously model and simulate single-site phase maintenance (Cellular) and multi-site maintenance activities (Phase). For example, under the current policy at Tinker, aircraft move to different facilities for each phase of the repair process: Pre-Dock, performed on the ramp; In-Dock PDM, performed in designated PDM docks; and Post-Dock, also performed on the ramp. This requires multiple moves that require time and money. What if there was a "One-Stop Dock" in which all maintenance could be performed? Where is the cut-off on the benefit from total to partial One-Stop Dock capability? It is being considered under the lean maintenance paradigm that all maintenance activities, except small portions of Pre and Post Dock activities, will be performed in a single dock?

5.1.4 Modeling Complex Constraints

RampPlan provides the capability of modeling the variety of physical and logical constraints imposed by the requirements of complex and dynamic MRO activities. For example, the tool must provide the capability to simulate the physical constraints on the accessibility of aircraft to bay doors. In some cases, an aircraft must move through multiple docks to gain access to a door. In Figure 3, the building, built in 1942, has nine docks and two doors. Only three airframes are next to a door. All the others are blocked by at least one airframe between them and the door. The simulation must enable the creation of such "Blocking Constraints" that would allow end users to make policy decisions about how to adequately address the depot performance limitations imposed by these constraints. For example, representation of the Blocking Constraint will enable answering questions such as "Do all the airframes move when the first one is ready or when the last one is ready to move?"

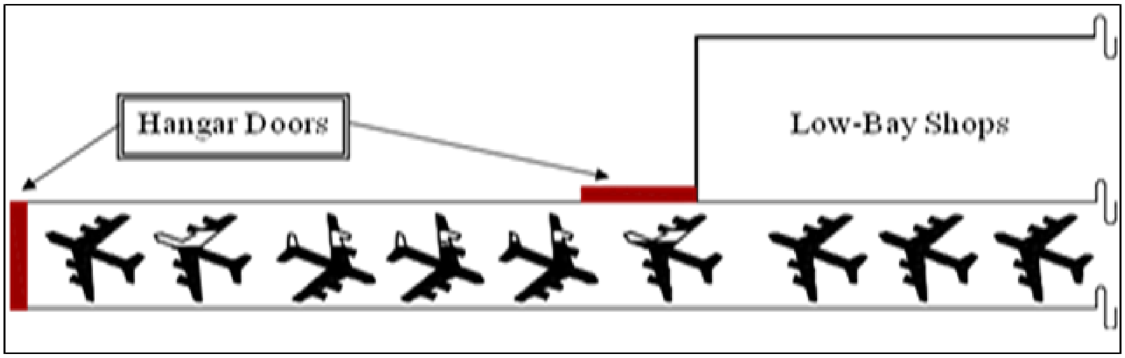

Figure 3: Hangar with Nine Docks and Two Doors

5.2 Ramp Configuration Simulation – RampPlan

This section describes some of the more immediate challenges for existing workload and configuration of the ramp to be met using the RampPlan simulation tool and establishing a Ramp Resource Configuration Plan (RRCP).

5.2.1 Configuring the Ramp for Leakers

Configuration policy today for the ramp tries to keep only fueling activities on the fuel pits and engine run activities on the permanent engine run spots. Under what conditions can these policies be modified to reduce the impact of leakers on the overall aircraft production schedule and resource requirement? The logic of the simulation model must include a policy for what work can be done where on the ramp with triggers for transitioning to a different policy based on levels of found work. How many leakers are there?

5.2.2 Multiplicity of Dynamic Workload Types

RampPlan must be able to evaluate a very dynamic workload, including the requirements of multiple aircraft types (e.g. KC-135, B-1, B-52, E-3, etc.) for PDM, with their unique and different maintenance schedules (Schedule for B-1. vs. Schedule for KC-135, etc.) and "drop-ins" called Unscheduled Depot Level Maintenance (UDLM). Each of these workloads has the depot characteristic of found work discussed in the first section of the paper, so the simulation tool must also be able to accommodate *work variability* requirements—that is, the ability to model and simulate work *discovered* or *found* on the aircraft during the maintenance cycle.

5.2.3 Ramp Transformation

The RRCP must accommodate the changes in dock capability derived from aircraft production facilities transformation. If a "One-Stop Dock" does actually happen, or a "Three-Bay Hangar" is built, how will those facility modifications/constructions impact the ramp configuration and utilization?

5.3 Advanced Experiment Management

RampPlan must provide sophisticated simulation experiment management capabilities. For example, the tool should have the ability to save experiments, to re-initiate/recreate the state from which the simulation was run, and to re-run the simulation using the original status information. These tools should also have the capability to save simulated (possible) world states and load them as the "current" world states.

6 CONCLUSION

This paper has outlined a number of key facets that delineate depot or MRO production from manufacturing production. In doing so, our focus has been on describing the problem domain well enough to understand what facets of MRO have the greatest impact on production and, therefore, should be included in planning models used for designing and studying the ramp resource configurations necessary to achieve the end state of the increased throughput in a phased, cost effective manner. Finally the value of using ontology models for performing simulation-based plan analysis was described.

REFERENCES

Benjamin, P., F. Boydstun, M. Graul, M. Painter, 2005. Toward Effective Depot Transformation: Leveraging Simulation to Enhance Transition Planning. *WSC '05: Proceedings of the 36th Conference on Winter Simulation, Orlando, Florida.*

Boydstun, F., M. Graul, P. Benjamin, and M. Painter, 2002. New perspectives towards modeling depot MRO. *WSC '02: Proceedings of the 34th Conference on Winter Simulation,* San Diego, California, 738-746.

DR. MICHAEL GRAUL is a senior research scientist at KBSI, holds a Ph.D. in Industrial Engineering from Texas

A&M University. Prior to joining KBSI, Dr. Graul worked as an industrial engineer, systems engineer, and lead statistician for several systems design and development projects within the geophysical, medical, and manufacturing domains. Since joining KBSI in 1992, he has been responsible for the design and development of a class of reconfigurable, self-maintaining simulation modeling technologies. Dr. Graul has initiated over 70 grass-roots BPR projects and has trained over 300 people in the IDEF family of system modeling methods.

DR. PERAKATH BENJAMIN is a Vice President at KBSI, manages and directs KBSI's R&D activities. He has over 16 years of professional experience in systems analysis, design, development, testing, documentation, deployment, and training. Dr. Benjamin has a Ph.D. in Industrial Engineering from Texas A&M University. Dr. Benjamin has been responsible for the development of process modeling, software development planning, and simulation generation tools that are being applied extensively throughout industry and government. At KBSI, Dr. Benjamin was the principal architect on an NSF project to develop intelligent support for simulation modeling that led to the development of the commercial simulation model design tool, PROSIM®.

DR. ARTHUR KEEN is a senior research scientist at KBSI, holds a Ph.D. in Industrial Engineering from Texas A&M University. He has over 20 years experience in knowledge representation and reasoning applied to intelligence, security, manufacturing, planning, and scheduling.

FRANK BOYDSTUN is an industrial engineer at the Oklahoma City Air Logistics Center, has managed process improvement of depot production processes for the last ten years. Ongoing process improvement efforts include simulation modeling, process knowledge capture, systems development/upgrading, numerous implementation projects, and sponsoring research by the Air Force into the nature of depot operations and the application of technologies to gain depot improvements. He holds bachelor degrees in English and Mechanical Engineering from Oklahoma State University