Proceedings of the 2002 Winter Simulation Conference
E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, eds.

NEW PERSPECTIVES TOWARDS MODELING DEPOT MRO

Frank Boydstun

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

Michael Graul
Perakath Benjamin
Michael Painter

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

ABSTRACT

There are subtle, and yet critical and unique differences that distinguish the depot maintenance, repair, and overhaul (MRO) domain from production manufacturing. These differences motivate the need for more efficient ways to capture the essence of the depot MRO domain dynamics. The authors provide an informal characterization of the depot MRO by highlighting some of the major differences. Along with this characterization, they propose a set of principles governing the physics of depot MRO operation. Finally, they describe the nature of idealizations needed to model and simulate this domain and a vision for future technologies that could more adequately and directly address these needs.

1 INTRODUCTION

The primary purpose of this paper is to give simulation analysts who are unfamiliar with the depot domain a foundation upon which to build successful modeling and simulation efforts. There are subtle, yet critical differences, that distinguish the depot domain from the manufacturing environment that often change how idealizations should be developed. To further aid those involved in simulating the depot domain, we discuss how these differences motivate changes in how simulation modeling will be accomplished in the future to more efficiently capture the essence of depot domain dynamics.

We begin by providing an informal characterization of the depot maintenance, repair, and overhaul (MRO) highlighting some of the critical differences between it and production manufacturing. We will use the terms “Depot MRO” and “Depot” interchangeably. Along with this characterization, we propose a set of principles governing the physics of the depot MRO domain operation. Finally, we describe the nature of idealizations uniquely suited for depot MRO modeling and simulation and describe a vision for future technologies that directly address these needs.

2 THE DEPOT DOMAIN – AN INFORMAL CHARACTERIZATION

We use the term *Depot MRO* in reference to the depot maintenance, repair, and overhaul (MRO) activity. This activity is part of the logistics life cycle, and is also included in the concept of sustainment, where sustainment means prolonging the effective conduct of warfighting tasks through logistics support. In the context of sustainment and logistics, the term depot MRO refers to the most radical activities for prolonging the useful life of the end items serviced. In addition, the more complex the system being worked in depot MRO, the more the following characterizations, concepts, and principles will apply (e.g., depot MRO of a large four-engine jet aircraft).

The purpose of depot-level MRO is to reverse the aging and wear-out process of the system. Early in the operational phase of the system life cycle, failure rates may be high, commonly called infant mortality, as the weaknesses of the engineering design and the production process are made evident through early failures. As the system is repaired, sometimes modified, failure rates approach a constant level. As the system nears the end of its design life cycle, failures increase at an accelerating pace. That is, both the number and type of failures increase. There are several levels of MRO activity, most aim at simply returning the system to operational status. Depot-level MRO is aimed at turning back the clock by not only doing routine long cycle maintenance and overhaul, but also making repairs and implementing modifications to the system that drive the operational reliability of the system back to a higher value, as illustrated in Figure 1. MRO of aging systems is a perpetual battle between the desire to bring the end item back to a *near new* state and the need to mitigate maintenance costs.

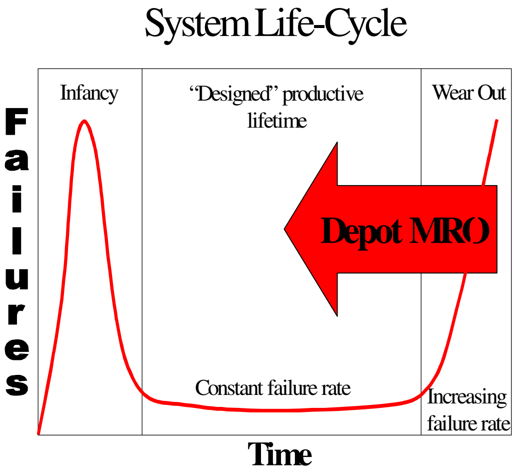

Figure 1: Depot MRO Purpose

2.1 A Contrast Between Production Manufacturing and Depot MRO

Key differences in the depot MRO domain that distinguish it from the production manufacturing domain include:

- The degree to which work content is known and the point at which it is identified.
- The degree to which the system relies on shop floor knowledge and adaptive flexibility.
- The presence of high variability in processing times.

The primary feature that distinguishes the depot domain from manufacturing is the inherent additive nature of work content variability. In the manufacturing domain, all work content is determined by design before the start of production. Consequently, the resources (e.g., skills, materials, equipment) can be found deterministically before workload is inducted or released to the shop floor. In the depot domain, 20-60% of the work [to be performed] is discovered during the repair cycle. Work content definition must take place within the context of ongoing repair operations. This contrast is notionally represented in Figure 2.

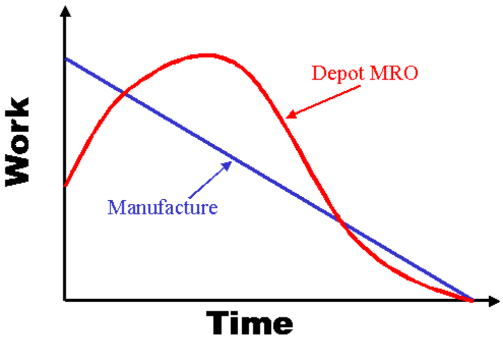

Figure 2: Depot Work Content Per Item

Work content is only partially known when the end item is first inducted and often continues to be redefined throughout the end item's depot sojourn. Initially, end items are inducted to the shop with a list of the known work that must be achieved. In general, the first set of tasks to be accomplished involves inspecting, diagnosing, or troubleshooting the system to find the additional requirement for maintenance, repair, or overhaul. Work needs discovered during either formal or informal inspections are referred to by terms like *over-and-above*, *unpredicted*, *stumble-ons*, and so forth.

Work uncovered during the depot MRO process may be classified into one of three categories. The first category will be called, *routine work*. An inspection may reveal that something not originally called out in the work plan for a particular aircraft is now required. Routine work has a reasonably well-known work content and well-defined work plan. The skills required are reasonably well developed, and the characterization of how long the task will take is fairly stable, although there are certainly cases where task duration is highly variable (e.g., removing corrosion). The second category of work will be referred to as *sometimes work*. This kind of work has been discovered and done before, but not enough for the work content and plan to be considered mature. More experienced workers are needed to get the work done in a reasonable time with consistent quality. The third category of work involves *new failures*. These are failures that require the engineering community to provide repair instructions. New failures can drive the issuance of a fleet grounding instructions and a Time Compliance Technical Order (TCTO)[this is Air Force terminology], that mandates inspecting the fleet for certain failure conditions within a short time period.

The implications of work content variability on the depot MRO system, as a whole, are dramatic in scope, impacting facility, labor quantity, skill mix, tooling, ground support equipment, parts, in addition to obvious ripples through planning, scheduling and engineering. The key to efficient depot operation is an effective diagnostic system and explicit mechanisms to leverage shop floor knowledge through immediate communication to affected support functions, to enhance responsiveness. This goes counter to conventional approaches, where the processes, systems, and methods successfully applied in the manufacturing domain are assumed to work equally well, and without modification, for depot operations. These systems in general are not designed to respond to additional state of the system knowledge originating on the shop floor, since no such event ever occurs in the manufacturing domain.

Since shop floor knowledge and discovery are at the heart of depot system operation, the shop floor is the fountain of knowledge for the true state of not only the work content, but also the system as a functional component of the fleet or force. Any dispatch and control technique for the depot shop floor must gather more than the status of execution of known work content, and must include not only the additive nature of the discovered routine work, but also the additive nature of the discovery techniques and the engineered repair instructions arising from new failures. This obviously goes toward the principles that must govern any information system used directly with the shop floor, and that is another topic. But let it suffice as a passing observation, that the shop floor information system must be as much a gatherer of new knowledge as it is a status tool against a plan.

A more efficient process would move traditional planning, scheduling, and materials management decisions closer to the shop floor — the point of discovery and primary locus of depot MRO knowledge — where information distortion and delays in the management and supply chain echelons can be minimized. Unlike in the manufacturing domain where the "plan is everything," in depot MRO [as General Eisenhower once characterized] "the plan is nothing – planning is everything."

2.2 Depot Involves Dealing with Increases in New Failures

Let us discuss some numbers to illustrate the concepts discussed thus far. The U.S. Air Force's KC-135 and E-3 platforms will serve as examples.

Over the last decade, the amount of work required to return KC-135 aircraft to certified operational status has increased by 70%. The KC-135 is the venerable 40's design, 50's manufactured, fuel tanker for midair refueling in the U.S. Air Force. During the year 2000, the KC-135 production line generated over 200 engineering support requests per month. These requests involved new failures discovered through inspection. Ten percent (10%) of requests for engineering support (i.e., new failures) originated during the last thirty days of the depot overhaul period when the aircraft is supposedly just about complete and ready to go!

There are eight possible outcomes of an inspection:

- Everything is okay.
- Direct labor is needed to tighten or align.
- Direct labor and indirect material are needed to paint, caulk, or refasten.
- Direct labor and direct part to replace a condemned part,
- Direct labor and direct part to replace a part carcass that is routed to a backshop for repair and return as a serviceable unit.
- New failure with request for engineering disposition that is somewhat similar to routine work.
- New failure with request for engineering disposition that is an emergency.
- Another inspection is needed.

Depot MRO inspections have a recursive nature to them, so you cannot be really sure all the inspections are finished until you have finished the last one. In 1994, the Air Force E-3 AWACS depot MRO package had 1008 inspections. At that time, the E-3 package was the smallest MRO package of the aircraft worked at Tinker. With eight possible outcomes for each inspection, there are 8^{1008} possible outcomes, or roughly 10^{903} possibilities. Let's pretend that the permutation space could be trimmed through various observations such that 99% of the space is eliminated. The remaining 1% is still on the order of 10^9 possible outcomes per aircraft. At the time, Tinker was doing some 80 aircraft per year!

2.3 Time to Repair – "The Mean is Not the Answer"

Within the depot domain the time to finish MRO is a function of the degree to which the system or item is discrepant (e.g., removing corrosion). Unfortunately, diagnosis of the system state, while performed by the technician, is rarely documented for use in calculating the [estimate of] time for repair. Consequently, the data that exists in the production systems provides more of an accounting view of repair and is relatively useless in describing the true time for repair. Moreover, a simulation analyst who attempts to utilize the time to repair data available in the production systems is likely to find that the "repair time" is a mean attempting to characterize a random variable with a wide variance [covering the span of part discrepancies and repair situations]. The solution is to collect repair times on the population of parts that exhibit a particular condition of failure – and group these items by the degree to which they are "broken". In lieu of this information it is interesting to note that many repair systems exhibit repair times that follow a log-normal distribution (Blanchard 1998); our data, from repair time estimates in the MRO of avionics items, seems to corroborate this evidence. In essence, without proper classification, the use of means for time to repair result in a poor estimate – and great care should be taken to properly classify the work into at least "light", "medium", and "hard-broke" categories prior to populating repair times in a given model.

Production manufacturing processes are designed to generate production durations with a mean processing time surrounded by a very tight variance. Repair environments are burdened by a much larger variance making the use of mean values suspect for either simulation or production scheduling and control. The larger variance is most often traced to information collection at a gross level (see Figure 3). Nevertheless, a form of mean processing time, called a *labor standard*, is routinely used for exactly these purposes.

Labor standards used in the depot are often *engineered* following general industrial engineering practices. The data derived through this analysis is then used to perform should cost analysis, thereby helping depot management

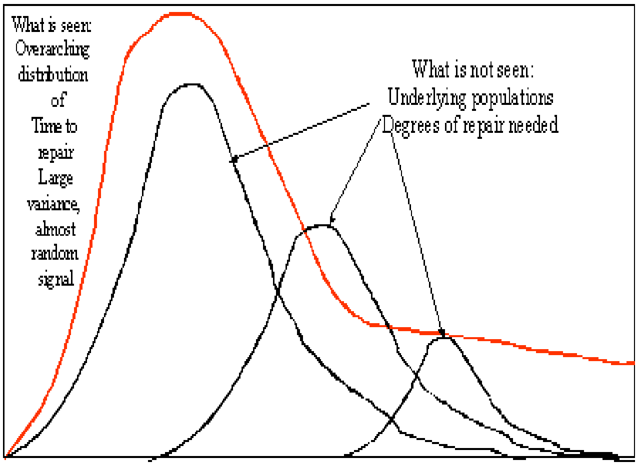

Figure 3: Hidden Repair Populations

determine what it will cost, and consequently what to charge, to perform a given task. In developing these labor standards there is little opportunity, though, to collect a reasonable number of observations from which to build a statistical description. Unlike the manufacturing environment, there are often multiple repair populations within the same end item class (see Figure 3). For example, some end items exhibit high levels of corrosion requiring extensive repair while others members of the same end item class may not. Shop supervisors and mechanics learn to recognize items that will likely take longer to repair than the labor standard allows. Driven by an efficiency-based performance measurement and evaluation system, this motivates *cherry picking* and other dysfunctional behaviors that further frustrate attempts to maintain useful labor standards.

Whether engineered using time and motion studies or derived from historical performance, labor standards are represented as a fixed numerical value. No variance in the possible processing time is accounted for in this value, whether from multiple repair populations or the learning curve. Consequently, using the same labor standards for schedule projection and shop status forecasting is highly unreliable. Differences of the depot domain are summarized in Table 1 below.

3 DEPOT DOMAIN PRINCIPLES

Although the depot domain differs significantly from the manufacturing domain, many of the key principles from the physics of manufacturing (Hopp and Spearman 1996) can be used, in conjunction with validated beliefs in the MRO domain, to understand its physics of operation. The following table summarizes the key depot system characteristics and the underlying principles that motivate depot domain dynamics.

Table 1: Depot Domain Differences

Concept	Description	Implications
Type of work	Routine, Sometimes, New Failures	• The process of discovering the work content can have drastic immediate implications to other entities far removed from the depot shop floor, but connected through functional use of the system • Maintenance Data History becomes a requirement –the history may provide insight into the reasons for failure.
Same but Different Work	Light, medium, hard broke	The same item for work can have many different magnitudes of work content, populations within populations
Location of work discovery	Location – where on the work-item was the work found.	The "point of discovery" becomes important in terms of how the discovered work relates to the remaining planned work. Discovered work can have both long lead resource requirements and absolute precedence [nothing else can be done] over subsequent work resulting in stopping all work progress.
Timing of work discovery	Timing – when in the repair process was the work discovered.	Discovery of work early in the process makes facility scheduling difficult, discovery of work late in the process makes part ordering and acquisition difficult.
Influence of variability	The information system houses "operation times", but processing time variability is not kept.	• Planning based upon operation times is no better than using "flow-time" estimates. • No clear way to distinguish between whether the time information represents a "flow-time" estimate or "process-time" estimate. • Means are used where other measures would be useful. Repair times may exhibit a lognormal distribution.

Table 2: Depot Principles

Depot System Goal	Underlying Principle or Validated Belief	Domain Oriented Description
Primary cognitive activities are focused on the product as seen by the repair agent.	*Law 11: Pay me now or pay me later* – since each end item is different, the variability is countered through knowledge asset management practiced informally by the supervisor. *Law 16: Self-interest* – some mechanics have honed their skills toward a particular task(s). *Law 17: Individuality* – depot operations are subject to worker capability and/or certification policies. *Law 20: Responsibility* – depot systems have evolved to a state where the worker is provided with only guidelines (template based operation definition). *KEYOBS1* – focus of the system should be the product from the point of view of repair, not a manufacturing view of the world.	The primary focus of manufacturing is the *process* of building the product; depot focuses first on the *product* for repair, and then on the process for repair. The focus is iterative, switching from product to process It is important to realize by the analysts focusing on the product and process, valuable insight into what information should be managed about the product and where that information will provide the most benefit will be gained.
BOM does not drive repair	*KEYOBS1* – currently the task drives repair; but there is not a formal link to a bill for repair (or task). *KEYOBS2* – Once the product is identified, the process to deal with that product will drive repair, hence the need for a flexible bill.	The manufacture BOM is largely irrelevant. It is the actual repair bill that is of concern, and this bill will vary from end item to end item even when the same task is being performed. The Maintenance Data History of the repair item in combination with the process for repair drives the repair operation.
Scheduled end-items in context with drop-ins	*Law 5: Capacity* – system enters a transient state where work release exceeds capacity and hampers its determination – too much unknown work content. *Law 6: Variability* – not only is the variance large and difficult to characterize for depot systems, but it is also a function of time.	The system must be able to allow for the insertion of workload at any level, from A/C end items, to component end items inserted into the shop from another depot.
~20% of final work content discovered [routine work and sometimes work]	*Law 8: Utilization* – increases in system utilization without other changes will ensure an increase in total repair time, in a non-linear fashion, accompanied with labor over-time and accelerated equipment maintenance/failure cycles. *Law 5: Capacity* – release rate exceeds planned resource capacity. *Law 7: Variability placement* – If the operations discovered are materials or key resource intensive, it may prove to be a bottleneck to more opportunistic operations (no precedence constraints) in queue.	With respect to an A/C circa 1994, roughly 20% of the operations performed during depot MRO are unplanned (against that tail), but are planned generically. Lead-time of material is normally the bottleneck, but delays in management decision can cause additional delays of a week or more.

Depot System Goal	**Underlying Principle or Validated Belief**	**Domain Oriented Description**
~15% of final work content discovered [new failures]	*KEYOBS1* – since the materials have not been planned (may not have an NSN), salvage is almost a sure bet, thus there exists the potential of bypassing the materials system completely. *KEYOBS2* – The mechanic and supply personnel may need some latitude to establish the 'best' product for repair requirement. *Law 7: Variability placement* – as high variability work content is discovered early in Depot MRO, and resolution attempted immediately, the variable nature of the work content will, on average, cause a longer flow time tying up key resources than if it is deferred. Material requirements should be identified and sourced, even if salvaged, but the actual work should be delayed as far as possible into the total cycle so that it does not preclude other planned, low variability operations.	These operations are new to engineering (no technical documentation), planning (no planning data) and the mechanic (no experience base). Maintenance Data History is essential, but more important is the knowledge and insight gained through the repair by the mechanic – there is clear need for knowledge asset management via IDEF3.
Work content unknown	*Law 6: Variability* – any increase in variability will increase cycle time. The need is to locate the sources of this variability and implement strategies and tactics for contingent scenarios, much like a football coaching staff prepares a game plan.	'Repair as required' is typically the statement on the low-level operation card. The actual nature of variability is due in part to the condition of the item for repair, the technical nature of the task, interruptions to that task, effectiveness of the mechanic, etc.
Process plans for repair are at a high level abstraction.	*Law 6: Variability* – genericizing is a great mechanism for simplification; unfortunately it completely avoids and even precludes the type of detailed information that is necessary to capture, control and eliminate variability.	The process operations for repair are developed as basic templates that support, but do not prescribe the work. Prescription of the work is normally found in the tech orders, and these are augmented by shop notes for best practice. The need is for simplified presentation coupled with detailed knowledge management.
Minimize disassembly	*KEYOB1* – clearly with this requirement comes the need for the mechanic and supply personnel to agree on the true 'product' to be ordered for repair. Moreover, the engineer and planner must write the task and operation to identify this (perhaps non-NSN) part, and detail the procedure for removal, repair and replacement. *KEYOBS2* – The less the working aspects of the system are disturbed, the less chance for damaging other components, less induced failures	One of the distinguishing characteristics of depot is that it attempts to get the end item back into service as fast as possible with only limited disassembly. The backshops operate in a more classic remanufacturing mode of operation in that they tend to disassemble the item completely route it for repair, then re-assemble for return to the depot shop or base of origin, or distribution through the supply system.
Most effective deployment of skills	*Law 7: Variability placement* – in an effort to guard key resources, the supervisor will tend to use key resources in consulting roles and only dedicate them to jobs which are not easily definable from either Maintenance Data History or best practice.	The skill supervisor is in a continual mode of identifying the repair situation and the attributes of that situation. The supervisor does not want to send more horsepower than is necessary, yet he certainly wants to abate the risk of not having the repair accomplished close to the scheduled norm.

Depot System Goal	Underlying Principle or Validated Belief	Domain Oriented Description
Protect knowledge base	*Law 16: Self-interest* – the supervisor will instinctively protect the knowledge base from outside factors such as RIFs by increasing the utilization of the resources in the system ('keep people working'). *Law 8: Utilization* – as the system utilization rises the average cycle times will increase as well, since there is no longer any way to accommodate variability in the system.	One of the most important issues facing the supervisor is the enhancement and maturity of his/her resource base. Once this has been accomplished there is an unwritten, but nevertheless real need to protect the knowledge base adhering to *Law16: Self-Interest*. Thwart loss of key resources and therefore the loss of core system knowledge.
Maximize capital asset diversity	*Law 17: Diversity* – ensures that the system has specialists as well as utility personnel who can be called on in a wide variety of situations. *Law 19: Burnout* – job diversity ensures that monotony is not a prominent issue; Query the Maintenance Data History for patterns of repetition before they impact the system in a negative way.	Better understanding of schedules and task sequences will allow for more diversity of work, and hence allow for skill development.
Part of inventory is salvage	*Law 4: Conservation of material* – not all of the sources for material are known to the systems above the shop floor. *KEYOBS1* – the system must identify the sub-systems or items actually pulled for replacement and distinguish these from NSN parts.	Depot sometimes must acquire components and even knowledge assets from other domains. Unfortunately none of this information is tracked on a formal, recurring basis (although there is a Cannibalization form).
Knowledge expands and ages with the system	*Law 5: Capacity* – capacity changes with capability. *Law 6: Variability* – cycle time risk through unknown work content can be abated with a seasoned knowledge source. *Law 16: Individuality* – workers have always played the part of knowledge worker, but never given a formal system to exercise that knowledge.	As systems continue to age, the knowledge of repair continues to accumulate. As these resources near retirement but the systems remain in use, it is inevitable that the cycle time to perform some operations will increase as the learning curve is scaled.

While these principles are phenomenologically based, the method of application and improvement requires the use of *constructive speculation*. The inference mechanism is neither wholly inductively nor deductively based. In fact, it is a marriage between inference from the domain [the physical world of MRO], inference from the derived belief system about MRO [namely, the depot domain principles], and inference from "outside" [removed from the context of MRO] belief systems and phenomenon.

4 DEPOT DOMAIN IDEALIZATIONS – ENABLING EFFICIENT SIMULATION MODELING OF MRO

Understanding depot domain concepts and principles helps analysts develop formal model idealizations that serve as a proxy of the domain in operation. To the extent that simulation modeling constructs allow one to model the unique features of this domain, simulation can be used to test our understanding of the fundamental principles governing its behavior and performance. One can come at it from the opposite direction as well. That is, simulation analysis may be used, in part, to discover the governing principles and further refine them.

Ultimately, however, understanding these principles is valuable only to the extent that it supports some useful purpose. Simulation analysts intuitively leverage this understanding to build useful idealizations that help them determine the sensitivity of overall system performance to changes in key system parameters (e.g., workforce size and mix, task processing time variability, induction policies), identify resource bottlenecks, determine throughput capacity, and so on. Our goal has long been to go several steps further. That is, we see expanded opportunities to integrate simulation technologies with improved operational data systems to provide for significantly higher levels of asset and status visibility, more efficient shop floor knowledge capture and dissemination, and closed-loop, real-time planning, scheduling, and control processes that uniquely support depot domain principles of operation.

This vision requires higher levels of integration between the planning, simulation, and the factory state monitoring functions displayed in Figure 4 below. Double-headed arrows indicate where decision-makers currently do or would like to make comparisons to gain added insight. We use the term *world state* to describe the set of objects, facts, assumptions, relations, and constraints known or presumed to exist in the domain at a given point in time. Rich forms of representation are needed to accurately capture this information. Production plans constitute the baseline against which production performance is measured. Simulation tools provide a means for testing alternative decisions and for developing forecasts of likely future events based on a characterization of the initial conditions defining the world state. Simulated future events provide a description of possible future world states. These events are part of the experiment results obtained through experiment scenario runs. These results are usefully compared against actual events and those that were explicitly planned.

Today, the visibility decision-makers have of these depot enterprise dimensions is sketchy and disjointed. There are tools that provide support for isolated portions. For example, some tools support the comparison of planned performance against actual performance. However, there is a general lack of tools that facilitate comparisons between simulation results and actual performance and for making comparison between world states, production plans, and simulation experiment results.

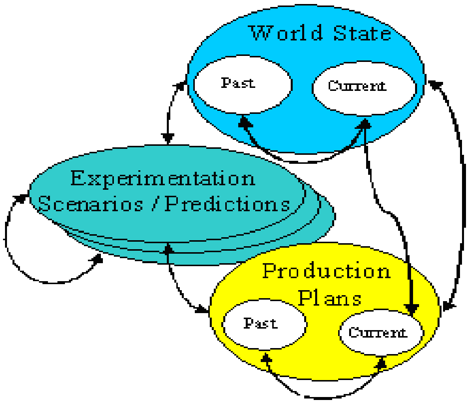

Figure 4: Total Asset Visibility in Depot MRO

Achieving the range of visibility reflected here — which includes visibility of past, current, planned, and potential future states — is most limited by shortcomings in the kinds of idealizations supported by today's simulation tools and environments. That is, the constructs provided by traditional discrete event simulation modeling tools often ignore or abstract away important characteristics and constraints that govern depot domain behavior. Clearly, there is also a general lack of detail and scope in the supporting enterprise systems that must be relied upon to provide data for simulation. These shortcomings, though, will likely be addressed only as the simulation tools themselves improve.

Our current understanding of the key issues affecting the design of future simulation tools and environments that will be applied in the depot domain includes the following observations:

- Discrete event simulation engines rely heavily upon the next event routing to schedule [in time] the next event on the calendar. The lack of time as a primary dimension and constraint motivates the need for a hybrid between production rule processing and discrete event simulation processing.
- Defining workload content for an end item a priori is not difficult; it is impossible to specify prior to the initiation of repair work on that item. Work content discovered while the work is underway ranges from 20-60%. Likewise, a discrepancy in one work-item may contribute to added workload in another work-item. Simulation engines must allow for the duration of repair to change with the state of the system variables. More generally, simulation tools must support the modeling of activities whose start and completion times are functions of the states of objects in the system rather than prescribed at run-time.
- Current simulation engines suffer from a need to both prescribe the processing or delay time a priori and/or utilize a rule firing mechanism that polls all processes for a given state change [unnecessary cycles]. Future simulation engines could benefit by taking advantage of an activity "poll-less" technology that reduces the time spent by the internal clock and calendar in determining the next best process to run.
- While workload planning can benefit from the capture and application of historical maintenance, demand, and resource usage data, scheduling is only viable if the real-time configuration state of the assets are maintained and used to drive the scheduling algorithms.

The underlying simulation engine should provide for intuitive rule specification and representation, including the following rule-based capabilities and element specifications:

- Logic based on fact- and object-instance slots;
- Rule-based change of slot state to activate control;
- Ability to dynamically generate attributes associated with an object instance;
- Ability to dynamically generate stacks associated with an object instance;
- Provide for object state transition specification, including: Input [conditions for entry], State -

[conditions for continued state membership], Termination [conditions for termination or unplanned exit from the process state], Transition [conditions necessary for successful transition from the process state], Exit [conditions for planned exit from the process state].
- Provide entity, resources/entities, operation, queue, process, machine, simulation-model, and operation idealizations that incorporate those object state specifications;
- Provide a rule development and syntax generation that follows the simple "list processing" style formulation used by CLIPS. Other rule specification capture and loader utilities may also be provided by companion process modeling or domain description capture tools.

The simulation engine should be optimized for events based on unknown work content. Standard time-based events should be handled through the event calendar by eliminating the polling of system statuses.

5 CONCLUSIONS

The focus of this paper has been to introduce the critical facets of the depot MRO domain that simulation analysts should be aware of prior to modeling, including the underlying principles that drive behaviors in the MRO domains. These principles should be applied toward the development of new tools and technologies to assist both the modeling practitioner and depot engineer in their quest to improve depot operations and build a lean depot capability.

Differences in the depot domain motivate the need for a new class of simulation tools designed to more closely emulate the constraints and resulting behaviors manifested in that environment. The traditional role of simulation may also be expanded to support more efficient depot domain operations. Most particularly, we see a future for closer integration between simulation technologies and the operational data systems used to manage and control depot domain operations. This will require changes in the range and detail of data managed by operational systems and corresponding extensions to simulation technology. These developments will be motivated by a recognition of the principles that govern depot domain physics and the need for additional constructs enabling simulation modelers to capture the essential dynamics of depot MRO.

6 REFERENCES

Hopp, W. J., and M. L. Spearman. 1996. *Factory Physics: Foundations of Manufacturing Management*. Chicago: Irwin Press.

Blanchard, B. S. 1998. *Logistics Engineering and Management*. New York: Prentice Hall

AUTHOR BIOGRAPHIES

PERAKATH BENJAMIN, 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®.

MICHAEL GRAUL, a senior research scientist at KBSI, has 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.

FRANK BOYDSTUN, 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.

MICHAEL K. PAINTER is a Program Manager and Senior Systems Analyst at KBSI. He manages business process reengineering and systems development efforts that focus primarily on improving depot maintenance, repair, and overhaul (MRO); supply chain and inventory management; and advanced planning and scheduling systems. He received a B.S. in Mechanical Engineering from Utah State University in 1985 and his MBA from Texas A&M University in 2002. He is a former Air Force officer, program manager, researcher, and technical consultant on a number of private and public sector efforts to develop and demonstrate technologies for large-scale process and information systems integration.