A RAND NOTE

MODELS TO ASSESS THE PEACETIME MATERIAL READINESS AND WARTIME SUSTAINABILITY OF U.S. AIR FORCES: A PROGRESS REPORT

J. H. Bigelow, K. Isaacson

October 1982

N-1896-MRAL

Prepared for

Office of the Assistant Secretary of Defense/Manpower, Reserve Affairs and Logistics

Rand
SANTA MONICA, CA. 90406

The research described in this report was sponsored by the Office of the Assistant Secretary of Defense/Manpower, Reserve Affairs and Logistics under Contract No. MDA903-81-C-0381.

The Rand Publications Series: The Report is the principal publication documenting and transmitting Rand's major research findings and final research results. The Rand Note reports other outputs of sponsored research for general distribution. Publications of The Rand Corporation do not necessarily reflect the opinions or policies of the sponsors of Rand research.

Published by The Rand Corporation

A RAND NOTE

MODELS TO ASSESS THE PEACETIME MATERIAL READINESS AND WARTIME SUSTAINABILITY OF U.S. AIR FORCES: A PROGRESS REPORT

J. H. Bigelow, K. Isaacson

October 1982

N-1896-MRAL

Prepared for

Office of the Assistant Secretary of Defense/Manpower, Reserve Affairs and Logistics

Rand
SANTA MONICA, CA. 90406

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

PREFACE

This Note reports on the first year's progress on a two-year project sponsored by OSD (MRA&L), under Contract No. MDA 903-81-C-0381. The ultimate goal of this project is to develop a prototype methodology for assessing the effects of varying certain logistics policies and resource levels on the peacetime materiel readiness and wartime sustainability of the U.S. air forces, in order that support resources may be better planned and justified during the annual Planning, Programming and Budgeting System (PPBS) exercises.

Not all policies and resources have been considered--that would be too much for a project of this size and duration. To date we have considered only policies and resources related to recoverable aircraft components, including stocks of the components themselves, and the means for transporting and repairing them. Billions of dollars are spent annually on these resources alone.

To meet the ultimate purpose of this project, we have followed two parallel paths: we have extended the capabilities of a pre-existing model, Dyna-METRIC; and we have developed a new model, AWARES--Assessment of the Wholesale And REtail System. Both models relate the operational performance of U.S. air forces in wartime scenarios to policies and resources related to recoverable aircraft components. Both models represent the component-relevant part of the logistics system as a network of "pipelines" that correspond to the states and processes a component can be in--attached to an aircraft, in repair at a base (intermediate level repair), in transit from base to depot, in depot-

level repair, on the shelf at a base in serviceable condition, etc. Components flow through this network in a time-dependent way that is ultimately driven by a user-specified wartime flying scenario.

The models differ in that Dyna-METRIC emphasizes the "retail" part of the logistics system, which includes activities close to the flight line such as base supply and intermediate-level maintenance. AWARES, by contrast, emphasizes the "wholesale" part of the system, which includes activities more remote from the flight line (but ultimately as necessary as retail-level activities), such as wholesale supply and depot-level repair. AWARES has no history prior to this project, but Dyna-METRIC has existed, in one version or another, for several years.

Rand sponsored Dyna-METRIC's initial development. Subsequently, considerable support was provided by the Air Force and the Navy. As a result, several versions of Dyna-METRIC existed. To meet an Air Force need, a version called 3.04 has been released for use as a standard version. That release has been made to the Air Force Logistics Management Center (AFLMC) for use in the worldwide system under development called the Combat Supplies Management System (CSMS). Other Air Force users interested in a stable, checked-out model also have had 3.04 released to them, and recently it was released to the Navy as well.

The 3.04 version includes considerable detail about the so-called retail portion of the logistics system. No detail is provided about the wholesale system though the retail system is linked to it through an order-and-ship time required to obtain replacement stock. The extensions to Dyna-METRIC made under this contract were aimed at improving the representation of the wholesale system, so that the model could

- Derive demands for repair and supply at the depots from the dynamics of the flying programs at the operating bases.
- Predict the operational performance at the bases given the resources and policies at the depot.

This study is being undertaken in parallel with two others, one sponsored by the Air Force, and one funded by the Navy. These parallel studies provide opportunities to test/apply our methodology during its development. In the Navy, both the Aviation Supply Office (ASO) and the Naval Air Logistics Center (NALC) have expressed interest in making use of both models. In the Air Force, we have already begun testing AWARES, and expect momentarily to begin testing Dyna-METRIC Extended, jointly with the Readiness Initiatives Group at the Ogden Air Logistics Center (ALC).

SUMMARY

This Note reports on the first year's progress on a project to develop a prototype methodology for assessing the effects of varying certain logistics policies and resource levels on the peacetime materiel readiness and wartime sustainability of the U.S. air forces, in order that support resources may be better planned and justified during the annual Planning, Programming and Budgeting System (PPBS) exercises. Our approach is step by step. Prior to this project we had developed methodology to investigate logistics resource and policy issues in the retail part of the system--i.e., the individual base and intermediate level maintenance facility. We are now extending this methodology to cover planning issues at the individual depot (e.g., resource needs of an individual shop or work center, or the effect of wholesale war reserve materiel), then moving to issues that cut across all depots (e.g., distribution of workload among depots, or between depots and contract facilities), and finally reaching PPBS level issues that involve tradeoffs between logistics and the non-logistic functions (e.g., procurement of aircraft versus procurement of spares for that aircraft).

S.1. TWO MODELS

Rand has developed two models that we expect can be used at the lower and (possibly) middle part of our spectrum of planning levels, one of which is an extended version of the pre-existing model, Dyna-METRIC, while the other is a new model, AWARES--Assessment of the Wholesale And

REtail System. Both models relate some support system resources to the performance of the operational forces during dynamic (wartime or peacetime) scenarios.

The resources considered in both models are those relevant to recoverable aircraft components, including spares, transportation, and the resources needed for repair.[1] The support system that uses these resources is divided into several functional areas (supply, maintenance, transportation) and echelons (flight line, base-level supply and intermediate level maintenance, wholesale supply and depot-level repair). Both Dyna-METRIC Extended and AWARES represent the various functional areas and echelons as a network of stockpiles and pipelines filled with spare components, and through which components move at rates dependent on transportation and repair resources.

Both models reflect an integrated view of the support system. All the functional areas and echelons of the support system must work together to provide support for the operational forces. For example, one's plan may assume either rapid repair and small stocks of prepositioned war reserves, or delayed repair and larger stocks. But one's plan for stockage cannot assume rapid repair (and hence small stockage requirements), and simultaneously one's plan for repair assume large available stocks (and hence a need only for delayed repair).

The above has stressed the similarities between the two models, but of course there are differences. Dyna-METRIC was designed to address problems at the retail echelons, i.e., base supply and intermediate

---

[1]Of course, there are support system resources that our models do not consider, such as POL and munitions.

level maintenance. Because Air Force bases generally deal with only one MDS[2], Dyna-METRIC was designed as a single-MDS model; thus some Dyna-METRIC features operate correctly only when conducting single-MDS analyses, for example, the calculation of the expected number of fully mission capable aircraft. Expected numbers of components in pipelines and expected backorders, however, are correctly calculated in multiple-MDS analyses. By contrast, AWARES is being designed to emphasize the wholesale echelon, and will treat multiple MDS's sharing resources in common.

Another major difference is in the purpose of the models. Dyna-METRIC primarily performs capability asessments, by estimating the probability that a specified collection of resources will support a specified number of aircraft flying a specified wartime scenario. It can also suggest stockage levels for components and subcomponents so that the probability of supporting the scenario meets a user-specified target. AWARES, on the other hand, is being designed to estimate resources--particularly resources at the wholesale echelon--that are required to support a specified scenario, and to perform tradeoffs, so that a user may explore the many different collections of resources that would enable a given number of aircraft to fly a desired scenario.

In order to be able to perform the desired resource tradeoffs and other analyses, we have taken two parallel approaches. To the extent that Dyna-METRIC Version 3.04 could be easily extended and modified, we have done so. The resulting model, Dyna-METRIC Extended, is described

[2]MDS means, roughly, type of aircraft. More precisely, it stands for "Mission-Design-Series". The analogous Navy terminology is TMS, for "Type-Model-Series".

in this Note. Many desired features, however, would have required extensive changes to Dyna-METRIC. Moreover, because the purpose of this project is to develop methodology, we expect more and more differences to accumulate between the old, unextended version of Dyna-METRIC and the methodology we are developing. Thus, we decided that in the long run it would be easier and more efficient to build another model containing less detail about the retail echelons and a better portrayal of the wholesale echelons.

One of the Dyna-METRIC features we elected not to implement in AWARES is the capability to simulate repair capacity constraints at intermediate-level maintenance, which has been used to represent automatic test equipment (ATE) in various studies using Dyna-METRIC. Also, we stopped the computations short of calculating the stochastic aspects of component failures, repairs, and the like. Either of these features could in principle be added to AWARES at a later time.

S.2. DYNA-METRIC EXTENDED

The standard Air Force implementation of Dyna-METRIC, Version 3.04, has been extended and modified to include a more detailed representation of the depot system. The new model, called Dyna-METRIC Extended, allows analyses of the interaction of the support system resources and policies at base and depot levels. The analysis of the interdependent functions and echelons helps to assure balanced support. Specifically, demands at the depot for maintenance and supply are derived from the flying programs executed at the bases. Also, depot repair times and shortfalls in depot supply impact the operational performance at the bases.

The major change made to Dyna-METRIC Version 3.04 was the transformation of the Centralized Intermediate Repair Facility (CIRF) submodel into a depot submodel. Components sent to a higher echelon of maintenance are now directed to the depot serving the component, rather than to the CIRF serving the base. Different repair times and condemnation rates may be specified for base and depot levels.

A second area where major extensions were made is the area of constrained repair. A limit may now be placed on depot throughput per component. The limited-server submodel available in Version 3.04 has been expanded to allow Automatic Test Equipment (ATE) to be stationed at both base and depot levels. (Formerly, ATE could be stationed at only one level.) Finally, a new management option has been installed. When this option is selected, and when the queues for constrained repair at the base exceed some limit selected by the user, the components with the lowest priorities are immediately sent to the depot.

A convenient new feature in Dyna-METRIC Extended is automatic time scaling. The model scales the inputs and outputs internally. The use of this feature reduces the main storage and run time requirements of the model, although some of the dynamic effects of the scenario may be muted. This feature allows the use of very long scenarios (e.g., six or 12 months), whereas Dyna-METRIC Version 3.04 restricts the user to scenarios no longer than 30 days. Because wartime transportation times from depot to theatre are on the order of 30 days, a 30-day scenario is too short for the depot to have an effect on the operational forces.

However, Dyna-METRIC Extended has some features which only operate correctly when conducting single-MDS analyses (i.e., considers only one

weapon system at a time), just as was true in earlier versions. In addition, all Dyna-METRIC versions have built into them a standard (s,s-1) inventory policy under which, whenever a component is lost from a location (e.g., condemned or sent elsewhere for repair), a replacement is requisitioned. These limitations effectively prevent Dyna-METRIC Extended from being used to investigate the total wartime workload that the depot is required to process, or the tradeoffs between several weapon systems that may have to be made if there are depot resource shortfalls.

S.2. AWARES--ASSESSMENT OF THE WHOLESALE AND RETAIL SYSTEM

AWARES will consist of two modules, a workload generator and a wholesale supply and depot-level repair module. A prototype version of the workload generator is currently operating. At the time of this writing, no version of the wholesale supply and depot-level repair module is in operation. The purpose of the workload generator is to calculate two time-varying quantities:

1. The maximum flow of broken components from the theater to the depot. The depot cannot possibly repair more than this quantity, so we call this the maximum workload;
2. The minimum required issues of serviceable components from wholesale so that the operational forces may perform their required mission.

The maximum workload is driven by user-specified flying programs for all MDS's at all flight lines that the user has included in his

inputs. If the user has included all MDS's at all bases, worldwide, AWARES will consider them (although the user might choose to aggregate or simplify such a large problem). Broken components removed from aircraft at the flight lines are sent back through the several support echelons, some fraction (possibly zero) being repaired at each one. Whatever components cannot be repaired at an intermediate level, and are not lost or condemned, will ultimately arrive--after transportation and administrative delays--at the depot for repair.

The minimum required issues are driven by three criteria:

1. The minimum number of aircraft needed to accomplish the flying programs;
2. Any additional requirement for airworthy aircraft for contingencies; and
3. Requirements for prepositioned war reserve materiel.

In AWARES, no echelon will requisition a replacement component until an order-and-ship time before it is anticipated that one of these three criteria will be violated. This inventory "policy," so different from the standard (s,s-1) policy, forms a lower envelope of requirements that would be levied on the wholesale echelon by any policy that can successfully support the operational forces (i.e., that satisfies the three criteria).

The wholesale module will calculate both the required wholesale stock levels and, given the stock levels, the minimum required depot repairs of each component. Any schedule of depot repairs is feasible if at least as many components are repaired as the minimum requirements,

and no more are repaired than specified by the maximum workload. In addition, if the amounts of depot resources consumed during each repair are known (man-hours, equipment hours, bench space, etc.), their maximum and minimum consumption can be calculated, and feasible schedules for their use developed. The wholesale module is also being designed to calculate certain diagnostics and sensitivities, so that if the required depot resources exceed those available, the user will be given indications of which factors (e.g., flying hours or transportation times) may be changed to alleviate the shortfall.

S.4. STATUS OF DYNA-METRIC EXTENDED AND AWARES

Both Dyna-METRIC Extended and a prototype version of AWARES exist and are installed at the Ogden ALC. Jointly with the Readiness Initiatives Group at Ogden, we are testing both the models: The Dyna-METRIC Extended test involves analyzing worldwide support for the F-16 aircraft, while the AWARES prototype is being tested with data from the Ogden landing gear facility. It is too soon to comment on results of these ongoing tests.

CONTENTS

PREFACE iii

SUMMARY vii

FIGURES xvii

Section

I. INTRODUCTION 1

1. Purpose and Approach 1
2. The Logistics System 4

II. DEPOT EXTENSIONS TO DYNA-METRIC 13

3. Description of Dyna-METRIC Version 3.04 13
4. Motivations for Extending Dyna-METRIC 17
5. Dyna-METRIC Extended 22
6. Status of Dyna-METRIC Extended 28

III. AWARES--ASSESSMENT OF THE WHOLESALE AND RETAIL SYSTEM 29

7. Motivations for AWARES 29
8. How AWARES Works 36
9. Status of AWARES 55

REFERENCES 57

FIGURES

1. The Vertically Integrated Support System 6

2. Complexities of a Multi-MDS, Multi-Echelon Logistics System 11

3. Hypothetical AWARES Results, I 33

4. Hypothetical AWARES Results, II 34

5. Two AWARES Modules 37

6. Example Function of Time as Represented in AWARES 39

7. Accounting for Depot Level Repair Time 46

8. Accounting for Wholesale Stocks 48

9. Freedom in Scheduling Depot Level Repair 49

10. Air Force Organic Depot Maintenance 51

11. Man-Hour Expenditure Assuming Nominal Stock Levels 52

12. Man-Hour Expenditure Adjusted for Increased Stock Levels 54

I. INTRODUCTION

1. PURPOSE AND APPROACH

This Note reports on the first year's progress on a two-year project sponsored by OSD (MRA&L). The ultimate goal of this project is to develop a prototype methodology for assessing the effects of varying certain logistics policies and resource levels on the peacetime material readiness and wartime sustainability of the U.S. air forces, in order that support resources may be better planned and justified during the annual Planning, Programming and Budgeting System (PPBS) exercises.

Not all policies and resources have been considered--that would be too much for a project of this size and duration. To date we have considered only policies and resources related to recoverable aircraft components, including stocks of the components themselves, and the means for transporting and repairing them. The fact that billions of dollars are spent annually on these resources alone indicates their importance.

Because the PPBS exercises constitute the highest, most comprehensive level at which the services do their planning, we wish to consider the entire logistics system, not merely a restricted part of it. This represents an expansion beyond past work with Dyna-METRIC by both Rand and the Air Force, which has for the most part concentrated on the retail part of the system (i.e., base supply, and organizational and intermediate maintenance). Dyna-METRIC considers the retail supply and repair processes in considerable detail, but has typically been used to

look at only a few bases, and only one MDS[1] at a time. By contrast, we wish to consider the total, worldwide forces simultaneously, including all bases and all MDS's, since there are some resources (e.g., depot manpower in some skill categories) upon which every MDS may draw.

We intend to relate these resources to the performance of the operational forces, not merely to intermediate measures of support-system effectiveness such as fill rates and backorders. The two operational measures used in Dyna-METRIC are sorties and available aircraft, and we have adopted them for the current study.

Finally, the models we are developing will be capable, like Dyna-METRIC, of dealing with the highly dynamic scenarios characteristic of wartime. This will enable the methodology to deal with peacetime in a dynamic fashion as well. And peacetime has its own dynamics. For example, the F-18 is entering the Navy inventory, the F-16 is being brought into the Air Force inventory, and the F-4 is being transferred over to the Air National Guard. Thus the amount of flying done by these MDS's is changing over time, which will cause changes in the numbers of components to buy and repair.

As stated above, we ultimately intend the methodologies developed to be useful in making and justifying resource allocation decisions in the annual PPBS exercises--i.e., at the highest planning levels. However, decisions made at high levels are worthless unless operations can be made consistent with them.[2] We have therefore adopted a "bottom

---

[1] MDS means, roughly, type of aircraft. More precisely, MDS stands for "Mission-Design-Series". In the Navy, the analogous terminology is TMS, for "Type-Model-Series".

[2] Of course, operations cannot conform exactly to plans, because factors may change between the time a plan is made and the time it is carried out. For example, a plan to buy certain items may go awry if

up" approach, in which we first investigate planning issues at a relatively low level, where consistency between plans and operations is easier to assure. As our methodology becomes more and more fully developed, we will test/apply it within the Air Force and the Navy. (Indeed, this project is being undertaken in parallel with two other projects, one with the Air Force, and one with the Navy, which are, among other things, serving as testbeds for our methodology.) A preliminary version of one part of our methodology is currently being tested/applied in a joint effort with the Readiness Initiatives Group at the Ogden ALC, using data from the Ogden landing-gear repair facility. We hope to initiate additional test/applications at higher level--i.e., at AFLC in the Air Force, and at comparable locations in the Navy--in the near future. As we continue to develop the methodology, and as we come to understand the problems peculiar to each level of planning, we will work our way step by step toward the PPBS level.

Different issues are addressed at different levels of planning. At the PPBS level, the issues are very broad, so that a methodology suited for use at this level must consider a full spectrum of factors. In our case, this means that the entire logistics system must be represented in our ultimate model, from the support provided to aircraft at the flight line to procurement of material from the original manufacturers. On the other hand, extreme detail is not useful at the PPBS level. The delay between formulating a plan at this level and executing it is long (2-7 years). Circumstances change enough in that time to invalidate detailed instructions, although broad guidelines will often remain valid.

---

the price or availability of the items changes. Thus, the executors of a plan must be given the flexibility to respond to changed or unforeseen circumstances.

By contrast, planning at lower levels addresses narrower issues in greater detail. Instead of a single model to represent the entire logistics system, there should be several mutually consistent models (or versions of the same basic model) that represent different parts of the system in varying amounts of detail. Indeed, in Phase I of this work we have developed two models for use at lower (than PPBS) levels of planning, one an extended version of the pre-existing Dyna-METRIC model, and the other a model called AWARES--Assessment of the Wholesale And REtail System. These two models emphasize different parts of the logistics system, and hence are suited for addressing somewhat different issues.

This Note has three parts. Part I is an introduction. We have already described the purpose and general approach of this OSD-sponsored project, and we will finish Part I with a discussion of the logistics system we are attempting to represent in our methodology. In Part II we motivate and describe the extensions we have made to Dyna-METRIC. For the convenience of the reader, we have also included a brief description of the version of Dyna-METRIC that we modified. In Part III we discuss the AWARES model, including how it works, how it might be used, and its current status.

2. THE LOGISTICS SYSTEM

The particular resources modeled in Dyna-METRIC and AWARES are those related to recoverable aircraft components. These include not only the spares that the system may own, but also repair resources

(manpower, test equipment, etc.), and transportation facilities. From our point of view, there is no significant methodological difference between an engine and a component, so we explicitly include engines in the part of the world we are concerned with. (Of course, this leaves out a lot of other support system resources, such as POL, munitions, etc. How many of these other resources we can eventually include will depend on how well the study progresses.)

The world of recoverable components may be represented as two interacting hierarchical structures. One, the indenture structure, describes the aircraft in terms of the components and subcomponents of which it is constructed. The other, the component support structure, describes the flow of components and subcomponents through the logistics system, which is composed of maintenance and supply functions, and the transportation which moves components from place to place within the system. Figure 1 depicts the combination of these two structures in a highly schematic fashion.

2.1. The Indenture Structure

Consider first the indenture structure: Aircraft are composed of components, which are in turn composed of subcomponents. If all components and subcomponents are operating satisfactorily, the aircraft is termed Fully Mission Capable (FMC). (In reality, it could be Not Mission Capable due to Maintenance (NMCM), even if no component is known to have failed--e.g., if a scheduled inspection is due. Preventive maintenance requirements for aircraft are not presently considered in our models, although we are investigating ways to incorporate them.)

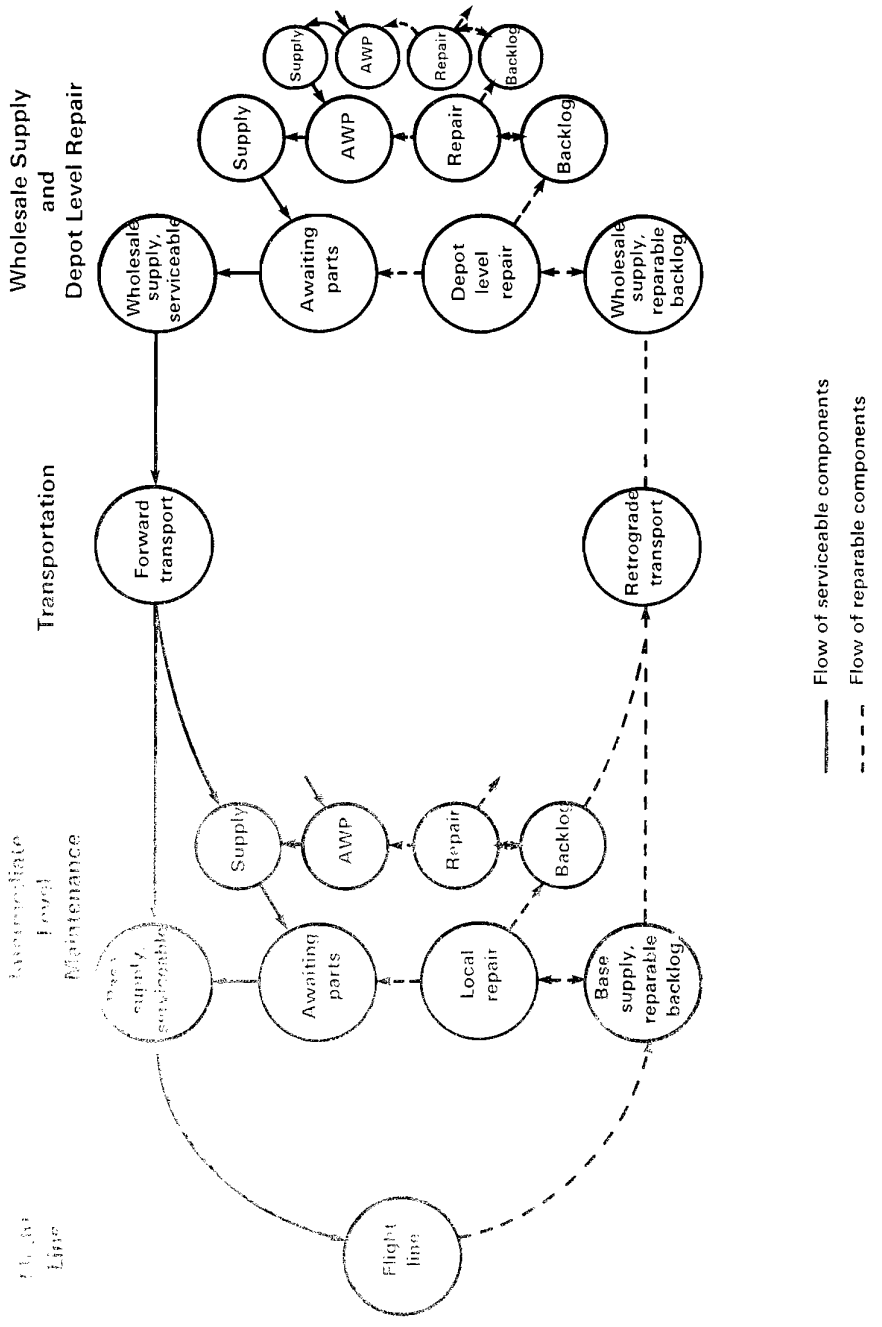

Fig. 1 — The vertically integrated support system

Failed components are modeled as being discovered, removed and replaced (if replacement stock is available) at the flight line (the leftmost bubble of Fig. 1), and sent to a shop at an Intermediate Level Maintenance (ILM) facility for repair (the double "stack" of bubbles at the center-left of Fig. 1). If no replacement is available at the flight line for the failed component, a "hole" is left in the aircraft, and until a replacement can be obtained from another location, or--if permitted--by cannibalizing another aircraft that is missing a different component, the aircraft will be left Not Mission Capable due to Supply (NMCS), and will be unable to fly any mission for which the missing component(s) is/are essential.

At the ILM, the failed component is scheduled into the repair process. During repair, it may be discovered that one or more of its subcomponents are defective. They will be removed, and the resulting "holes" in the parent component will be filled by replacement subcomponents, if available, or by cannibalizing other components at the ILM, if they are available and such actions are permitted. If the appropriate subcomponents are not available, the parent component must remain in AWaiting Parts (AWP) status until subcomponents can be found.

Meanwhile, the defective subcomponent may itself enter the repair process at the ILM, and if it has sub-subcomponents, the process can extend to a lower level of indenture. There is theoretically no limit to the number of indenture levels that can be considered, but we are unaware of any instance in which more than five levels have been considered (in planning at any level), chiefly because of the difficulty of obtaining the necessary data.

(Note the similarity between an aircraft and its components at the flight line, and a component and its subcomponents at the ILM. In both cases there is a need for replacement stock; cannibalization is a potential source of supply; and the penalty for having too little supply is a non-operable hulk--an NMCS aircraft in the one case, and an AWP component in the other.)

It is important to distinguish between the indenture structure as described by engineering drawings of an aircraft, and that implied by the behavior of flight crews at the flight line. For example, the engineering drawings of the C-5A nose landing gear show that a component called an arm assembly is a subcomponent of the nose strut. But the flightline maintenance crew will often remove the arm assembly directly from the aircraft; they do not necessarily remove the entire strut and send it to the ILM to have the arm assembly taken off. This distinction between two kinds of indenture is recognized in the terminologies used in both the Air Force and the Navy; there are Line Replaceable Units (LRU's) (in the Navy, Weapon Replaceable Assemblies, WRA's), that are removed and replaced at the flight line, and Shop Replaceable Units (SRU's) (in the Navy, Shop Replaceable Assemblies, SRA's), that are only detached from their parent components at the ILM. For our purposes, the indenture structure defined by the maintenance crews' behavior is the one of interest.

2.2. The Component Support System

The second hierarchical structure is the component support system. This system controls the flow of components and subcomponents inside and

outside the theatre and determines whether replacement stock is available to cover failed components and subcomponents.

Figure 1 shows a three-level, or three-echelon, component support structure, consisting of three echelons connected by transportation links. On the far left is Echelon 1, the flight line. This receives its most immediate support from the second echelon, the intermediate level maintenance facility (ILM), which is separated from the flight line by administrative (and perhaps transportation) delays. Any support the ILM cannot provide must be given by the third echelon, the wholesale system, which is separated from the ILM by longer delays, including both transportation and administrative delays.

Ordinarily, this is an accurate description of the component support structure as seen from an individual base. The maintenance crews on the flight line are supported by a (usually co-located) ILM, and any task beyond the capability of the ILM to perform is referred to the wholesale, or depot-level, echelon. In the Pacific Air Force (PACAF) command, however, most of the ILM capability of each of the several bases has been combined at a Centralized Intermediate Repair Facility (CIRF). Because some capability remains at the base, this has the effect of adding one more echelon to the component support structure for bases in PACAF.

In the Navy, one can also find a four-echelon example. Little engine maintenance can be done aboard a carrier. Most engines needing repair must be sent to a shore-based Air Intermediate Maintenance Department (AIMD). If the engine cannot be repaired there, it is referred to the depot-level facility, called a Naval Air Rework Facility

(NARF). For engines, therefore, a carrier sees four echelons of support: the flight deck; the shipboard AIMD; the shore-based AIMD; and the NARF.

Other exceptions to the three-echelon description are possible, even if they have yet to be tried in practice. What support structure is most appropriate to support highly dispersed deployments, or rapidly changing ones? Nor may we forget that there are many MDS's stationed at many flight lines, supported by many ILM's, some of which are themselves supported by other ILM's, and all of which are ultimately supported by wholesale supply and depot-level repair. Moreover, no two bases need be supported by the same number of echelons. This highlights the need for our methodology to allow for a flexible component support structure. Figure 2 illustrates the complexities that must be dealt with in a multi-base, multi-MDS methodology.

2.3. The Integrated View

The world shown in Figs. 1 and 2 constitutes a system for supporting the operational forces, and not merely a collection of echelons or functions. If adequate and timely support is to be provided, the different parts of that system must work in harmony. To illustrate:

- The amount of stock set aside in base supply as war reserve materiel (WRM) must be adequate to cover the incremental wartime order-and-ship and local repair pipelines. If WRM is laid in under the assumption that resupply will begin on Day 30, but the transportation facilities are planned to be tied up

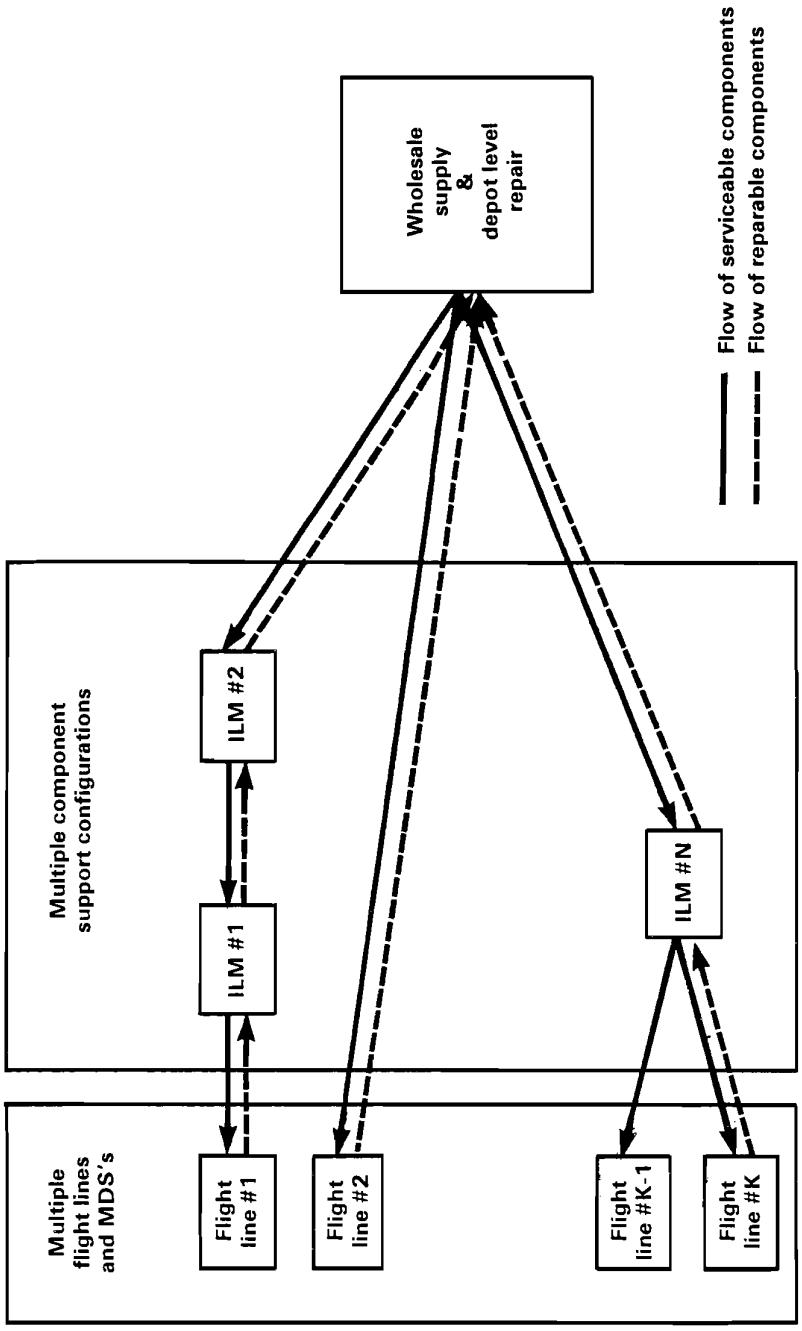

Fig. 2 — Complexities of a multi-MDS, multi-echelon logistics system

until Day 40, later generations may talk about the ten days that lost the war.

- Conversely, if there are adequate transportation facilities to permit resupply by the twentieth day, then 30 days of WRM are excessive. Why not use the resources tied up in the excess stock to buy additional airframes?
- For a sustained war, it may be necessary to have enough test equipment at the depot to repair, say, all the F-15 avionics components returned from the theater. But if the peak rate of F-15 avionics failures persists only a short time, for example due to a short-lived surge in flying activity, it may be better to allow a backlog of unrepaired components to build up temporarily, and to compensate by assigning special transportation priorities and resources to shorten the transportation pipelines.

Many other illustrations will no doubt occur to the reader. The lesson is clear: if the support system is to provide adequate support to the operational forces, the various parts of the system must act in concert. They must be integrated. The appropriate stock levels for individual parts cannot be determined individually, nor can their individual operating policies be decided separately, each under the assumption that the other parts will be resourced and will operate "as planned."

II. DEPOT EXTENSIONS TO DYNA-METRIC

3. DESCRIPTION OF DYNA-METRIC VERSION 3.04

Dyna-METRIC is a computer model designed to estimate the capability of a given component support system configuration with given component-related resources, to support aircraft flying a given (usually wartime) scenario. Dyna-METRIC calculates the expected disposition of each component across all pipelines on requested days--for example, the numbers of components in intermediate level repair, or awaiting parts (AWP), or in base supply, or installed in aircraft. Relying on a rather powerful and general mathematical result called Palm's Theorem (see Ref. [1]), Dyna-METRIC also estimates the entire probability distribution of the number of each component in every pipeline. By comparing the number of any component in base supply with the number authorized, one can determine the backorder distribution. Similarly, by comparing the numbers of components installed in aircraft with the total number required to make the aircraft fully mission capable (FMC), one can determine how many aircraft are likely to be grounded in a "not mission capable supply" (NMCS) condition.

Dyna-METRIC Version 3.04 was released to Air Force and Navy installations in August 1981. This release combined capabilities that were available in previous versions of the model with others developed to support a variety of special studies at Rand and in the Air Force. Version 3.04 will be integrated into the Air Force standard Combat Supplies Management System (CSMS) as the combat capability assessment

model. This version is the baseline model which was expanded to include features which would facilitate depot-level analyses. To provide some context for readers not familiar with that model, this section summarizes its capabilities and limitations. Particular emphasis is given to those features which will most aid or hinder the analysis of the impact of depot resources and policies upon the abilities of the bases to perform their missions.

3.1. The Physical System Modeled by Version 3.04

The physical system modeled in Version 3.04 is essentially the system described in Sec. 2 above, except that it is somewhat specialized. The indenture structure is specialized in that only two levels of indenture are allowed--i.e., components are indentured to aircraft, and subcomponents are indentured to components. In Version 3.04, sub-subcomponents cannot be represented.

The component support structure is likewise specialized. Each flight line has its own co-located ILM, although by entering the proper data the ILM may be made to have no effect. The flight line may be further supported by a Centralized Intermediate Repair Facility (CIRF), at the discretion of the user. However, the wholesale, or depot-level echelon appears in only rudimentary form, and may not faithfully forecast dynamic dislocations in the system due to variations in depot and wholesale resources and activities. Repair at the depot is not explicitly analyzed. The depot impacts the performance of bases and CIRFs through the resupply time required to obtain replacement stock. Resupply time represents forward transportation time, plus any repair

and handling time lags. Resupply times may vary between peacetime and war. Further, resupply may be cut off for some period during the conflict. Most of the modifications and extensions discussed in Sec. 6 will replace this simple depot submodel with a more complex and realistic one.

Built into Dyna-METRIC Version 3.04 (and also into the extended version) is the standard (s,s-1) inventory policy used by supply officers in both the Air Force and the Navy. When an item is declared Not Reparable This Station (NRTS), and sent to a higher maintenance echelon, a requisition is placed simultaneously for a replacement from the same facility. If serviceable stock and transportation are available, a replacement begins transit to the lower echelon. Otherwise, a replacement will be sent when a serviceable unit becomes available at the higher echelon (received either from an even higher echelon, or from repair facilities at that echelon) and when transportation becomes available.

Where an item is sent when it is NRTS'd depends on whether the base that NRTS'd it is supported by a CIRF. If so, the item is sent to the CIRF; if not, it is sent directly to the depot. The repair process at the CIRF is assumed to be the same as the repair process at any ILM: items entering repair will either be repaired and returned to stock, or will be NRTS'd to the depot. The bases served by the CIRF, because they no longer possess ILM, have a reduced repair capacity. This assumption is reflected in the implementation of the priority-repair, limited-server submodel that is embedded in Dyna-METRIC Version 3.04: a base served by a CIRF may not be assigned Automated Test Equipment (ATE). At

such a base, any failed components which are to be tested on ATE are NRTS'd to the CIRF. The only bases allowed to possess ATE are those not served by CIRFs.

The above structure, in combination with input data detailing flying programs and component and subcomponent characteristics, is used by Dyna-METRIC to construct probability distributions that describe the disposition of components throughout the system on requested days. Specifically, the model determines how many of each component and subcomponent are expected to be in base and CIRF repair, how many of each component and subcomponent are expected to be in transit or backordered from a higher echelon, and how many of each component are expected to be AWP. From these data the model derives the probability distributions describing pipelines at the bases and the probability distribution of NMCS aircraft, which are used in turn to generate the various reports and analyses which can be requested from Dyna-METRIC. Mathematical descriptions of the pipeline calculations, and the use of the probability distributions, are contained in Refs. [2] and [3]. The Version 3.04 implementation of the model is further described in Ref. [4].

3.2. Reports Available from Dyna-METRIC Version 3.04

A variety of reports dealing with individual components, available aircraft and sortie generation are available from Dyna-METRIC. The reports can be divided into two major groups: reports on performance and reports on requirements for stock or repair capability to achieve performance goals.

Reports on performance include:

1. Lists of problem components ordered by probable impact;
2. Expected backorders of each component at each base;
3. Detailed pipeline tables indicating the expected disposition of each component throughout the theatre;
4. Resource workloads for those components assigned to ATE;
5. Expected non-FMC aircraft by base under both full cannibalization and no cannibalization policies;
6. Expected FMC sorties each base can generate on selected days.

Requirements reports include recommended component and subcomponent stock purchases to achieve individual item ready rates or an overall NMCS goal, and a report on the required number of ATE to avoid queueing delays. Reference [4] includes a more detailed discussion of the reports, including several examples.

4. MOTIVATIONS FOR EXTENDING DYNA-METRIC

Although Dyna-METRIC Version 3.04 considers wholesale supply, the discussion in the previous section indicated that treatment inadequate as a tool for conducting detailed depot-level analyses. The main limitation of the Version 3.04 treatment of the depot is the assumed insensitivity of the depot to the demands placed upon it, because the depot is modeled as an ample source of stock a resupply time away from the theatre. This section discusses how a model like Dyna-METRIC can be used to study the impact of depot resources and policies on base-level combat capabilities, and then demonstrates the inappropriateness of

using Dyna-METRIC Version 3.04 for conducting many of these analyses. Section 5 discusses the extensions and modifications made to the model to provide it with a more complete representation of the depot system.

4.1. Issues to be Analyzed Using the Dyna-METRIC Methodology

One of the most important capabilities of the Dyna-METRIC family of models is that of interrelating the different echelons. Demands for repair and supply at the higher echelons can be derived from the dynamics of the flying program executed at the lowest echelon. Further, resources and policies at the higher echelons may have an impact on the operational performance at the lowest echelon. Using a model like Dyna-METRIC allows us to relate depot repair times to the ability to generate wartime sorties. If there are shortfalls in depot capability, the effects of those shortfalls on the bases' ability to achieve their flying requirements may be explored.

Being able to interrelate all the functions at base and depot levels has the decided benefit of providing for balanced support. Many advantages accrue from interrelating base and depot performance. For example:

- Better predictions of base-level performance can be made given information about depot responsiveness.
- The impacts of different operational force wartime scenarios on depot workloads can be explored.
- Scheduling for priority repair at the depot can be based on the knowledge of which components are degrading performance at the bases.

- Different stockage policies may be evaluated in terms of their impact on such measures as operational aircraft availabilities, a more pertinent measure than, for example, the expected fill rate.

Analysis which considers the interaction of multiple echelons of repair and supply is necessary to design consistent stockage policies. Such policies are currently designed separately for each echelon, and are designed to minimize intermediate measures such as total backorders, rather than higher level measures like expected non-FMC aircraft. This may result in an inappropriate distribution of resources, with some components overstocked and others understocked. Because Dyna-METRIC computes the full probabilistic distributions of demands for replacement items over time by location, stockage policies can be designed to economically satisfy such of these demands as are required to achieve the flying program. Rather than having the depot system working separately to attain goals only marginally related to base-level performance, the system will be able to maximize such wartime-relevant measures as sorties generated per day or expected non-NMCS aircraft.

4.2. Shortcomings in Using Dyna-METRIC Version 3.04 for Depot-Level Analyses

Many of the above issues may be at least initially addressed by using the Version 3.04 implementation of Dyna-METRIC. In that version, there are two possible approaches for conducting depot-level analyses: using the rudimentary depot submodel provided, or using the more complex CIRF submodel to represent a depot. Each of these approaches is discussed in turn

First consider the depot submodel embedded in Version 3.04. As discussed in Sec. 3, the depot is modeled as an ample source of stock located a resupply-time away from the theatre. The resupply time is determined by which component is being ordered, and is not influenced by which base or CIRF places the order. This is a reasonably accurate model of the depot when the following two assumptions hold:

- The bases and CIRFs are approximately the same distance from the depots. A single overseas theater being supported by a depot system located entirely in the continental United States would probably satisfy this assumption. When this holds, the real resupply time is not impacted by which base or CIRF orders the component.
- The depot has sufficient stock and maintenance capacity so that the time required to receive a component from the depot is independent of the forecast dynamic demands made on the depot.

Because of budgetary and resource constraints, the second assumption will typically be invalid. Thus, for most analyses we will want to conduct, this first approach will be inadequate.

The second approach involves using the CIRF submodel to represent a depot. The resupply time using this approach depends on the base where the component failed and not on the component itself. Because the CIRF/depot is assumed to have limited maintenance and supply resources, the receipt of a component or subcomponent by a base may require more time than just the forward transportation from the CIRF/depot:

Once the CIRF/depot on-hand stock is exhausted, no component or subcomponent will be shipped until the CIRF/depot either repairs one or receives one from the supplier after an acquisition lead time. Many useful analyses may be conducted within this framework if we may assume:

- o Depots serve families of bases rather than specific groups of components and subcomponents.
- o Either repair at base and depot is unconstrained (the repair cycle time for a component is independent of the number of demands made for repair at the base or depot), or the base has no repair capacity at all, while repair at the depot is constrained.
- o Repair-cycle times at depots are the same as at bases.

Although analyses of the depot have been made using this second approach, the above three assumptions have served to limit their scope. In particular, to capture the effects of the Technical Repair Center (TRC) concept, we would like to model depots which serve specific families of components and subcomponents, so that the depot to which a component or subcomponent is NRTS'd depends on the item and not on the base where its failure was discovered. The second assumption makes the limited-server constrained-repair submodel basically unusable, since in general the base will be located a distance from the depot and have some repair capacity of its own. Also, the depot repair-cycle time will generally be different from that of a base. Most of the extensions and modifications made to Version 3.04's capabilities facilitate the relaxation of these assumptions. In the next section, we will discuss

the changes that were made in order to produce the extended, or depot, version of Dyna-METRIC.

5. DYNA-METRIC EXTENDED

The depot, or extended, version of Dyna-METRIC is derived from the Version 3.04 implementation described in Sec. 3. The modifications and extensions which have been made to deal with the shortcomings discussed in Sec. 4 fall into four major groups:

- o The CIRF submodel was modified to become a depot submodel.
- o A limited-server constrained-repair submodel was added at the base level. (In Version 3.04, the limited-server submodel could be used at base level only if the base was not served by a CIRF.)
- o Extensions have been made to allow the evaluation of long scenarios.
- o Additional reports are available.

5.1. Creation of a Depot Submodel from the CIRF Submodel

The depot submodel in Dyna-METRIC Extended is very similar to the CIRF submodel in Version 3.04. Demands for repair and supply are driven by departures from the base repair process:

A component leaving base repair will either enter awaiting parts status (AWP), be returned to base serviceable stock, be NRTS'd, or be condemned. NRTS'd and condemned components generate demands on depot supply when the NRTS or condemnation decision is made. NRTS'd items enter the pipeline to the depot. When they arrive at the depot a retrograde transportation time later, they can be inducted for repair.

Using the projected demands for repair, the model predicts departures from depot repair. Some of these departures will represent depot condemnations for which the depot will order replacements to be delivered from the supplier an acquisition lead time later. Others will enter AWP for some period before being returned to serviceable stock, while yet others will be returned to serviceable stock immediately.

The model compares the serviceable stock levels with demands made on depot supply to determine shortages or backorders at the depot. These shortages are allocated to the bases on the basis of relative cumulative demands, and are represented at the bases as unfulfilled demands on depot supply.

Components in the extended version are assigned to specific depots. Thus, when a base NRTS's or condemns a component, the depot which is to supply and possibly repair the component is determined entirely by the component, and is completely independent of the base generating the demand. The transportation parameters involved, including forward transportation, retrograde transportation, initial availability of forward transportation, and mid-scenario cutoff of forward transportation, are determined by the base where the component failed and the depot to which the component is assigned. Determining transportation in this manner allows the bases to be dispersed geographically in relationship to a system of dispersed depots.

In the Version 3.04 CIRF submodel, the second echelon of repair was represented as though it had support processes similar to those at base level. These assumptions have been relaxed in the depot submodel,

because the repair resources may be different from those available at a base. Parameters that may vary across echelons include the repair cycle time and the condemnation rate. The demand rate for subcomponents is also different. At the base it is determined by flying hours; at the depot it is determined by the number of component arrivals for repair. In the latter case, the user defines a fraction for each subcomponent indentured to a given component which determines how often, on the average, failed subcomponents will be discovered while the given component is in repair at the depot.

Retained from the Version 3.04 CIRF submodel is the repair process: Components entering depot repair are either eventually returned to serviceable stock or are condemned and replacements are ordered from the supplier. Also retained is the stockage algorithm, which recommends a level of resources based on the expected depot pipeline for each component and the variance-to-mean ratio of the pipeline for that component. The mathematics for this approach are discussed in Ref. [3].

5.2. Extended Constrained Repair Capabilities

An entirely new Dyna-METRIC feature is a limit on the depot throughput of each component. This is useful when the user wants to prevent the model from supposing an unrealistic throughput using the unlimited-server assumption. Using this feature, he may constrain the throughput of each component per time unit to whatever level seems most reasonable given depot-level resources and capabilities.

The major extension in the area of constrained repair is the ability to use the limited-server model at both base and depot level.

Recall that in Version 3.04 there had to be no repair capacity at the bases in order to analyze limited-servers at the CIRF. This was because base ILM, including ATE, was assumed to be located at the CIRF. When changing the CIRF submodel to a depot submodel, the assumption that bases had no ILM no longer held. Now, the base first attempts to repair the component before sending it to the higher echelon, unless the component has been indicated in the input to be not base reparable.

Modeling constrained repair at both levels allowed the introduction of a new management option--intermediate level maintenance overflow. When this option has been selected, and when the queues for constrained repair at the base level exceed some limit selected by the user, those components with the lowest priorities as determined by the model (generally those with the smallest pipelines which are keeping down the fewest aircraft) are NRTS'd to the depot without first attempting to repair them. This increases demands on depot repair and supply in a manner usually not considered when projecting depot workload.

The constrained repair features of Dyna-METRIC should be used with caution. Because in most cases Dyna-METRIC can consider only one MDS at a time, depot resources should be reduced to represent only that fraction which will serve the MDS being analyzed. This is also true of the stock levels input by the user to represent stock-on-hand at the start of the conflict. In some instances, such as the facility at the Warner-Robins Air Logistics Center (ALC) for repairing F-15 avionics components, resources at the wholesale echelon are dedicated to a single MDS and this limitation of Dyna-METRIC should cause no difficulty. In other instances, such as the landing gear repair facility at the Ogden

ALC, resources are shared across many MDS's and it may be very difficult to determine an appropriate fraction of those resources to apply to each. (Note that in this second case Dyna-METRIC will still correctly predict base pipeline contents and expected demands on wholesale supply under an (s,s-1) inventory policy. But Dyna-METRIC will not allocate wholesale shortages to several bases or MDS's, and hence cannot correctly predict available aircraft at the different bases.)

5.3. Extensions for Analyzing Long Scenarios

Dyna-METRIC Version 3.04 was initially designed for scenarios of approximately thirty days. Because the transportation times between bases and depots are estimated to be at least thirty days, the depot will have no impact during such scenarios (although it can affect force readiness following such a scenario). Redimensioning Dyna-METRIC to deal with scenarios long enough to be of interest when doing depot-level analyses causes the model to require a prohibitive amount of main storage and run time. There are two features of Dyna-METRIC which can be used to facilitate the analyses of longer scenarios. These are called the restart feature and the time-scaling feature.

The restart feature works by running the model for a moderate amount of time and then writing out the contents of the pipelines to intermediate storage. The model may then be restarted using the stored pipelines as initial pipelines, and can run for some period of time, write out the pipelines, and so on. This will run no faster than the appropriately dimensioned Version 3.04, but has two important advantages. First, less main storage is required, making it feasible to

run arbitrarily long scenarios even on small computers. Second, parameters such as repair cycle times and condemnation rates can be changed each time the model is restarted. This provides a facility for time-varying demand repair and resupply parameters which is otherwise not available. When those parameters are constant, this feature is less convenient than the time-scaling feature, because the user must initiate several model runs and manipulate some data between runs.

The second feature for dealing with long scenarios is time-scaling. This feature automatically rescales the data to fit the number of time-units that are internally available to the model. This is done internally and is not noticeable to the user. The model's output is rescaled before being reported. The advantages are that the model requires less main storage and a shorter run time when doing long analyses, and that a long analysis can be done with only one run. The disadvantage is that the dynamics of the scenario tend to be smoothed out. For analyses where periods of the scenario are extremely dynamic, the restart feature should be used.

5.4. Additional Reports

In addition to the standard Dyna-METRIC reports describing base performance (expected not-FMC aircraft, expected sorties, etc.) and recommended stock purchases for components and subcomponents at bases and depots, a number of new reports have been added. These include component by component reports on pipelines and demand-related data, and a report on depot workload broken out by ATE type.

The component by component reports include expected base and depot pipelines, broken out by segment (in repair, on order, in transit, AWP), cumulative demands for depot repair and supply, expected shortfall of depot serviceable stock and expected shortfall of base stock. Much of this data is also reported for subcomponents.

The depot workload report gives a day-by-day summary of demands for repair on each type of ATE, and for each depot. Both cumulative and daily demands are reported.

6. STATUS OF DYNA-METRIC EXTENDED

Dyna-METRIC Extended, as described in the previous sections, has been designed and implemented, and its testing and verification are currently in progress. Model installation has been completed at the Ogden Air Logistics Center, and a test application of the model will be conducted by analyzing worldwide support for the F-16 aircraft. After that test application has thoroughly exercised the extended model, it will become available to other interested users.

III. AWARES--ASSESSMENT OF THE WHOLESALE AND RETAIL SYSTEM

7. MOTIVATIONS FOR AWARES

As discussed in Part II above, Dyna-METRIC Extended is well suited to the analysis of some issues that involve the wholesale part of the system. But it will not address them all. The major limitation of Dyna-METRIC stems from the fact that it assumes that an (s,s-1) inventory policy will be followed, regardless of whether this results in adequate support at any flight line. This makes Dyna-METRIC a good tool for assessing how well the (s,s-1) policy, coupled with the stated stock levels, will support the operating forces. But Dyna-METRIC cannot easily be used to investigate whether another inventory policy might allow improved operational performance. In addition, the inventory policy clearly has a substantial impact on how many of each component the wholesale system must issue, and hence on how many the depot must repair. This is an aspect of the logistics system that Dyna-METRIC is no suited to investigate.

To overcome these limitations, we judged it preferable to construct a new model, rather than to further modify Dyna-METRIC. We have called this new model AWARES--Assessment of the Wholesale And REtail System. At this writing, the initial version of AWARES has not been fully programmed. Part of it, which we call the "workload generator," is complete and in test, while the remainder, the "wholesale supply and depot-level repair module," is still being programmed. And none of AWARES is in final form. We expect to continue its development throughout this project.

AWARES and Dyna-METRIC both incorporate essentially the same view of the component support system (i.e., components flowing through pipelines) and of the indenture structure of aircraft. But AWARES does not assume an (s,s-1) policy will be followed. Instead, the user must specify explicit goals of support, in terms of flying hours (or, equivalently, sorties), non-NMCS aircraft, and available stocks of serviceable components at each base. Then AWARES determines the minimal policy that will meet these goals, where minimal means that any change for the worse--delays in transit, smaller shipments of components, etc.--will cause the goals not to be met. By contrast, an (s,s-1) policy may sometimes more than meet the goals, and may sometimes fall short.

This minimal policy is additive. That is, one can determine the minimum required shipments of components for one MDS at one base, then separately determine the minimum required shipments for a second MDS at a second base, and finally add the two. Or one may determine the total required shipments for both MDS's in a single step. The results of the two approaches are the same. This means that AWARES can simultaneously consider several MDS's. By contrast, when several different MDS's all require the same component, and stocks are insufficient to fill all orders, Dyna-METRIC must determine which MDS's shall suffer the shortage. This, the so-called "common-item problem," is what has limited Dyna-METRIC's consideration of multiple MDS's.

In a sense, therefore, AWARES calculates what is required to meet specified perfomance goals rather than (like Dyna-METRIC) assessing what

goals can be met using a given collection of resources. But AWARES can be used for the latter task if it is used iteratively, as follows. First, calculate what is required to meet the goals. Next, compare the required resources with those available. If there are resource shortages, either the performance goals must be scaled back or other parameters must be improved, such as transportation times, repair times, or NRTS fractions; and the cycle must be repeated until the user achieves a satisfactory answer. (Parameter changes such as these could represent changes in transportation priority for certain components, or addition of new kinds of repair capability at an intermediate level maintenance facility, or some other policy change.) AWARES is being designed to calculate various diagnostic quantities and sensitivities to aid the user in locating especially critical parameters, and later versions will possess more of this capability.

Initially, we are using AWARES to assess the capability of depot-level repair to support the operational forces during wartime, looking variously at an individual shop or work center, or at an entire ALC or NARF. This will shed light on the thorny question of determining the appropriate capacity--particularly organic capacity--of depot-level repair facilities. Thus, the first question we ask is, how much capacity must a depot shop or work center have if it is to carry out its wartime task?

Now assume we have somehow answered this question, and we find that the required capacity exceeds that available. Our next question is, what effects do the shortages have on the operational forces? To answer this question we need diagnostics indicating which wholesale resources

are short, at what time in the scenario, and which weapon systems will have to cut back their flying if the shortfalls are not remedied. A particular resource shortage may limit the activity of only one weapon system--e.g., the F-15 avionics test equipment at Warner-Robins--or it may affect several weapon systems--e.g., the capacity of the landing gear facility at Ogden. We wish to know by how much, and when, the activity of the various affected weapon systems must be cut back, and we wish to be able to trade off cutbacks in the activity of each affected weapon system against the others.

As a companion question, we also ask what the depot might do to alleviate the problem. Conceivably, for example, it could help for the depot to schedule the repair of critically short components ahead of the others--what we call "priority repair." Nor must the shortage necessarily be remedied by actions taken at the depot. Increased stock levels, reduced transportation times, or a reduced NRTS rate at the ILM's might also solve the problem, and any of them could conceivably be easier and less expensive to implement. That is, we want to be able to conveniently vary the mix of spares, transportation, and repair resources, to see what range of support packages could do the job. We are designing AWARES to make it easy to trade off different kinds of resources, so that the user can explore various different ways to provide support. We regard this capability as important for planning and resource allocation, and for suggesting "work-arounds" when one part of the support system fails to perform according to plan.

Figures 3 and 4 present a hypothetical example to give an indication of how this process might work. We input desired sorties as

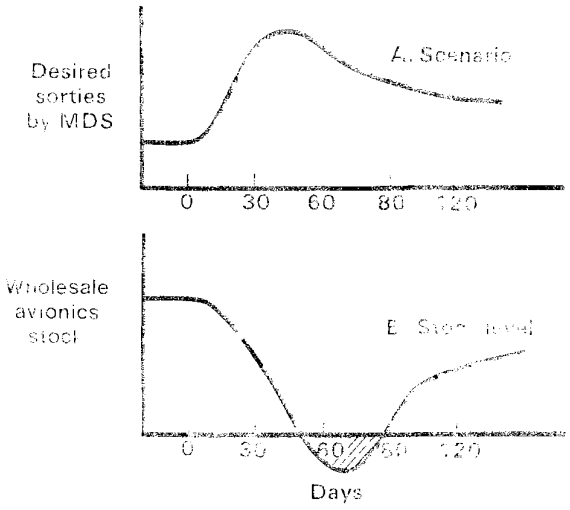

Fig. 3 — Hypothetical AWARES results, I

a function of time, as shown in Curve "A" in Fig. 3, as well as inputing all the logistics scenario parameters, which are not shown here. Suppose that when we run AWARES, we find that a particular avionics component behaves as in Curve "B." Note that demands on the wholesale system for this component exceed supply between approximately Days 50 and 90. We can tell this by the fact that wholesale avionics stock drops below zero in this period.

We wish to investigate how we might cope with this deficit. Curve "C" in Fig. 4 shows the most straightforward way, simply changing the stock level. If we increase it enough, we can achieve the desired

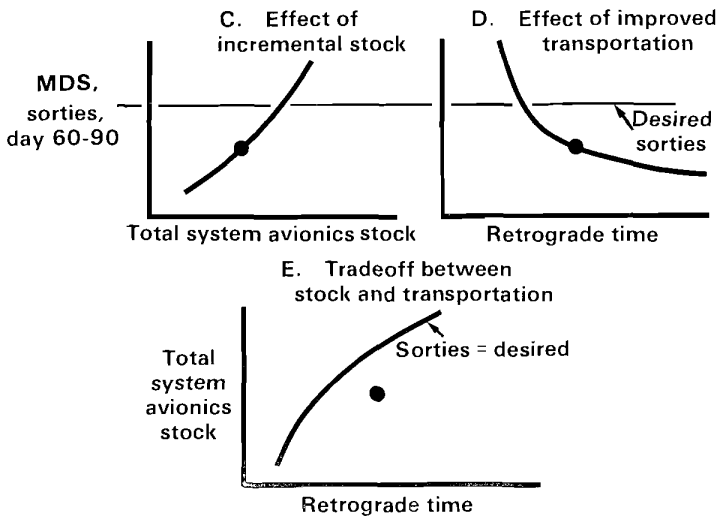

Fig. 4 — Hypothetical AWARES results, II

sorties in the critical period, which is shown as the horizontal line cutting across both Curves "C" and "D." Curve "C" also shows, if no remedial action is taken, how many sorties will have to be given up.

But adding stock is not the only possible remedial action. We could also improve transportation. If retrograde transportation were quicker, carcasses would arrive at the depot sooner, so they could be repaired and issued to the theater sooner--assuming no other bottleneck prevented this. (Equivalently, we could reduce the depot repair time.) Curve "D" shows the effect of changing the retrograde transportation time on the number of sorties flown in the critical period.

But we don't have to choose either remedy by itself. We can combine them, as shown in Curve "E." The solid dot represents the situation in the initial AWARES run. Moving up from this point represents adding stock to the system, and if enough stock is added, we will be able to fly all the desired sorties. Moving to the left from the point represents reducing the transportation time, and again, if we move far enough, we can achieve the desired sorties.

We can move in other directions as well, simultaneously adding stock and reducing the retrograde time, and movement in many such directions will eventually bring us to a point where we can fly all the desired sorties. The set of simultaneous adjustments to stock and retrograde time that allows all the sorties to be flown forms the curve labelled "sorties = desired" in Diagram "E." It is these sort of tradeoffs among resources that we intend AWARES to be able to produce, easily and conveniently.

One point must be made. All of this flexibility has its price. It is impossible to consider all bases individually, and all MDS's flying out of them, and all components, and still consider the many tradeoffs we wish to consider. Thus, we must deal with an aggregated description of the support system. It is still unclear just how we will aggregate, and how much, in order to retain the important features of the support system. This is one of the more important research topics we will address as we continue developing AWARES throughout the next year of the project.

8. HOW AWARES WORKS

Now we will discuss how the present version of AWARES actually works. (We take the liberty here of writing as though all parts of AWARES do work, whereas, as mentioned above, this is only true of the workload generator. But we describe here how the wholesale supply and depot-level repair module is being designed to work.) In AWARES, the logistics system is modeled essentially as it is described in Sec. 2 above. The indenture structure is extremely general, as an unlimited number of indenture levels can be accommodated. The component support structure is likewise general, as a flight line may be supported by as many or as few echelons as the user desires. Moreover, the component support structure can be different for each component and subcomponent. This allows a base to send some components to an ILM for repair, and others directly to a depot. Or, two components may be sent to different depots.

In spite of this generality, it is convenient to continue to view the support system in the highly schematic fashion of Fig. 1, and further to group all but the wholesale echelon into a single module, called the "workload generator." The workload generator and the wholesale module can then be described separately. This division of the system into two modules is shown in Fig. 5.

8.1. The Workload Generator

The purpose of the workload generator is two-fold. First, it must calculate the flow of reparable carcasses into the depot. No larger depot workload could be realized than by inducting these carcasses into

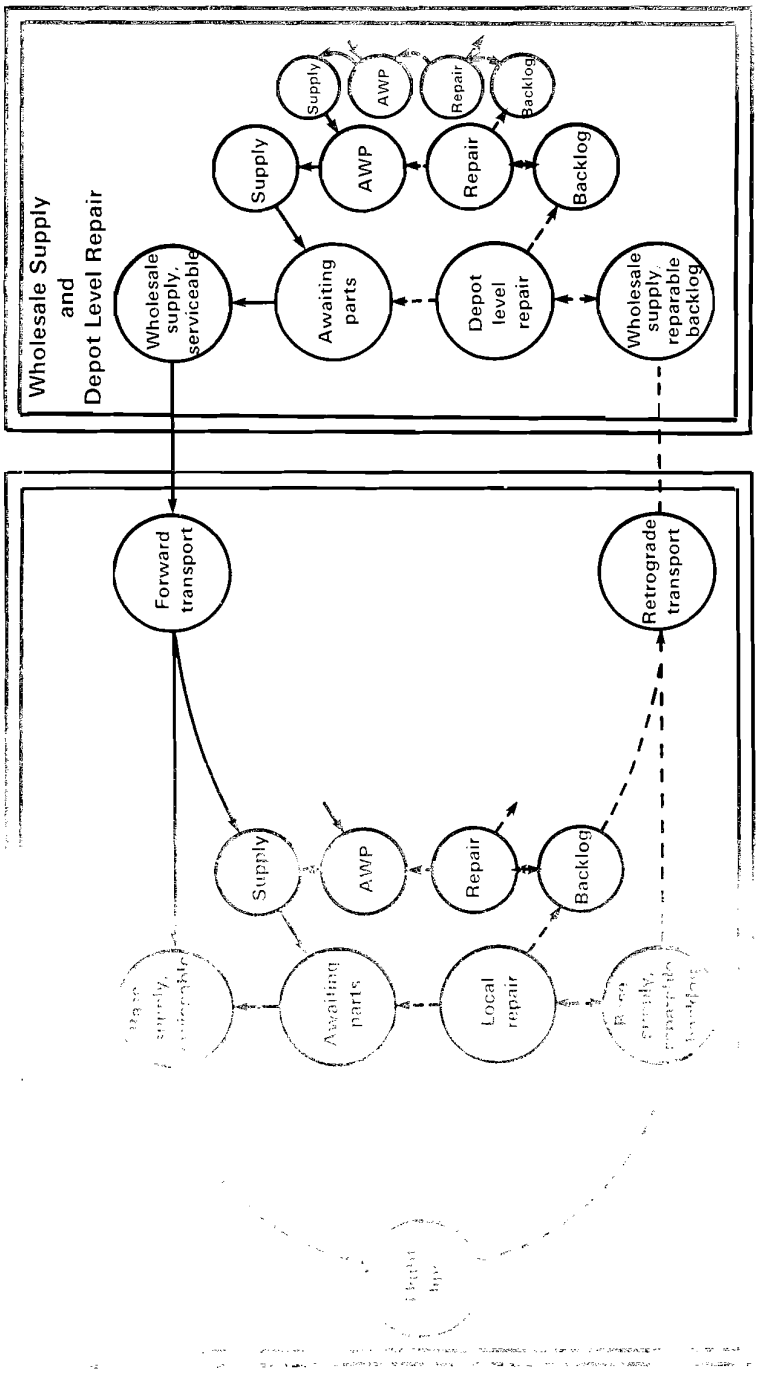

Fig. 5 -- Two AWARES modules

depot-level repair immediately upon their arrival at wholesale, and hence we call this the maximum depot workload. Second, the workload generator must calculate the required issues of serviceable components from wholesale. As mentioned above, these issues are determined to be the minimum necessary to meet certain user-specified performance goals, and not simply the issues needed to satisfy the standard (s,s-1) inventory policy. (This expresses the minimum demands on the wholesale echelon, and not necessarily the minimum depot-level repair workload, since stock on hand and procurements may serve to fill some of the requirements. We discuss this point later.)

8.1.1. Maximum Workload

The ultimate drivers of maximum workload are the desired flying programs for each MDS flying from each flight line. (Recall that AWARES deals with multiple MDS's and multiple flight lines.) These flying programs, multiplied by removal rates, yield total removals of a component from the flight line. To obtain the flow of carcasses to wholesale, these removals must be reduced by the fraction that will be locally repaired or condemned, delayed by a retrograde transportation time, and finally accumulated over all MDS's at every flight line. If there are losses during transport, they must be subtracted as well.

Virtually every factor involved in this calculation can vary during the scenario. AWARES is a dynamic model, and does not assume steady state. Figure 6 shows how an example function of time is represented in AWARES. This particular example is flying hours, but it could as easily be demands per flying hour, the fraction NRTS'd, or any of a number of

other functions. The function is piecewise constant, where the constant pieces are allowed to persist for as long or as short a period as desired. Thus, a scenario might have a highly dynamic segment, in which changes in various factors occurred daily (or more often), and also have a protracted, near-steady-state phase, in which factors seldom changed.

As mentioned earlier, AWARES considers the indenture relation among components. To generate the maximum workload of a subcomponent, it is necessary first to generate the rate at which all of its parent components pass through test at the ILM. Some fraction of the subcomponents indentured to these parents will be discovered to be faulty, and will be tested and repaired, or NRTS'd, according to their own times and fractions. The calculations are virtually identical to those for the parent components, except that the role of the aircraft is taken by the parent component, and the role of the parent component is taken by the subcomponent.

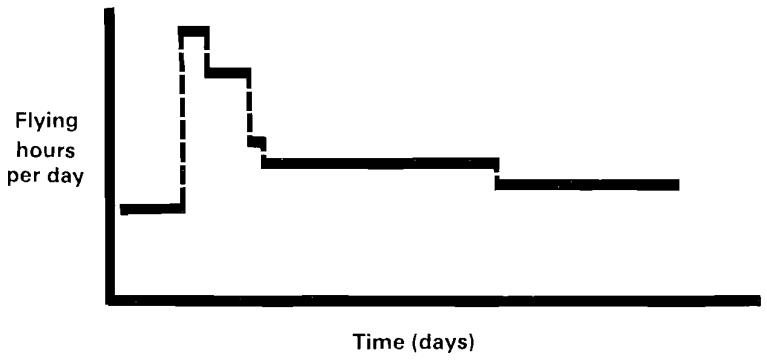

Fig. 6 — Example function of time as represented in AWARES

We have omitted two features from AWARES that are possessed by Dyna-METRIC. First, after calculating the expected contents of each pipeline at the requested times during the scenario, Dyna-METRIC then estimates the probability distributions of the number of items in each pipeline. AWARES stops short of this final step, and calculates only the expected pipeline contents. Second, Dyna-METRIC offers its users the ability to represent limitations in the capacity to repair certain components, such as might result from an insufficiency of automatic test equipment for avionics components at an intermediate-level maintenance facility. AWARES represents repair as a process that requires a user-specified amount of time (the time may be specified to change during the scenario), but whose capacity is not limited. Either or both of these features could be added to later versions of AWARES.[1]

8.1.2. Minimum Requirements

The calculation of requirements for serviceable components is very similar to the calculation of reparable generations just discussed. What AWARES calculates is a minimum requirement. This is quite

[1] In its present application, that of estimating the capability of a depot-level shop or work center to support the operational forces in wartime, we do not feel it necessary that AWARES possess these features. Rather than programming AWARES to calculate probability distributions for the contents of all pipelines worldwide, we will make off-line after-the-fact estimates of variation in the content of selected pipelines of interest, in particular the depot repair pipeline. And, so long as we are primarily interested in the effect of limited ILM repair capacity on the work the depot is required to perform, we think we can adequately represent capacity limitations by judiciously increasing repair times and/or NRTS rates at the ILM. Because these features are not in AWARES, the model requires less data, and is easier and less costly to run. Of course, in future applications these features may prove to be as important in AWARES as they have been in Dyna-METRIC, in which case we will incorporate them in a later version of the model.

different from the usual (s,s-1) inventory policy, in which any item removed from the aircraft must be replaced by a serviceable item, and any item that base supply returns to wholesale must be replaced. Instead, AWARES tries to link the requirements for serviceables more directly and closely to the ability of the operational forces to accomplish their mission.

The minimum requirements are driven by three factors:

1. The minimum number of aircraft needed to accomplish the flying programs;
2. Any additional requirement for FMC (more precisely, for non-NMCS) aircraft for contingencies; and
3. Requirements for prepositioned war reserve materiel (WRM).

At the flight line, AWARES will allow aircraft to become grounded for lack of spares until there are only the minimum number of non-NMCS aircraft, as determined by Criteria (1) and (2), and it will delay this event as long as possible by judicious cannibalization. Once this point is reached, AWARES will fill the critical holes in aircraft just fast enough to keep the required number of aircraft flyable. AWARES makes this calculation for each component individually, and while it may be necessary to begin delivering one component to a flight line by, say, Day 15, deliveries of another component may not be needed until, perhaps, Day 30.

Stepping back one echelon to the ILM, we find that similar considerations apply. At the least, the ILM must provide the minimum required serviceables to the flight lines it supports. There may also

be a stockage requirement at the ILM (Criterion (3)). As we have mentioned before, AWARES is currently an expected value model, and does not model the stochastic nature of component failures. One reason for an extra stock requirement at the ILM can therefore be a safety level.

But the main reason for stock in excess of that required by Criteria (1) and (2) is that we don't really know what scenario we are in. We are going along in peacetime, flying training flights, and suddenly, without (much) warning, a conflict begins. Or we are in the midst of a limited conflict, and it suddenly escalates. At the ILM, these uncertainties are met by maintaining WRM stocks. (Indeed, the services calculate their required WRM stocks by positing a "most demanding" wartime scenario, and estimating the stocks of components needed to support their forces during that scenario.)

These requirements at the ILM are met first from serviceable stock available locally, second from stock that can be repaired locally, and, only as a last resort, by receipts of serviceables from wholesale. Thus, receipts of wholesale can be calculated by subtraction: Receipts from wholesale equals requirements of the flight line plus WRM requirements minus local serviceable stock minus local repairs.

The final step in this process is to advance the receipts of serviceables at the ILM by the appropriate order and ship time, and to accumulate the results for all ILM's. The result is the minimum required issues of serviceable components from the wholesale system, as a function of time. If there are losses during transport, these issues must be increased to cover them.

As mentioned earlier, AWARES will treat indentured components, and its method for doing so is similar to the method for treating the parent component. However, the parent component occupies the same role for the subcomponent as the aircraft did for the parent component. As AWARES tests the parent component, broken subcomponents will be discovered, and the parent components will enter awaiting parts status. AWARES will allow more and more parent components to enter AWP status, until there are only the minimum permissible number of serviceable parent components at the ILM, as determined by Criterion (3) above, and, if more than one subcomponent is indentured to the parent, AWARES will delay this event as long as possible by judicious cannibalization. Once this point is reached, AWARES will fill the critical holes in the parent components just fast enough to maintain the minimum number of serviceable parents at the ILM. An important implication of this procedure is that parent components are never requested from wholesale as long as any are in AWP status. Thus, not only are all demands made as late as possible and as small as possible, they are also made for components at as low a level of indenture as possible.

The demands for subcomponents generated in this fashion are met first from serviceable stock available locally (in excess of the minimum stock specified by Criterion (3)), second from stock that can be repaired locally, and last by receipts from wholesale. Included in the stock that can be repaired locally are subcomponents in AWP status, held there for sub-subcomponents. In short, the indenture relation can be carried as deep as desired.

Before leaving this subject, some comments are in order. First, AWARES calculates minimum requirements, not requirements as they would be determined by following an (s,s-1) ordering policy, or any other "real" policy. By doing so, AWARES calculates what resources are necessary to support exactly the scenario under consideration, and nothing extra. Following this principle has the considerable advantage that a requirements calculation by AWARES is consistent with an AWARES capability assessment. If we followed another principle, we could calculate required resources, and then find that the resources could be cut (up to a point) without compromising the scenario.

Second, we discussed earlier how the minimum requirements are driven by the specification of flying hours, required non-NMCS aircraft, and required prepositioned WRM. The flying hours are directly related to the scenario one is currently investigating, but the other two drivers are not. Instead, they reflect some notion of what other potentially more demanding scenarios might lie around the corner.

Third, we think the concept of minimum requirements is potentially useful--and in fact is frequently used, if not recognized as such--in resource planning and resource allocation issues. We think it is less useful as a day-to-day management principle, and we are not advocating it as a way to run the support system. In any scenario, including the most extreme scenario, some resources will be "slack"--i.e., not fully utilized--some of the time. In AWARES, these resources will tend to gravitate to the wholesale echelon. If there is spare stock, for example, reparables will flow to wholesale, but serviceables will not be sent forward until needed. But in practice, a different distribution of

spare stock may be preferred, that leaves some at wholesale, but sends some forward to the engaged forces.

8.2. The Wholesale Module

At this point, AWARES has calculated for us the flow of carcasses to the wholesale system, and the minimum required issues of serviceable components from wholesale. These two quantities, both expressed as functions of time, are inputs to the AWARES wholesale-supply and depot-level repair module. This module will take the two input functions and, by a combination of repairing carcasses and drawing upon serviceable stock at wholesale, will ensure that serviceable components can be issued at a rate at least equal to the minimum required. Note that we speak of this module in the future tense. The workload generator, which we described above, is already implemented, at least in prototype form. The wholesale module is still being put together.

The purpose of this module is two-fold. First, it will estimate the wholesale resources necessary to accomplish this task, in terms of available spares (i.e., non-prepositioned, or "other" war reserve materiel, OWRM) at the start of the scenario, and the repair resources such as manpower (by skill), facilities (by shop), and equipment. Second, in the event that the required resources exceed those available, the module will calculate diagnostics that indicate which resources are short, which parts of the workload those resources are related to, and how scheduling or other management changes might help to alleviate the shortages.

8.2.1. Required Depot Repairs of a Component

To see how this module might work, let us suppose that the workload generator has calculated both the carcass arrivals and the required serviceable issues of a particular component. In Fig. 7, carcass arrivals are cumulated from some time in the scenario before any interesting action starts, and plotted as a function of time (the left curve). If every carcass is inducted into repair the moment it arrives, then the carcasses will emerge as serviceable components a repair time later. The cumulative repairs can be found by shifting the right-hand curve one repair time to the right.

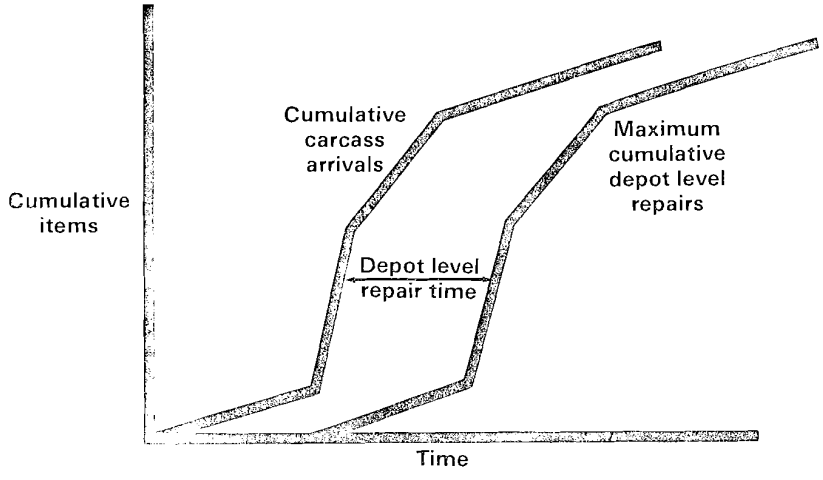

Fig. 7 — Accounting for depot level repair time

Other adjustments are possible. For example, if there is a reparable backlog at the start of the scenario, the left curve--and hence the right curve as well--can be shifted upward by an amount equal to the backlog. If the retrograde transportation time is reduced, both curves are shifted to the left. If the depot repair time increases for some reason, then the right curve--but not the left--is shifted to the right. If some of the carcasses are condemned, the right curve--not the left--must be scaled down by a factor equal to one minus the condemnation fraction. By these means, it is possible to easily calculate the maximum cumulative repairs of this item for a host of different support system parameters.

We can do similar things with the minimum serviceable issues from wholesale. The upper curve in Fig. 8 represents the cumulative serviceable issues of the item we are considering, cumulated from the same point in time as were the carcass arrivals in Fig. 7. But if there are any items in stock at the start of the scenario, it will not be necessary to receive serviceables out of repair as soon as this curve says that serviceables must be issued. The first issues can be from the stock initially on hand. To adjust for this, we lower the curve by an amount equal to serviceable stock, obtaining the lower curve. Note that where this would result in the lower curve falling below zero, we have truncated. We have a name for the lower curve: Minimum Cumulative Required Depot-Level Repairs.

As with Fig. 7, we can make other adjustments here. If procurements are scheduled to arrive at various points in the scenario, we can lower different parts of the upper curve by different amounts.

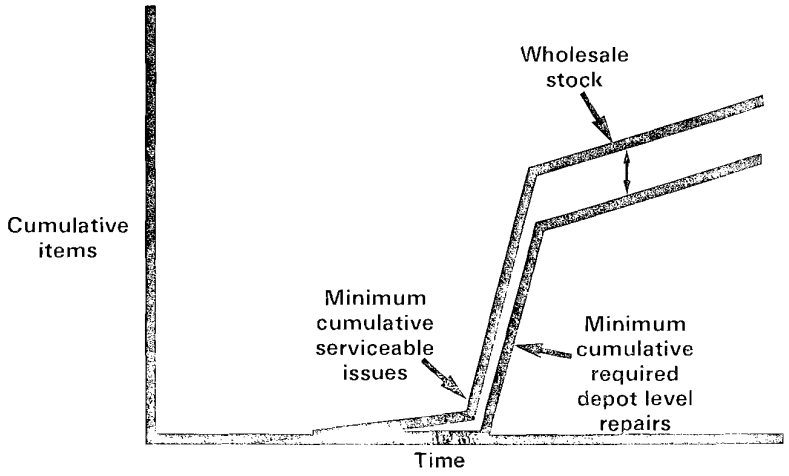

Fig. 8 — Accounting for wholesale stocks

If the order-and-ship time changes, both curves will shift, to the left if the O&ST increases, to the right if it decreases.

Now we put the maximum possible cumulative repairs and the minimum required cumulative repairs together in Fig. 9. In this example, we are in luck. The minimum required repairs curve everywhere falls below the maximum possible repairs curve, meaning that at no time will our requirements for repairs outstrip our available carcasses. If this had not been the case, we would have had to adjust the parameters of the model to make it true. As described with the help of the previous Figs. 7 and 8, this could involve changes in transportation or repair times, stock levels, condemnation rates, or perhaps other parameters.

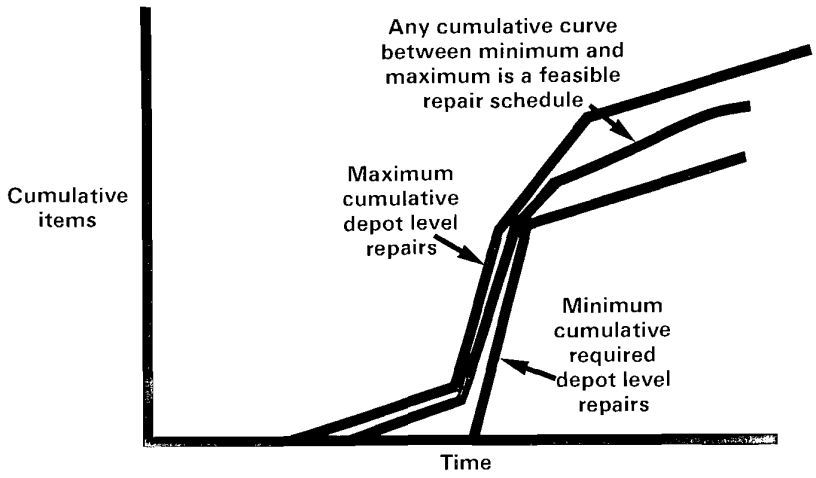

Fig. 9 — Freedom in scheduling depot level repair

But once the minimum required repairs curve is safely made to lie below the maximum possible repairs curve, we see something else. Any cumulative curve that lies between the two is a feasible repair schedule for this item. The fact that it lies below the maximum possible repairs curve ensures that there will be enough carcasses available to meet the schedule. The fact that it lies above the minimum required repairs curve means that serviceables will be repaired and turned into supply fast enough to meet the need for serviceable issues from wholesale.

8.2.2. Depot Resource Expenditures

This same idea can be used for investigating repair resources such as manpower, shop space, and equipment. But a step must be taken first to convert the carcass arrivals and the serviceable issues into the maximum and minimum cumulative expenditures of the resource, respectively. To do this, one must learn how much of each resource is consumed by the repair of each item. And in order to develop the cumulative resource expenditure curves for a resource, one must accumulate over all items whose repair uses that resource.

Figure 10 suggests one of the ways to categorize depot-level repair resources for the Air Force (it could be done similarly for the Navy), in this case by shop. The five ALC's appear as columns. The boxes under each ALC represent highly aggregated shops, and give a pretty good overview--although not complete--of the kinds of work done at each ALC, such as landing gear at Ogden, avionics at Warner-Robins, etc. Different kinds of workload come in from the left. In the center-left, components generated at the flight lines are entering, and there are several categories--airframe, landing gear, avionics, weapons, and engine components. At the top left, whole aircraft arrive for programmed depot maintenance, PDM, and they give rise to a flow of components--denoted by the dashed lines--into the various "shops" at the different ALC's. At the bottom left, whole engines arrive for overhaul, and they, too, give birth to components.

Figure 11 shows some hypothetical curves of cumulative man-hour expenditures (it could be any resource), curves that we will suppose were obtained as described above. The uppermost of the cumulative

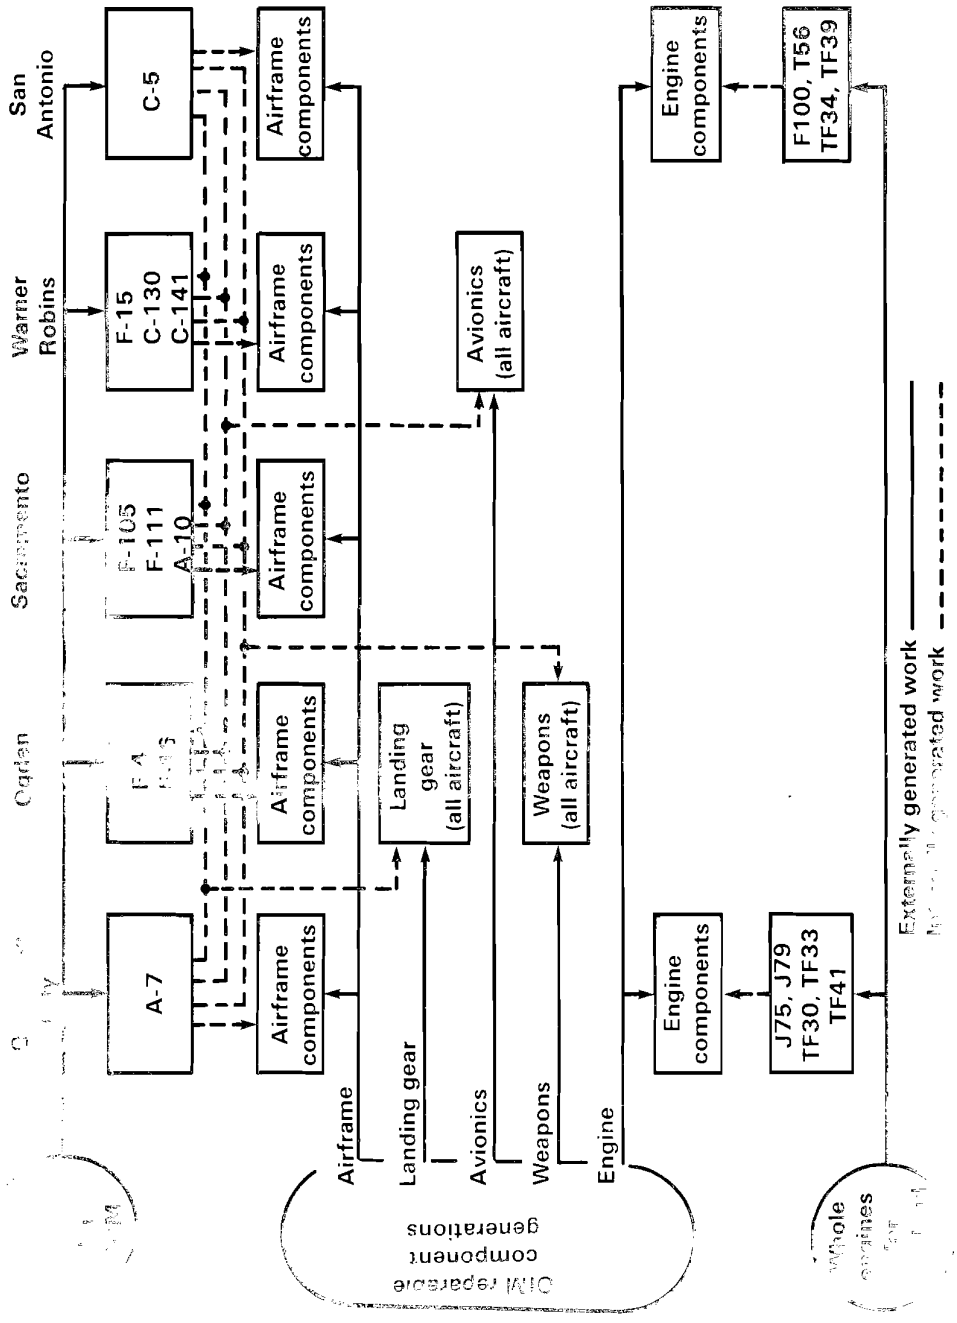

Fig. 10 — Air Force organic depot maintenance

curves represents the man-hours that would be expended if every carcass were inducted into repair as soon as it arrived at the depot. This is therefore the maximum possible cumulative expenditure curve, in that it is not possible to spend more man-hours by any point in the scenario than is shown by this curve.

The lowest cumulative curve represents the man-hour expenditures that would have to be expended if depot repairs were held to a minimum. To calculate it, we first decide how much stock of each component is available at the depot, and adjust the minimum required issues curve

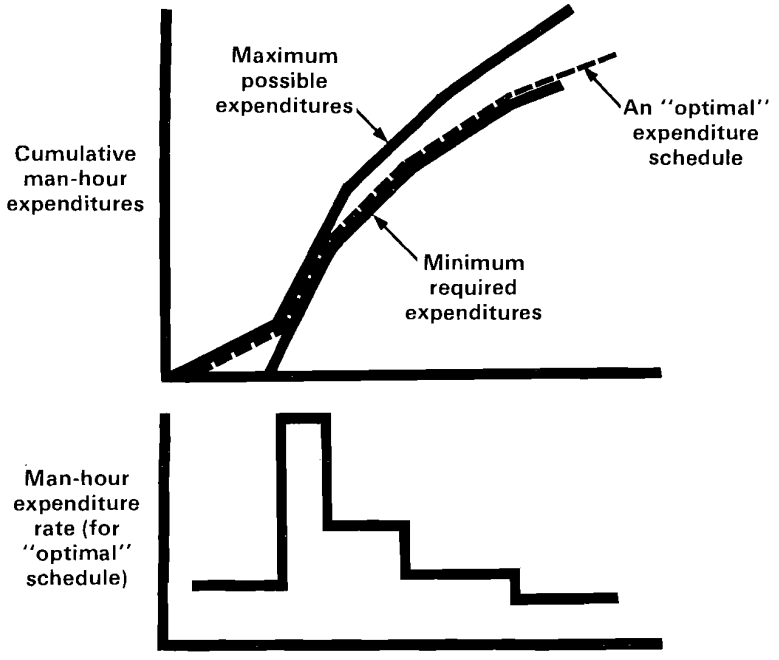

Fig. 11 — Man-hour expenditure assuming nominal stock levels

accordingly (see Fig. 8) to obtain the minimum depot repairs for each component. These we convert to man-hours and accumulate over all items, as we did with the maximum depot repairs curves. Note that if there is sufficient wholesale stock for each item so that the required serviceable issues can be met, then the minimum cumulative man-hour expenditure curve must necessarily lie below the maximum curve.

As we had some flexibility in scheduling the repairs of individual components, so we also have some flexibility in scheduling man-hour expenditures. Any cumulative curve lying between the minimum and the maximum represents a feasible schedule. In the dashed line we have chosen a particular schedule that we call "optimal" because its maximum rate of man-hour expenditures has been made as small as possible. The rate of man-hour expenditures is, of course, equal to the slope of the line; thus, in the "optimal" schedule we have kept the maximum slope to a minimum. The lower graph in Fig. 11 shows the man-hour expenditure rate throughout the scenario for this "optimal" schedule.

As we discussed earlier, by changing transportation times, repair times, stock levels, or other parameters of the system, we can cause the maximum and minimum cumulative curves for individual components to shift either vertically or horizontally. These shifts have their effect on the cumulative man-hour expenditure curves, and in Fig. 12 we show the (hypothetical) effect of increasing the wholesale stock of one or several components. The maximum cumulative man-hour expenditures remains as it was in Fig. 11, but the minimum curve is shifted downwards. This moves the two curves farther apart, offering greater flexibility in scheduling depot-level repairs. A new "optimal" schedule

can therefore be developed that requires a lower peak rate of man-hour expenditures than did the old schedule.

Together, Figs. 11 and 12 illustrate how AWARES could be used to perform tradeoffs among resources, in this instance a tradeoff between stock and repair capacity. The required repair capacity is proportional to the peak required man-hour expenditure rate. Thus, we must maintain a capacity almost twice as great for the situation depicted in Fig. 11 as we must for the situation of Fig. 12. However, the situation of Fig. 12 requires a greater investment in stock. It is not possible to determine a priori which situation is the more economical.

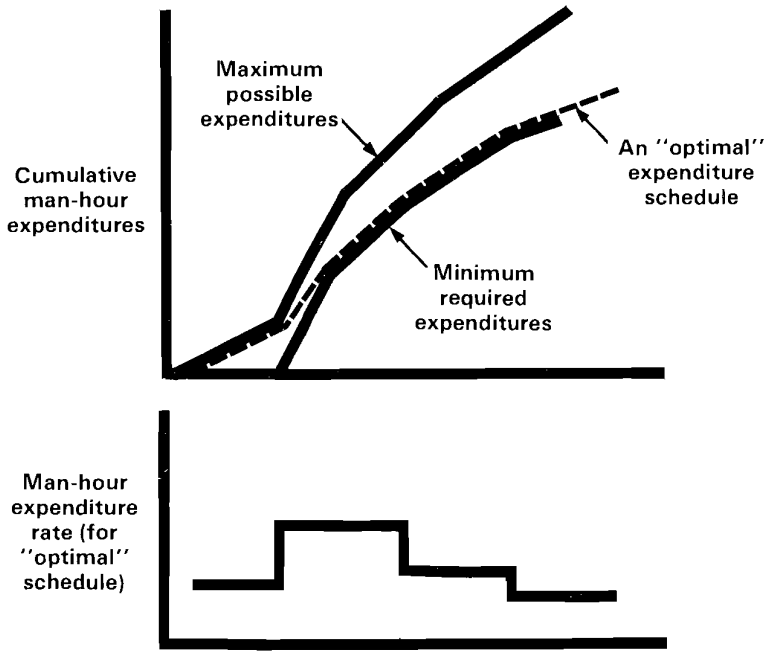

Fig. 12 – Man-hour expenditure adjusted for increased stock levels

9. STATUS OF AWARES

We are presently engaged in an exercise to look at repair resources other than carcasses available. Jointly with the Readiness Initiatives Group at Ogden, we are trying to prove this methodology for the Ogden landing gear facility. A preliminary version of the workload generator is installed and working at Ogden. It calculates the required serviceable issues and the carcass generations. We are in the process of building a preliminary version of the wholesale supply and depot-level repair module. Enough is implemented to calculate the minimum wholesale stock necessary to make the minimum required repairs curve fall below the maximum possible repairs curve. We are presently adding a feature that will aggregate over several items, so it can estimate total manpower resources in the landing gear facility.

We are collecting data on the items that are repaired in the landing gear facility, and have begun collecting information on the man-hours expended in repairing each item.

Some preliminary AWARES runs have been made with the partial data we have collected so far. We intend to report some results of this work in a future publication.

REFERENCES

1. Crawford, G.B., Palms Theorem for Nonstationary Processes, The Rand Corporation, R-2750-RC, October 1981.

2. Hillestad, R. J., and M. J. Carrillo, Models and Techniques for Recoverable Item Stockage When Demand and Repair Are Nonstationary, Part 1: Performance Measurement, The Rand Corporation, N-1482-AF, May 1980.

3. Hillestad, R. J., Dyna-METRIC: A Mathematical Model for Capability Assessment and Supply Requirements when Demand, Repair and Resupply are Nonstationary, The Rand Corporation, R-2785-AF, April 1982.

4. Pyles, Raymond, "The Dyna-METRIC Readiness Assessment Model: Motivation, Capabilities and Uses," The Rand Corporation, unpublished working paper, February 1982.