file Copy

(2)

AD-A212 826

R-3389-AF

Dyna-METRIC Version 4

Modeling Worldwide Logistics Support of Aircraft Components

Karen E. Isaacson, Patricia Boren,
Christopher L. Tsai, Raymond Pyles

May 1988

40 Years

RAND

PROJECT AIR FORCE

The research reported here was sponsored by the United States Air Force under Contract F49620-86-C-0008. Further information may be obtained from the Long Range Planning and Doctrine Division, Directorate of Plans, Hq. USAF.

ISBN: 0-8330-0890-0

Published by The RAND Corporation
1700 Main Street, P.O. Box 2138, Santa Monica, CA 90406-2138

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER R-3389-AF	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Modeling Worldwide Logistics Support of Aircraft Components		5. TYPE OF REPORT & PERIOD COVERED interim
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) K. Isaacson, P.M. Boren, C. L. Tsai, R. A. Pyles		8. CONTRACT OR GRANT NUMBER(s) F49620-86-C-0008
9. PERFORMING ORGANIZATION NAME AND ADDRESS The RAND Corporation 1700 Main Street Santa Monica, CA. 90406		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Long Range Planning & Doctrine Div, (AF/XOXFP) Directorate of Plans, Ofc DCS/Plans & Operations Hq USAF, Washington, D. C. 20330		12. REPORT DATE May 1988
		13. NUMBER OF PAGES 212
14. MONITORING AGENCY NAME & ADDRESS(If different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

no restrictions

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Models
Logistics
Logistics Management
Combat Readiness
Operational Readiness
Spare Parts

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

See reverse Side

DD FORM 1473 1 JAN 73

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

This report describes Dyna-METRIC Version 4, a computer model that relates logistics resources and policies to wartime readiness. Developed for the use of logisticians to improve wartime logistics support, Dyna-METRIC assesses the effects of wartime dynamics and repair constraints and provides operational performance measures, problem detection, and spares requirements. Version 4 consists of five programs that provide for three echelons of interaction (including the depot-to-theater link) and three levels of components (for which demand processes, repair processes, and spares levels may vary). Dyna-METRIC portrays component support processes as a network of pipelines through which aircraft components flow as they are repaired or replaced within a single theater. Using the sum of all the pipeline segments, Dyna-METRIC determines the complete probability distribution for the number of parts in repair and on order. Combining the distributions for all components provides the estimate of aircraft availability and sorties.

R-3389-AF

Dyna-METRIC Version 4

Modeling Worldwide Logistics Support of Aircraft Components

Karen E. Isaacson, Patricia Boren,
Christopher L. Tsai, Raymond Pyles

May 1988

A Project AIR FORCE report
prepared for the
United States Air Force

40 Years 1948-1988
RAND

PREFACE

The RAND Corporation developed new analytic methods for studying the transient behavior of component-repair/inventory systems under time-dependent operational demands and logistics decisions like those that might be experienced in wartime (Hillestad and Carrillo, 1980). These methods underwent a series of extensions that culminated in the Dyna-METRIC model described here.

Dyna-METRIC evolved through a series of projects at RAND that addressed the evaluation of alternative policies for improving readiness and supportability. The dynamic queueing equations were developed and applied to the problem of determining appropriate levels of spare engines for the C-141 aircraft to support planned wartime surges in flying activity (Berman, Lippiatt, and Sims, 1978). Indentured component features and aircraft availability performance measures were added to study alternative logistics policies for carrier-based aircraft squadrons (Rice, 1979). Further features were added to enable the analysis of multiple bases of aircraft and their dependence on transportation of spare parts, in support of the Project AIR FORCE study, "Responsive Intra-Theater Transportation for Spare Parts." The model was extended to include both constrained and unreliable repair resources (test equipment) for another Project AIR FORCE study, "Supporting Modern Tactical Avionics." Centralized repair and resupply capability features were added to allow Dyna-METRIC to be embedded in the experimental Combat Support Capability Management System being tested in the Pacific Air Forces under the Project AIR FORCE study of the same name. All such enhancements resulted in a version of Dyna-METRIC called 3.04; to meet an Air Force need, it was released for use as a standard version (Pyles, 1984).

Dyna-METRIC was further extended to better represent the wholesale system, deriving demands for repair and supply at the depots from the dynamics of the flying program at the operating bases and predicting the operational performance of the bases given the resources and policies at the depot. This was done under the "Wartime Sustainability Project" for the Office of the Secretary of Defense.

A complete redesign of the model was supported by the "Improving Readiness Assessment and Management Study" under Project AIR FORCE. The redesign made the model more reliable and efficient. Several new features were incorporated, including:

- An additional level of component indenture.
- An additional echelon of component repair and supply.
- The ability to analyze multiple aircraft types.
- Substantially more detail in the description of component pipelines.
- Expanded report capabilities.

The redesigned model is Version 4 and has been released as a standard version. The Management Sciences Division (XPS) of the Air Force Logistics Command at Wright-Patterson AFB is currently responsible for this version.

This report is intended for users of the Version 4 Dyna-METRIC model and for others who wish to understand the model's underlying mathematics. A more rigorous mathematical description may be found in Hillestad (1982) and Hillestad and Carrillo (1980). This report provides a comprehensive description of the model's motivation, capabilities, use, and mathematics.

This document was prepared as part of the Project AIR FORCE Resource Management Program, specifically the studies of "Improving Readiness Assessment and Management" and "Uncertainty/Variability."

SUMMARY

This report describes Dyna-METRIC Version 4, a computer model that **relates logistics resources and policies to wartime readiness.** Developed for logisticians to use to improve wartime logistics support, Dyna-METRIC assesses the effects of wartime dynamics and repair constraints and provides operational performance measures, problem detection, and spares requirements.

Version 4 is expressly suited to **conducting worldwide analyses of logistics support for aircraft components.** It provides for three echelons of interaction (including the depot-to-theater link) and three levels, or indentures, of components (for which demand processes, repair processes, and spares levels may vary). It also contains a submodel that assesses the effect of constrained repair resources. Output reports include capability assessments (for full- and partial-cannibalization assumptions), lists of problem components and subcomponents, depot workload report, and recommended spares levels.

As with other models, Dyna-METRIC cannot represent the real world exactly. Not explicitly modeled (though workarounds exist) are lateral supply, ordering policies for economic order quantity (EOQ) items, flight line constraints, and multiple types of collocated aircraft.

Repair productivity is considered unconstrained unless repair capacities are explicitly stated. The use of repair cycle times instead of computing the dynamic queue for repair reflects an assumption, central to the model, that removals are predictable. In particular, the model assumes that removals are proportional to either sorties or flying hours, and that the mean demand rate is known, as well as the degree of variation about that mean. However, several studies of the demand process (Crawford, 1987; Stevens and Hill, 1973) have shown that these assumptions do not hold—demand levels over time and across bases are extremely variable, and not amenable to prediction.

Even when repair capacities (and thus queues for repair) are explicitly represented, the full effect of uncertainty on performance is not captured, in part because of the inadequate representation of constrained repair in this version of Dyna-METRIC. Also, the value of management adaptations to cope with the uncertainties is imperfectly portrayed. Research is currently underway to develop models that better address these issues.

Version 4 consists of five programs. PART partitions the input data set for processing a handful at a time, thus reducing the amount of computer space needed to run the model. The partitioned data set passes to ECHO for error checking (verifying input values and formats) and summarizing in six tables. PIPE computes the expected pipeline contents. MOD compresses and modifies those pipeline contents. Except for several stockage reports that come from PIPE, REPORT writes the output, including detailed pipeline segment reports, performance reports, problem parts lists, and stockage computations.

Sample analyses with a hypothetical Z-455 spaceship illustrate the use of Version 4. Initially we consider a simple, peacetime, steady-state analysis of a single fleet of spaceships at one base. Then we expand to several fleets at three bases with varying logistics support from other locations. Using this dynamic wartime case, we incorporate each of the input data groups into separate excursions, demonstrating their usage and their effects on performance.

The basic mathematical computation underlying the model is that of the expected number of components being processed by each function and echelon. Dyna-METRIC portrays component support processes as a network of pipelines through which aircraft

components flow as they are repaired or replaced within a single theater. Pipeline segments are characterized by a delay time that arriving parts must spend in the segment before leaving. Some delays (e.g., local repair times) vary by component; others (e.g., base-to-depot transportation time) vary by base. The expected number of components in each segment, then, depends on the rate at which demands occur and the time the components spend in each segment.

Using the sum of all pipeline segments, Dyna-METRIC determines the complete probability distribution for the number of parts in repair and on order. Combining such distributions for all components provides the estimate of aircraft availability and sorties. The probability distributions are also used to compute spares requirements and to identify problem parts.

As this is a technical report, much detail about using the model is provided in appendices. They contain information about compiling, installing, and running the model, as well as list error messages, stop codes, and input formats. They also include modeling tips and an index to the input data fields.

ACKNOWLEDGMENTS

Numerous Air Force and RAND analysts have made contributions to this and earlier versions of Dyna-METRIC. Of particular note in the development of Version 4 was the support received from Curtis Neumann, Michael Niklas, and Barbara Wieland of AFLC/XPS. We also thank Irving Cohen of RAND, whose timely suggestions and criticisms have improved both the model and this report. Without him, there would be no Dyna-METRIC. Finally, our thanks to Manuel Carrillo and Marc Robbins, whose painstaking and careful reviews have made this a better document.

CONTENTS

FIGURES

TABLES

GLOSSARY

AWP (awaiting parts): a repair status indicating that a component's repair cannot continue until one or more serviceable subcomponents become available.

cannibalization: the practice of transferring a serviceable component from one aircraft to repair another. The aircraft must already be unserviceable because of another component failure, and the needed serviceable component cannot be obtained from local supplies.

cirf (centralized intermediate repair facility): a shop that repairs components from one or more remote bases.

component impact: an approximation of the expected number of aircraft rendered NFMC by shortages of a particular LRU, computed by dividing the LRU's expected number of backorders by its QPA.

condemnation: a decision or status indicating a component or subcomponent is irreparably damaged.

constrained repair: a status of repair resources indicating that the repair facility is limited such that an arriving component joins a queue of other components also awaiting service.

ECHO: the error checking and data echo program, the second of the five Dyna-METRIC programs.

FCFS (first come, first served): a repair process whereby the first component arriving for repair is the first to be serviced.

FMC (fully mission capable): an aircraft status indicating that the weapon system can accomplish any of its wartime missions.

lone base: a base that is not supported by a cirf.

LRU (line replaceable unit): a component typically removed from the aircraft at the flight line, rather than in a back shop.

MOD: the pipeline file modifier, the fourth of the five Dyna-METRIC programs.

NFMC (not fully mission capable): an aircraft status indicating that the weapon system's ability to accomplish at least one wartime mission has been degraded.

NRTS (not reparable this station): a decision or status indicating that a component cannot be repaired at a specified facility.

PART: the data set partitioner, the first of the five Dyna-METRIC programs.

peacetime: a period of steady-state activity preceding the wartime scenario, encompassing all time before and including day 0.

PIPE: the pipeline computation, the third of the five Dyna-METRIC programs.

pipeline: a network of repair and transportation processes through which reparable and serviceable parts flow as they are removed from their higher assemblies, repaired, and requisitioned from other points of supply.

pipeline segment: a single process in the pipeline characterized by part arrivals over time, a delay time, and part departures over time.

PMC (partially mission capable): an aircraft status indicating that the weapon system can perform at least one wartime mission, though perhaps in a degraded mode.

QPA (quantity per aircraft): the number of a particular component or subcomponent physically mounted on an aircraft. Not the same as quantity per application (except for LRUs).

REPORT: the report writer, the fifth of the five Dyna-METRIC programs.

SRU (shop replaceable unit): a subcomponent of an LRU, typically removed from the LRU in the shop.

sizing parameter: a constraint to the size of a Dyna-METRIC input data set, set when the model is compiled.

subSRU: a subcomponent of an SRU, including bits and pieces that are often consumed during repair of the SRU. A subSRU may be reparable itself.

VTMR: Variance to mean ratio.

wartime: the (dynamic) scenario of interest that begins on day 1.

I. INTRODUCTION

The Dyna-METRIC model (Hillestad and Carrillo, 1980; Hillestad, 1982; Pyles, 1984) was developed to provide five new kinds of information that logisticians could use to improve wartime logistics support within a single theater:

1. operational performance measures
2. effects of wartime dynamics
3. effects of repair capacity and priority repair
4. problem detection and diagnosis
5. spares requirements.

This report describes the latest implementation of that model in comprehensive detail, including its capabilities, limitations, implementation, use, and mathematical foundations.

Since its development, Dyna-METRIC has been adopted for use as a standard assessment tool within the U.S. Air Force. The Weapon System Management Information System (WSMIS) of the Air Force Logistics Command (AFLC) has used Version 3.04 to perform assessments of planned and actual stock support to single aircraft mission design series in standard, single theater operation plans (DRC, 1983; DRC, 1987). Derivatives of the model have been developed for special purpose, single-unit analyses (Carrillo, 1982; Clarke, 1982; Ogan, 1983; Gage and Ogan, 1984). Similarly, model variations have extended the Dyna-METRIC mathematics to estimate wartime depot workload and stock requirements (Bigelow and Isaacson, 1982) and to estimate worldwide peacetime and wartime stock requirements. But no previous version of the model has been able to conduct worldwide assessments of how the logistics functions and echelon system interact to enhance or constrain wartime capability.

Such analyses are important. Without them, logisticians in different organizations (e.g., base supply, base repair, depot supply, depot repair, depot distribution, and transportation) cannot be sure that their well-intended actions may not actually interact with others' actions to degrade the forces' wartime capabilities. For example, a manager's decision to divert spare parts from one theater to another may only marginally enhance the receiving theater's capabilities but have a disastrous effect on the theater whose parts were diverted. Alternatively, the manager who decides to increase peacetime flying without requisite investments in spare parts repair and procurement may cause the force to degrade its readiness and sustainability by using its wartime stocks excessively in peacetime. When such decisions are made, managers must look beyond the immediate problem—to see the entire logistics support system—to discover whether a proposed cure may not be worse than the disease.

Version 4 of Dyna-METRIC was developed expressly to assess worldwide logistics support for aircraft components, including the depot-theater interactions. Thus one can assess how diverting spares from one theater to another might affect support in both theaters, or how the base/theater/depot repair processes, stock levels, transportation processes, cannibalization[1] policies, and wartime plans interact to affect the forces' combat capability. Moreover, the

[1]Cannibalization is the practice of transferring a serviceable component from one aircraft to another. Such an action is possible only when a serviceable component needed to repair one aircraft cannot be obtained from local supplies and another aircraft is already unserviceable because some other component has failed.

extended model assesses how indentured[2] components' support may affect combat capability by causing or alleviating delays in awaiting parts in the repair process at each echelon.

As indicated in Table 1, the improved model incorporates several other enhancements. Two optional computations were added: the depot workload computation suggested by Bigelow and Isaacson (1982), and the computation of base-level stockage with a no-cannibalization constraint. Also, the computation of fully mission capable aircraft and sorties has been improved to account for the inability of the flight line to cannibalize specific components.

Most of the enhancements, though, provide more detailed representations of the logistics system for individual components, particularly in the areas of:

- component demand processes (permitting time-varying demand factors, sortie- or flying hour-based demands, and onshore and offshore demand factors),
- stock deployments (over time at all echelons and deploying peacetime reparables for repair after some delay),
- repair processes (permitting NRTS[3] actions to occur at the beginning or end of the repair process, different repair times at different echelons, and different repair processes and scope of repair at different echelons).

The enhancements in the representation of constrained repair reflect, though do not satisfy, the need for a more explicit representation of uncertainty in demand and repair, and of the actions management might take to overcome unanticipated demands. As we know from studies of the demand process (Stevens and Hill, 1973; Crawford, 1987), the mean removal rate for components, and the variation about that mean, change over time and are difficult to predict. Undoubtedly, there will be more removals than expected for some components, and fewer for others. Spending more money on more spares is not necessarily a satisfactory protection against uncertainty, because it is difficult to know which components will require the extra spares. A better strategy for coping with uncertainty might be to use more flexible resources such as repair. The constrained repair submodel is an initial attempt to analyze the effect of repair and the adaptations available to management in the use of repair.

Prototype models in research use at RAND, developed under the so-called "Uncertainty Project," more adequately represent uncertainty (which includes an increase in the variance to mean ratio, or VTMR) and a fairly wide range of management adaptations (Isaacson and Boren, 1988). It is expected that subsequent versions of Dyna-METRIC will include such capabilities. In the meanwhile, Version 4 does permit the use of VTMRs greater than one and the priority repair management adaptation that applies in particular circumstances. Users are encouraged to experiment with the use of these capabilities. Better representations are likely to emerge eventually from RAND and other research institutions.

This report contains information about the model that may be of interest to various audiences. Accordingly, different readers will find some parts more useful than others. Section II describes the model's general capabilities and limitations so logisticians can understand the model's applicability to their specific problems. Section III describes the five programs that make up the model. Section IV is a user's guide, written from the viewpoint of the hands-on user, who must understand the model's operation in detail. Section V contains the mathematics of the model.

---

[2]Indentured components are subcomponents of other components. They are often *consumable* parts, meaning they are replaced in the repair of other higher order components, or they may be reparable themselves.

[3]Not reparable this station, an indication that a repair facility was unable to repair a given component, typically because of technical limitations.

Table 1

DIFFERENCES BETWEEN VERSION 3.04 AND VERSION 4

Version 4 better suited for conducting worldwide analyses:
- Multiple depots explicitly modeled
- Multiple identical bases need be entered (and analyzed) only once
- A larger number of components may be simultaneously analyzed
- Automatic time-scaling, so that arbitrarily long scenarios may be analyzed without recompiling

Version 4 analyzes more types of components:
- A level of indenture below SRUs called subSRUs
- NRTS and condemnation decisions may be before or after repair
- SRU repair may proceed in parallel with LRU repair
- Maintenance types besides RR and RRR (for example, JEIM)
- Sustained demand rates
- Demand rates per flying hour or per sortie

Version 4 analyzes more types of aircraft:
- Base-dependent quantities per aircraft
- Can specify how many of a multiple quantity LRU must be working for the aircraft to be mission capable
- Onshore and offshore bases (with different demand rates)

More detail with Version 4:
- Option so that unachievable sorties not flown (though sorties on partially mission capable aircraft are still flown)
- Maximum turn rate records defined by base
- Peacetime reparables' arrival time explicitly modeled
- Separate base, cirf, and depot repair times and NRTS rates
- Condemnations modeled for bases, cirfs, and depots
- SRU depot replacement fractions
- Level of repair may be stated for all components
- Optional LRU-dependent flying hour programs
- Cutoffs forward only (as in Version 3.04), or forward and retrograde

Version 4 has expanded report capabilities:
- Simplified problem parts reports, including problem subcomponents
- Performance for partial cannibalization of LRUs
- Demands on repair/supply at each echelon on each time of analysis

Version 4 has expanded constrained repair capabilities:
- Constrained repair at all echelons (rather than the highest)
- Separate wartime depot daily production limit that may be optionally based on peacetime depot daily production
- LRUs may be assigned to multiple types of constrained repair

Differences in stockage algorithms:
- Version 4 can stock LRUs to a no-cannibalization criterion
- Version 4 does not do tradeoffs between SRUs (Version 3.04 did so by pretending the SRUs were LRUs)
- Version 4 algorithms for purchase of LRUs do not assume SRUs are fully cannibalizable unless the user has so specified

II. SCOPE, CAPABILITIES, AND LIMITATIONS

This section presents the scope of the model in terms of the operations and logistics system for aircraft components. Although this material focuses on the Air Force world and deals with aircraft, Dyna-METRIC has been used to analyze a variety of weapon systems, including helicopters and tanks.

This discussion continues with how Dyna-METRIC represents that logistics system. It describes the model's capabilities (in assessing the system's performance in a dynamic wartime scenario and in helping identify factors that may limit performance) and limitations (arising from the assumptions and approximations that underlie the model's mathematics).

SCOPE

Dyna-METRIC models one or more types of aircraft at one or more bases located in one or more theaters of operation for a period of time that may range from several days to several years. The model predicts the effect of the logistics support system on the bases' ability to execute their assigned flying programs.

Bases are the operating locations to which aircraft are assigned, whether they be permanent bases or temporary facilities to which aircraft are deployed. The scenario commences with each base's aircraft either *fully mission capable* (FMC) or degraded by some flying program that we assume has been underway for a substantial period. Aircraft can operate out of a base on a fly-out, fly-back sortie program (as fighter aircraft typically do), or on a fly-in, fly-out program (as with cargo aircraft flying a circuit). In either case, broken parts arrive with incoming planes; in the case of cargo aircraft, removals of failed components may be more likely at some bases than others.

Different bases can support different types of aircraft, but any particular base supports only one type—aircraft are assumed to be identical within a base. However, a real base supporting several aircraft types can be modeled by locating each type at different bases served by a common centralized intermediate repair facility (cirf).

Aircraft can be flown on different missions at different times. For example, a base may fly air-to-air missions for some initial period and subsequently fly ground attack missions.

The flying programs to be executed by the various bases may vary over time. The number of aircraft can increase with the deployment of new units and decrease because of attrition or the reassignment of aircraft. The number and length of sorties may vary each day, as can the maximum single aircraft sortie rate, which limits the number of sorties that one operational aircraft can fly in a single day. With this flexibility, the model can accommodate most conceivable flying programs, including the peacetime and wartime scenarios.

The Aircraft

Aircraft are assumed to have an indentured component structure: they are composed of *line replaceable units* (LRUs) that are composed of *shop replaceable units* (SRUs) that are composed of what we call *subSRUs*. SubSRUs include bits and pieces that are consumed during SRU repair as well as other components that may be repaired either locally or at a higher echelon.

Dyna-METRIC sees the entire aircraft as a collection of LRUs. Certain major aircraft components, such as engines, are generally not called LRUs, but these components should be treated like LRUs in the model.

Aircraft availability is modeled as a direct function of the availability of the aircraft's LRUs. SRUs affect availability only through their own availability to support the repair of their parent LRUs, and subSRUs are one step further removed in their effect on aircraft availability.

A given LRU may appear on an aircraft one or more times. Any shortage or hole of an LRU may ground the aircraft. If several of an LRU are on a plane, they can all be classified as essential or they may be classified as wholly or partially redundant, in which case more than one unit must fail before the aircraft is rendered *not fully mission capable* (NFMC).

LRUs may also be classified as essential or nonessential to a particular mission that the aircraft can execute. For example, a plane with a broken radar unit might be incapable of executing an air-to-air mission but be capable of ground attack.

Finally, the model accommodates the possibility that there may be limited differences in the components on the aircraft at a single base. This situation may arise when components are being phased in or out, or when some of the aircraft are specially equipped.

The Logistics System

The logistics component support system is assumed to be a five-echelon hierarchical structure:

1. flight lines (from which aircraft operate) are supported by
2. local base repair shops, which are supported by
3. cirfs, which are supported by
4. depots, which may be supported by
5. various suppliers of components.

Not all echelons are required. For example, omitting cirfs connects bases directly with depots, as shown in Fig. 1.

Reparable components essentially move upward in this hierarchy. Reparable parts are removed from the aircraft at the flight line and are serviced at base level. If not reparable there, they are transported to a cirf and serviced; if not repaired there, they are sent on to the depot. Parts at any level can also be *condemned* as irreparably damaged; thus, consumable items may be analyzed by assigning them condemnation rates of 100 percent. Stocks of serviceable spare parts may be held at any level, and over time these spares are sent down the hierarchy to replace the reparable ones that have been sent up.

The repair capabilities of each level can be modeled in considerable detail. Repair for LRUs can be *unconstrained*, where maintenance is assumed to begin as soon as a component arrives at a repair facility, or *constrained*, where an arriving component joins a queue of other components also awaiting service. Components are selected from the queue based on a priority scheme that minimizes maximum backorders rather than first come, first served. How long a component waits for service depends on how many aircraft are NFMC relative to other components and on how heavily loaded the repair facility is.

Repair capabilities at the various levels may change over time. This facilitates the modeling of deploying units and the mobilization of repair facilities for wartime operations.

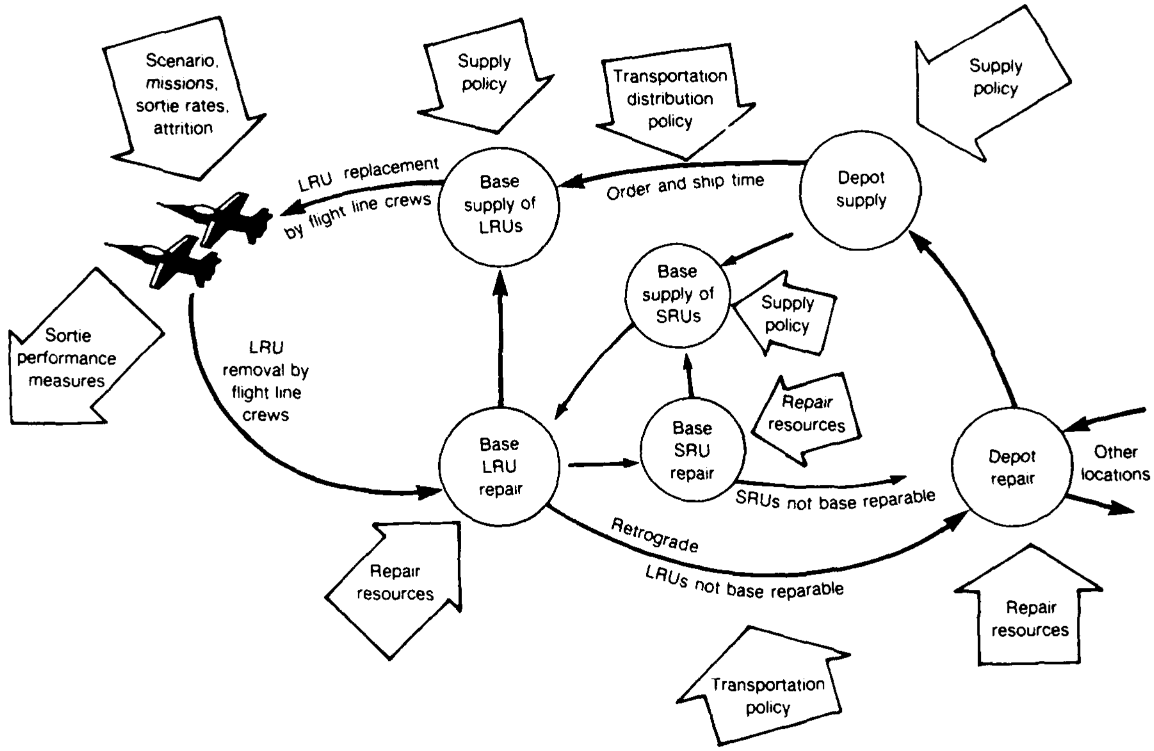

Fig. 1—Aircraft Logistics Support Network

Inherent in Version 4 is a *worldwide* perspective. Demands for repair at several depots may arrive from several types of aircraft flying in several different theaters. The depots' repair facilities and spare stock pools are effectively shared by all forces. In general, the model resolves allocation questions on a earliest requisition basis.

How the Model Represents the Logistics System

Dyna-METRIC portrays component support processes as a network of pipelines through which components flow as they are repaired or replaced. Arriving components must spend a (random or deterministic) delay time in each pipeline segment. Delay times, such as local repair times, vary by component; others, such as intratheater transportation times, depend on the base being assessed. Further, components may be "frozen" in their segments and not flow, as in the transportation segment when a transportation cutoff is in effect.

Failed components enter the pipeline network at the bases' flight lines. Each base has several aircraft, the use of which generates component failures and thus demands for replacement components. Each base has a flight line support capability that removes and replaces those components, drawing serviceable spares from local supply as needed to repair aircraft.

A base may also have component repair shops that test failed components and return them to serviceable condition. For units deploying to new bases, that repair capability may be available only after some delay, while the repair facility is being deployed and set up.

Once an LRU has been removed from the aircraft, it is either repaired or condemned (in which case a replacement is ordered from the next higher echelon). Items repaired at local shops are returned to local stock. Items unable to be repaired locally are declared NRTS. A base sends NRTS items to either the base's cirf or the component's depot for repair while ordering replacements from there; a cirf sends NRTS items to a depot and requisitions it for replacements; a depot cannot send NRTS items anywhere, so it condemns them and orders replacements from an outside supplier (unless resupply is cut off in the scenario). Upon receiving a requisition, the repair facility immediately sends a serviceable spare if available, or sends one as soon as possible after all previous requisitions for the same component have been filled. Once the component arrives for repair, it is repaired and returned to that facility's own stock so that it can be issued to satisfy the next demand.

As LRUs are processed at the various facilities, failed SRUs may be discovered. The SRU repair and resupply network is essentially the same as that for LRUs, as is the repair and resupply network for subSRUs.

The key equation in Dyna-METRIC computes each component's expected pipeline contents—the expected number of each component that will be in each segment of the pipeline network. The computation is based on the planned time-dependent aircraft flying activity or (optionally) the achievable *partially mission capable* (PMC) and FMC time-dependent aircraft flying activity. The model computes the flying-dependent removals caused by that activity and then, using the time-dependent availability and delays associated with transportation and repair at bases, cirfs, and depots, and the likelihood that the component will be NRTSed or condemned, determines the expected contents of each pipeline segment. The segments are then totaled to forecast the total pipeline size (the expected quantity on order and in local repair) as seen by each base.

The expected total pipeline size is the key parameter for a probability distribution that describes the number of components in the network, as seen at each base's flight line. That is, the expected total pipeline size is used to determine the probability that there is one component removed from the base's stock or its aircraft, the probability that there are two components, the probability that there are three components, and so on (Hillestad and Carrillo, 1980; Hillestad, 1982).

If the operational demands and the logistics system's characteristics were constant over time, the estimation of the expected pipeline quantity and the probability distribution would be easy. Assuming that individual component removals are independent of each other (removing one component neither causes nor prevents the removal of another), Palm (1938) showed that the pipeline quantity has a Poisson probability distribution, with a mean equal to the average failure rate times the average repair time.

Unfortunately, operations and logistics seldom achieve steady-state, especially in wartime. Not only do the operational demands change during wartime, but the logistics system restructures itself. Stock is redeployed, repair equipment and personnel are redeployed, and transportation is reallocated over time. Thus the steady-state equations do not accurately forecast pipeline quantities in a dynamic wartime scenario.

Hillestad and Carrillo (1980) demonstrated that Palm's result could be extended to the dynamic wartime situation. In their formulation, the time-dependent component removals due to operational demands (e.g., daily demands over some time interval) are combined with the time-dependent repair or transportation capability (e.g., the probability that an item entering the pipeline segment at time s will still be in the pipeline segment at time t) to estimate the expected pipeline quantity size over time. They also extended Palm's original result to show that the pipeline distribution would be Poisson—even under conditions of time-varying demands and repair.

Dyna-METRIC combines each component's dynamic demands and repair process times to estimate the expected pipeline quantity for each pipeline segment. The dynamic demands for segments after the first segment are derived from the dynamic departures from the preceding segment. For example, the number of LRUs entering the base-to-cirf pipeline is the number of departures from the base repair segment times the NRTS rate. The model computes expected pipeline quantities for each item's base, cirf, and depot repair segments and for the transportation segments between these locations. SRUs *awaiting parts* (AWP) at each location are computed from the number of subSRUs in stock and under repair, and (likewise) AWP LRUs are computed from SRUs in stock, in repair, and AWP. Backorders at depots and cirfs are computed from quantities in stock, in repair, AWP, and on order. Those backorders are allocated to bases using a first come, first served rule. The expected base pipeline for components, then, consists of items in local repair and on order from higher echelons (in transit and backordered).

CAPABILITIES

The model's capabilities fall into three categories: performance measures that assess the logistics system, diagnostics, and spares requirements computations.

Performance Measures

Given descriptions of the scenario, aircraft, and logistics system, Dyna-METRIC provides various measures of performance. A performance shortfall generally cannot be assigned to a single cause or solution. Poor component reliability, slow or inadequate repair capabilities, ineffective transportation, and too limited a supply of spares may combine to cause the problem; and the solution may require changes in any or all of these areas.

Besides traditional component-oriented logistics statistics such as backorders, Dyna-METRIC provides higher combat-oriented capability measures related to the force's ability to generate sorties. The combat measures include aircraft availability and daily sortie generation capability. For each base, the model reports the expected number of available aircraft at any specified time and confidence level. For example, Dyna-METRIC might report that on day 5 of a scenario, a given base could expect (on average) 16 available aircraft but only 13 available with 95 percent confidence.

The model also estimates the expected number of sorties a base can generate on any specified day. It assumes that a base never overflies the program specified in the scenario (though the base may fail to achieve its program because of a shortage of available aircraft), so the predicted sortie generation capability is always less than or equal to the scenario's flying program. Thus the model's daily sortie estimates reflect both requested sorties and available aircraft.

Higher order performance measures are quite sensitive to whether LRUs can be cannibalized from one aircraft to repair another. Aircraft availability and sortie generation are typically much higher under a full-cannibalization policy than under one of no cannibalization.

The model allows the user to label each LRU as cannibalizable or not, and then computes aircraft availability and sortie generation first by using this information, then assuming a policy of full cannibalization.

From the expected base pipeline value, the model derives the probability that a given number of components are in repair or on order at each base. Using these total pipeline probability distributions for each component, and the component's available stock at each base, the model next forecasts how the LRUs in repair and on order would (probabilistically) generate backorders (or aircraft holes) for each component at a given time. It then distributes those holes across aircraft for the two cannibalization policies. For full cannibalization, Dyna-METRIC assumes that all component holes at each base are instantly consolidated on the fewest possible aircraft, making as many FMC aircraft as possible. For partial cannibalization, holes of LRUs flagged as not cannibalizable are assumed to occur randomly across the aircraft at each base. Holes of cannibalizable LRUs are then consolidated onto the aircraft that are already down for noncannibalizable LRUs. Leftover holes are consolidated onto as few of the remaining aircraft as possible. In each case, Dyna-METRIC derives a full probability distribution of the number of degraded aircraft from which the capability assessment report is directly obtained--in particular, the expected number of NFMC aircraft and the expected number of FMC sorties.

Diagnostics

Dyna-METRIC has a special report of particular interest when the projected performance is unsatisfactory. It identifies those LRUs that are most likely to be a problem for at least one base and sorts them by the number of aircraft they are likely to ground. For these LRUs, it reports:

- The number of aircraft they will probably ground.
- The number of aircraft they would ground if the base-level spares were most effectively redistributed.
- Where in the logistics system the LRUs are tied up (e.g., queued for repair at the cirf, in transit from the depot, awaiting serviceable SRUs at a base).
- Which SRUs and subSRUs are tied up and where, if they limit LRU availability.

Requirements Computations

Three requirements computations are incorporated in the model. The stockage algorithm optionally computes stock with single-component fill rate goals, or with FMC aircraft goals for full- or no-cannibalization policies. The depot workload requirement computes the maximum and minimum workload necessary for a depot surge to meet its expected requisition levels for each component.

The pipeline probability distributions are used to compute stockage requirements. In this mode, Dyna-METRIC recommends additional LRU, SRU, and subSRU stock to achieve an FMC goal at the lowest cost. Two general strategies are employed: buying spares to assure that each component *individually* achieves a target FMC goal (disregarding other components), or buying spares so that all LRUs *jointly* achieve the FMC goal. The first strategy *falls short* of the goal of the second.

To assure only that each LRU does not violate the FMC goal with the stated confidence level, the model uses the LRU's individual pipeline probability distributions to increase the

LRU's stock level until the stated confidence is achieved for that component alone. But to assure that all LRUs jointly achieve FMC aircraft goals (regardless of cannibalization policy), the model first makes sure that each LRU achieves the goal individually, then it buys more LRUs across the full range of parts to achieve the overall goal. This latter computation employs a marginal analysis technique: It determines how much closer to the goal it would be with an additional unit of LRU 1, or of LRU 2, etc.; then it adds a unit of the LRU that had the best benefit/cost ratio. It continues to add LRUs in this manner until attaining the goal.

The stockage algorithms assume that initial stock levels entered by the user represent a sunk cost. Thus existing stock is retained throughout the analysis, and only marginal stock additions are made to improve performance. If these algorithms are run with no input stock, the actual stock mix (and costs) would probably be different.

For LRUs at cirfs and depots, and for SRUs and subSRUs anywhere, the model adds enough stock for each part to individually achieve its target FMC goal. The number of aircraft associated with a cirf is the sum of the aircraft at the cirf's bases. The number of aircraft associated with a depot is the total number of aircraft at all its bases. For each component and subcomponent, the model adds enough stock so that the probability of achieving the FMC goal is greater than the requested confidence level. Stock levels for SRUs, subSRUs, and cirf and depot LRUs are computed before stock levels for base LRUs are computed. This reduces the number of on-order and AWP LRUs everywhere in the system, as well as the requirements for base LRU stock to offset excess backorders and AWP LRUs.

Finally, Dyna-METRIC can compute the maximum possible wartime depot repair workload (the expected daily arrivals for depot repair), the minimum required wartime depot workload (the minimum number of LRUs that must be inducted daily into depot repair to satisfy expected depot requisitions), and the amount of LRU stock needed at the depot to offset repair and retrograde transportation delays under dynamic wartime conditions.

LIMITATIONS

Dyna-METRIC's representation of the real world is limited by its mathematical assumptions, use of approximations, and program implementation constraints. Overcoming the limitations that are due to mathematical assumptions would require new mathematical breakthroughs. Overcoming limitations that are due to approximations could be done, but that would dramatically affect the model's computational resources (CPU time and memory). As for implementation constraints, computers have always had limits on their ability to evaluate the equations essential to mathematical models.

Unconstrained Repair

In Dyna-METRIC's simplest uses where constrained repair is not modeled, the mathematics underlying the model require two key assumptions about demands, transportation, and repair processes:

1. Demands for repair are proportional to flying hours or sorties and arrive randomly according to one of two well-known arrival probability distributions: Poisson or negative binomial. Further, the time-dependent mean and variance of that distribution is known.

2. Repair and transportation times have known probability distributions that are independent of the demand history.

These assumptions permit the computation of the pipeline probability distributions. However, neither assumption is always true. Several studies of the demand process (Stevens and Hill, 1973; Crawford, 1987) have noted that demand levels over time and across bases are extremely variable and difficult to predict. Further, a shop's fixed repair resources (manpower, equipment, etc.) have the potential both to become overloaded if demands surge and to speed up repair for some critical components by delaying repair of less critical components.

These assumptions may cause the model either to underestimate or to overestimate the logistics system performance if repair resources are not explicitly modeled. If the demand and repair processes do not deviate radically from these assumptions, Dyna-METRIC should be fairly accurate. Otherwise, one should supplement unconstrained repair analyses with analyses where the repair resources are explicitly modeled.

Lateral Supply

The assumption that demands, repair, and resupply functions are independent also prevents Dyna-METRIC from directly assessing the effects of lateral supply across bases. Essentially, lateral supply would have the same effect as expedited resupply from a higher echelon. Because the effective resupply time would depend on the history of demands, repairs, and resupplied items, lateral supply violates the model's underlying mathematical assumptions.

A workaround exists for this situation. If cirfs are not being used in an analysis, one can model several related bases as being supported by a cirf. Then some of the theater's stock can be relocated to the cirf to be requisitioned and shared across all the bases to simulate lateral supply. However, the cirf will "ship" to satisfy the oldest requisition, rather than the most urgent, so this workaround does not capture much of the benefit of a responsive intratheater support system.

Different Aircraft Types

To compute the effect of the component pipeline probability distributions on FMC aircraft availability, Dyna-METRIC assumes that the aircraft deployed at each base are nearly identical. That assumption is critical to the computation of both full and partial cannibalization of FMC aircraft, so it affects the reported available FMC aircraft, sorties achieved, and the stock needed to achieve FMC aircraft goals.

Again, a workaround exists if the cirf feature is not being used in the analysis. One can represent each real base with multiple aircraft types as several bases with a common cirf housing the base's stocks for all aircraft. By setting the transportation times between the base and cirf to zero, one can assess how both unique and common components' support affects the capabilities of multiple aircraft types.

Constrained Repair

Dyna-METRIC uses a deterministic, expected value computation to compute the expected pipelines for constrained priority repair, so it only approximates the real-world random arrival and random repair processes. Further, it applies the resulting component pipeline distributions as though they were independent. However, the component's pipeline distributions cannot be independent because the repair capacity constraints cause LRUs to undergo additional queueing (because of other LRUs' demands), and priority repair attempts to equalize all component's pipelines. The constrained repair computations only approximate probable logistics system performance under quite restrictive conditions:

- One component's demands dominate all others, and
- All components' combined demands exceed the shop's capacity.

Such conditions assure that the independence assumption and the stochastic arrival and repair processes are irrelevant.

High combined demands occur frequently in wartime assessments when flying activities surge dramatically. However, high demand patterns seldom occur in peacetime, and peacetime assessments would probably underestimate the effect of random arrival patterns on queueing. Dyna-METRIC provides an optimistic assessment of performance of peacetime queueing.

Typically, one component's demand rates greatly exceed those of all others repaired by the same resources. That component usually has the largest pipeline, the greatest backorders, and highest repair priority. When one component always has the greatest backorders, most of the shop's capacity is devoted to repairing it; variations in the pipelines for other components would not affect FMC aircraft (assuming cannibalization). If some components' backorders grow until they nearly match those of the worst component, the model provides an overly pessimistic assessment of performance. The problem LRUs report detects such a pessimistic assessment. In either case, one should use a detailed stochastic discrete event simulation like Dyna-SIM (Miller, Stanton, and Crawford, 1984) to estimate the effects of unsaturated repair or roughly balanced priority repair.

Ordering Policies for EOQ Items

Some spare parts are so small or inexpensive that they are ordered in quantities greater than one at a time. Dyna-METRIC's mathematics apply precisely only to those cases where the order quantity is one. The mathematics are only approximately accurate for larger economic order quantity (EOQ) policies. As the order quantity increases, the pipeline variability also increases.

An approximate workaround is to increase the demand variance-to-mean ratio proportional to the square root of the order quantity. The pipeline variability then reflects the expected variability due to the order quantity.

Approximate Additive Pipelines

For computational efficiency, Dyna-METRIC does not compute the joint probabilistic effects of backorders and AWP quantities with related pipelines. Rather, the expected values of these quantities are added to the appropriate pipelines as though they were also Poisson or negative binomial distributions.

Of course, this is not rigorously correct,[1] but it has been satisfactory in practice. Tests of the approximation show that only modest errors are introduced in the computations of total base component breakdowns or NFMC aircraft when the expected backorders or AWP quantities are small (less than one). As these quantities increase, the errors appear to decrease.

[1]To treat this rigorously, the model must convolve the related probability distributions, a task that would greatly increase computer time.

Flight Line and Operational Constraints

Flight line resources and operational constraints affect the sortie rates that an FMC aircraft can achieve. These factors are beyond the scope of the Dyna-METRIC model, so they do not appear explicitly. Nevertheless, their effects can be estimated in other models or analyses and incorporated into Dyna-METRIC's sortie rate parameter. Thus one can estimate how those processes affect wartime sorties.

Precision and Accuracy

Unlike the mathematics upon which it is based, the computerized model cannot carry out its computations with infinite precision. It must contend with the fact that computers cannot represent numbers that are very large or very small. Dyna-METRIC computes extremely small probabilities and multiplies or sums them in various ways. Often a computed probability is smaller than the computer can represent. Summing these small numbers (or actually zeroes) leads to numeric instabilities that may affect results.

Dyna-METRIC partially compensates for this effect by using logarithms where possible. Generally, this problem arises only with consumables that have extremely high demand rates (several demands per sortie) and very long resupply times. Removing the offending component from the analysis usually corrects the problem. Such components are analyzed more appropriately outside the rigorous confines of Dyna-METRIC.

III. PROGRAM DESCRIPTION

Dyna-METRIC Version 4 consists of five computer programs, varying in length from 1000 lines of FORTRAN code to nearly 20,000 lines:

1. PART, the data set partitioner;
2. ECHO, the error checking and echo program;
3. PIPE, the pipeline program;
4. MOD, the pipeline data file modification program; and
5. REPORT, the report writer.

This section discusses the purposes of each program and describes how the programs operate and interface with each other. Readers requiring more detailed information on how to actually run the model should consult App. A (file requirements).

PART: DATA SET PARTITIONER

Dyna-METRIC Version 4 handles a large number of components by processing them a handful at a time. This reduces the amount of computer space required to run the model because the only data necessary at any given time are those describing the partition of components currently being processed.

The constraints that determine which components must share a partition are fairly simple. All SRUs indentured to an LRU, and all subSRUs indentured to those SRUs, must be in the same partition as the parent LRU. If two LRUs are assigned to the same type of constrained repair, they must also be in the same partition.

When PART is compiled, limits are specified for the maximum number of LRUs, SRUs, and subSRUs that may share a partition.[1] PART examines the Dyna-METRIC input data set and divides the components into partitions, trying to violate neither the compiled size constraints nor the relationship constraints specified in the previous paragraph. If this is impossible (e.g., an LRU has 50 indentured SRUs, but the limit on SRUs per partition is set to 30), PART satisfies the relationship constraints and violates the size constraints, writing a message specifying the smallest size constraints with which it could comply. The user must either recompile the model to allow sufficiently large partitions or delete a sufficient number of components to fit the existing size constraints.

The program works as follows. Administrative and scenario data are read from the unpartitioned data set and written to the partitioned data set. PART skims the remainder of the unpartitioned data set to determine which SRUs are indentured to which LRUs, which subSRUs are indentured to which SRUs, and which types of LRUs have been assigned to which types of constrained repair.

For each type of constrained repair, the assigned LRUs (and their SRUs and subSRUs) are placed in the same partition. If an LRU is assigned to multiple types of constrained repair, the associated partitions are merged. For each partition so determined, the records describing the LRUs, SRUs, and subSRUs are written to the partitioned data set.

[1] Appendix B describes the model's sizing parameters.

After dealing with LRUs that are assigned to constrained repair, the program proceeds to the remaining LRUs. If an LRU has too many SRUs, or if there are too many subSRUs indentured to those SRUs, the LRU, SRUs, and subSRUs are placed in a partition by themselves, their records are written to the partitioned data set, and a warning message is written. If the LRU and its subcomponents do not violate the size constraints, PART identifies additional LRUs (with their SRUs and subSRUs) that may share the partition without violating the size constraints. When no further such additional LRUs can be identified, the partition is full. PART writes out the associated records to the partitioned data set and proceeds to the next LRU that has not yet been assigned to a partition.

Because partitioning can be a fairly time consuming process, the partitioned data set is frequently saved and modified directly when any of the original parameters require alteration.

ECHO: ERROR CHECKING AND ECHO PROGRAM

ECHO examines the partitioned Dyna-METRIC data set produced by PART to ensure that all model parameter values are within valid ranges. It also echoes the input data in easily read tables, writing warning messages as errors are found. After writing the data tables, the program either ends with no stop code (implying no errors were found) or with a stop code of 20 (meaning at least one error was identified).[2]

PIPE: PIPELINE PROGRAM

PIPE reads the partitioned Dyna-METRIC data set produced by PART and checked for errors by ECHO. It writes a file of expected pipeline contents that is compressed by MOD and then used by REPORT to write the standard Dyna-METRIC reports.

PIPE also computes SRU and subSRU stock levels for all locations and LRU stock levels for cirfs and depots. Optional reports written by this program (rather than by REPORT) include the computed stock levels just mentioned, the cost of that stock, a table of daily demands on repair and supply at all echelons, and a report on minimum and maximum depot workloads.

Computing pipeline contents is particularly involved. PIPE computes pipelines in the following order:

LRUs at the base. Compute the administrative and repair segments, saving the demands for cirf or depot repair (due to NRTS actions) and the demands on cirf or depot supply (due to NRTS actions and condemnations). Save data on the discoveries of failed SRUs.

LRUs at the cirf. Compute the retrograde segment to the cirf (based on demands for cirf repair generated by the bases), the cirf administrative segment, and cirf repair segment (saving the demands on depot repair and depot supply). Save data on the discoveries of failed SRUs.

LRUs at the depot. Compute the retrograde segment to the depot (based on demands for depot repair generated by the bases and cirfs), then the depot administrative, repair, and on-order segments. Save data on the discoveries of failed SRUs.

SRUs at the base. Compute the repair segment (based on the discoveries of failed SRUs at the base). Save the demands for cirf or depot repair and the demands on cirf or depot supply. Save data on the discoveries of failed subSRUs.

---
[2]Appendix C describes error and warning messages. Appendix D lists stop codes.

SRUs at the cirf. Compute the retrograde segment to the cirf (based on demands for cirf repair generated by the bases), the cirf administrative segment, and cirf repair segment (based on the same demands and on the failed SRUs discovered at the cirf), saving the demands on depot repair and depot supply. Save data on the discoveries of failed subSRUs.

SRUs at the depot. Compute the retrograde segment to the depot (based on demands for depot repair generated by the bases and cirfs), then the depot administrative, repair (based on the same demands and on failed SRUs discovered at the depot), and on-order segments. Save data on the discoveries of failed subSRUs.

SubSRUs at the base. Compute the repair segment (based on the discoveries of failed subSRUs at the base). Save the demands for cirf or depot repair and the demands on cirf or depot supply.

SubSRUs at the cirf. Compute the retrograde segment to the cirf (based on the demands for cirf repair generated by the bases), the cirf administrative segment, and cirf repair segment (based on the same demands and on failed subSRUs discovered at the cirf), saving the demands on depot repair and depot supply.

SubSRUs at the depot. Compute the retrograde segment to the depot (based on demands for depot repair generated by the bases and cirfs), then the depot administrative, repair (based on the same demands and on failed subSRUs discovered at the depot), and on-order segments.

Depot subSRU backorders. Compute the depot subSRU backorders, determining the number that affect SRUs at the depot and the number that affect subSRU orders from bases and cirfs. Compute the on-order subSRU segments for cirfs and lone bases.

Cirf subSRU backorders. Compute the cirf subSRU backorders, determining the number that affect SRUs at the cirf and the number that affect subSRU orders from the cirf's bases. Compute the on-order subSRU segments for bases served by cirfs.

AWP SRUs. Determine how many SRUs are AWP at each location.

Depot SRU backorders. Compute the depot SRU backorders, determining the number that affect LRUs at the depot and the number that affect SRU orders from bases and cirfs. Compute the on-order SRU segments for cirfs and lone bases.

Cirf SRU backorders. Compute the cirf SRU backorders, determining the number that affect LRUs at the cirf and the number that affect SRU orders from the cirf's bases. Compute the on-order SRU segments for bases served by cirfs.

AWP LRUs. Determine how many LRUs are AWP at each location.

Depot LRU backorders. Compute the depot LRU backorders, determining the on-order LRU segments for cirfs and lone bases.

Cirf LRU backorders. Compute the cirf LRU backorders, determining the on-order LRU segments for bases served by cirfs.

After doing the computations for all LRUs in the partition not assigned to constrained repair, the model turns to the case of constrained repair. Here, LRUs sharing a repair resource must be processed simultaneously because what happens to one type of LRU probably affects the others. Once the repair pipelines are computed, each LRU is treated independently, as were the earlier LRUs.

After all partitions of component data are processed, PIPE is finished with its basic computations. If stockage options were selected, it reports the value of the additional stock. The file of pipeline values it has generated is saved and processed by MOD for use by REPORT, which writes the standard Dyna-METRIC reports.

MOD: PIPELINE DATA MODIFIER

MOD has several functions. It compresses and modifies the pipeline file created by PIPE and produces the specially formatted file of pipeline values used to restart PIPE (based on either the pipeline file or a user-supplied data set of pipeline values).

Its operation is directed by options 16 and 22. If neither option is selected, the only operation performed is the compression of the pipeline file. Option 16 prepares an additional file of pipeline values in a special format to reinitialize PIPE. Option 22 stops the pipeline file from being compressed; instead it constructs a file of pipeline initialization data based on user-supplied data.

REPORT: REPORT WRITER

This program writes most of the Dyna-METRIC reports, based on pipeline data originally computed by PIPE and then compressed by MOD. These include performance reports (for each base at each time of analysis), lists of problem components, and detailed pipeline segment reports. In addition, REPORT computes economical base LRU stock levels to achieve full-cannibalization or no-cannibalization aircraft availability goals.

The major part of REPORT is the actual preparation and writing of the output reports. Reports are prepared in the following order for each time of analysis:

Detailed pipeline segment report. Option 15 gives the expected disposition of all components at each location.

Performance report. Option 11 uses the backorder probability distributions for each LRU to construct the probability distributions for available aircraft under both full cannibalization and partial cannibalization policies. From the available aircraft distributions, performance measures (e.g., expected unavailable aircraft) are directly derived.

Problem LRUs list. Option 8 selects LRUs with a sufficiently high probability of being a problem (or actually, an insufficiently low probability of *not* being a problem). REPORT sorts them by their effect on aircraft availability, and writes a table of actual-vs.-target availability. For the worst 20 LRUs, it graphs each component's effect on availability. It also produces a detailed pipeline segment report for all the problem LRUs, their worst SRUs, and the worst subSRUs indentured to those SRUs.

Compute base LRU stock levels. Selected with option 3, 4, or 17. First, if option 3 is chosen, REPORT buys stock to meet individual LRU readiness goals. Next, if option 4 is selected, it adds stock using marginal analysis until a full-cannibalization aircraft availability goal is achieved. Then, if option 17 is selected, it uses marginal analysis to add stock until a no-cannibalization aircraft availability goal is met. It also determines the value of the additional stock.

Recompute performance based on new stock levels. Selected with option 12 and any stockage option. The problem LRUs list is also redone if option 8 is also selected.

After writing the reports for all times of analysis, REPORT is basically done. The value of additional stock is reported for all stockage options. If option 9 is selected, REPORT writes out the final base LRU stock levels.

IV. USER'S GUIDE

This section reviews the various uses of Dyna-METRIC with an emphasis on input data requirements and output reports. Throughout we employ a small-scale example to serve as an illustrative device. Initially this example is simple and straightforward, but it increases in complexity as additional topics are introduced. The data set structure, data requirements, and output reports come to resemble those of a fairly elaborate modeling exercise.

While reading this section, the user should refer to the input specifications in App. E for the precise location of each data field.

THE PEACETIME/WARTIME TRANSITION

One of Dyna-METRIC's key attributes is its representation of nonstationary (time-dependent) processes. This is especially useful when modeling the transition from a steady-state peacetime environment to a dynamic wartime environment. By convention, *peacetime* refers to an infinitely long period of steady-state activity preceding the (usually dynamic) *wartime* scenario of interest. If time in the scenario is measured in terms of days, then peacetime encompasses all time before and including day 0, while wartime begins on day 1 and continues thereafter.

PEACETIME STEADY-STATE EXAMPLE

Consider the purely hypothetical case of the Z-455 Prospector, a futuristic spaceship engaged in collecting asteroidal debris for sale to the Earth-bound as souvenirs. Highly capable, and yet elegantly simple in design, the Z-455 is composed of only ten LRUs:

1. navigational computer
2. maneuvering thruster
3. fuel cell
4. docking mechanism
5. targeting computer
6. asteroid blaster
7. particle collector
8. sensor
9. control panel
10. cellular mobile telephone.

Except for the telephone (no two of which are ever quite the same), these LRUs are shining examples of modern, modular construction and are freely interchangeable among different ships.

Because of their distant field of operation, it is impractical to station Z-455s at Earth facilities. Instead, they are based on an orbital Materials Processing Center (called MPC) located on the innermost edge of the asteroid belt. In addition to serving as a temporary repository for accumulated asteroidal debris, MPC provides logistics support for its Z-455s in the form of limited component repair, stockpiles of spare parts, and resupply linkages to Earth.

This initial analysis evaluates the peacetime performance of a single fleet of 20 Z-455s based at an MPC. The following requirements and restrictions are imposed on Z-455 activity:

- On average, ships must undertake two expeditions per day;
- Expeditions are four hours in duration;
- Turnaround considerations limit any one ship to three expeditions per day.

The MPC itself remains fully operational at all times, and replacement of condemned (irreparably damaged) components is always available from some supplier on Earth. For the moment, that supplier is not explicitly represented.

Input Data Set

Regardless of the situation being modeled, every Dyna-METRIC input data set contains the seven record groups detailing:

- administrative data (TOP)
- option selection (OPT)
- base descriptions (BASE)
- aircraft levels (ACFT)
- sortie rates (SRTS)
- maximum sortie rates (TURN)
- LRU descriptions (LRU).

In addition, this example uses the record groups associated with flying hours per sortie (FLHR) and stock levels (STK). Figure 2 shows these rather modest input requirements.

Administrative Data (TOP). Administrative data are described in the first three records of the data set. There is no header record. The first record contains the title of the exercise, the second record includes the mandatory Dyna-METRIC version identifier, the third record lists the times of analysis (the times for which output reports are produced).

Much of the information requested on the second record is inapplicable to this exercise. Transportation cutoff direction (one way or both ways) has no relevance because cutoff never occurs (transportation between the MPC and Earth is always available). Likewise, it is unimportant to this steady-state analysis whether transportation and repair delays have deterministic or exponential distributions. (The results in either case are identical since they depend only on the means of the distribution.) By default, the MPC experiences no administrative delay in component processing due to such activities as paperwork and accounting. Explicit mission labels are omitted, so default labels ("M1," "M2," etc.) will differentiate the mission types. Z-455s have only one mission (to pulverize asteroids and collect the pieces), so no confusion arises from accepting the default label.

The third record contains the times of analysis for which the model generates output reports. These reports give the status of Z-455 operations at the end of each time of analysis. Times are typically specified as days of a scenario (although any unit of time is acceptable as long as it is used consistently throughout the data set). Day 1 will provide the "peacetime" status because the steady-state nature of the situation ensures that results will be the same throughout all time.

Option Selection Records (OPT). The model options determine the nature of Dyna-METRIC's computations and output reports. For capability assessments, two options are particularly important: option 11 (performance report) and option 8 (problem LRUs list).

```
columns
         1         2         3         4         5         6         7         8
12345678901234567890123456789012345678901234567890123456789012345678901234567890
PEACETIME ANALYSIS OF Z-455
                   VERSION 4.4
   1
OPT
     11 15   .90
      8  5
BASE
MPC
ACFT
MPC   20
FLHR
MPC  4.0
SRTS
MPC  2.0
TURN
MPC  3.0
LRU
NAV COMPUTER          1    2  1          .00050  2.0       .60
NAV COMPUTER    X                               21.   21.   20000
MANEUVER THRUS        1   18 12          .00007  0.5       .30
MANEUVER THRUS  X                               14.   14.    8000
FUEL CELL             1   10 10          .00180  4.0       .95
FUEL CELL       X                               14.   14.    2500
DOCKING MECH          1    2  21         .00006  0.5       .10
DOCKING MECH    X                               14.   14.    5000
TARGET COMPUTER       1    1  1          .00040  2.0       .65
TARGET COMPUTER X                               28.   28.   30000
ASTEROID BLASTR       1    2  2          .00175  2.0       .75
ASTEROID BLASTR X                               28.   28.    9000
PART COLLECTOR        1    2  1          .00030  2.0       .20
PART COLLECTOR  X                               21.   21.    6000
SENSOR                1   36 36          .00020  2.0       .95
SENSOR          X                               35.   35.   15000
CONTROL PANEL         1    2  2          .00055  2.0       .70
CONTROL PANEL   X                               21.   21.   10000
CELL MOB PHONE        1    2  2          .00410  4.0       .99
CELL MOB PHONE  X                               21.   21.  750000        1
STK
NAV COMPUTER     MPC      5
MANEUVER THRUS   MPC     20
FUEL CELL        MPC     15
DOCKING MECH     MPC      5
TARGET COMPUTER  MPC      5
ASTEROID BLASTR  MPC      9
PART COLLECTOR   MPC      5
SENSOR           MPC     30
CONTROL PANEL    MPC      5
CELL MOB PHONE   MPC     20
```

Fig. 2—Input Data Set, Peacetime Analysis

The requested performance report will be based on a target of 15 percent of the fleet being NFMC. Also included is the number of Z-455s that are FMC with a .90 level of confidence. As many as five LRUs with the potential for degrading performance beyond an acceptable level will be noted on the problem LRUs list, along with their expected effects on the entire fleet.

Base Descriptions (BASE). These records describe the resupply characteristics, component repair availability, and cirf transportation linkages of operating bases. The paucity of data on this record reflects the simplicity of the example. There is no cirf in this scenario, so all cirf-related fields remain blank. Likewise, SRU-related fields remain blank because the LRUs of the Z-455 have no indentured SRUs. Default (blank) values for the remaining fields are appropriate, as discussed below.

Dyna-METRIC always considers resupply to be available during peacetime but unavailable from the start of the scenario until the resupply start time. Because this start time is set to day 0, there is no initial period of unavailable resupply. An additional resupply cutoff starts at day 0 and lasts for zero days, again allowing continuous resupply in wartime. Also, there is a switch that specifies whether parts ordered during peacetime may arrive during cutoff periods. The switch is unimportant here because there are no periods of resupply cutoff.

Dyna-METRIC also considers all types of repair to be available in peacetime. To keep repair immediately available at the start of the scenario, RR and RRR (maintenance type) repair starts on day 0 with peacetime reparables "arriving" at MPC the same day.

Sustained demand rates refer to an optional set of component demand rates introduced at an intermediate stage of a wartime scenario and carried through to its conclusion. They are inappropriate to this steady-state appraisal, so their start time remains blank. Note that unlike other start times, a blank or 0 here does not mean that sustained demand rates go into effect on day 0, but rather that they do not go into effect at all.

The identical base count gives the number of bases in the scenario identical to the one under consideration. Use of the identical base feature can prevent a duplication of effort on the part of both the user and the computer by allowing one base to act as a surrogate for any number of others from which it is indistinguishable. A blank, 0, or 1 indicates a unique base.

Finally, the onshore switch identifies the MPC as an offshore base. The onshore or offshore distinction carries with it no implied geographical connotation, it merely classifies bases into two types for which component demand rates differ (perhaps to account for variations in environment or climate).

Aircraft Levels (ACFT). These records give the number of aircraft assigned to each base during both peacetime and wartime. The MPC is assigned 20 ships in peacetime, which persists throughout the scenario.

Flying Hours per Sortie (FLHR), Sortie Rates (SRTS), and Maximum Sortie Rates (TURN). These three record groups describe the requirements and restrictions governing flying activity. The FLHR group gives the number of flying hours per sortie at each base. FLHR records are not mandatory, but they are necessary for bases whose aircraft fly sorties of a duration other than one hour. The SRTS record group gives the average number of sorties per aircraft per day that each base must generate. The TURN record group gives the *turn rate* (maximum number of sorties per day for a single aircraft) at each base. This example shows spaceships at the MPC averaging two 4-hour expeditions per day (during both peacetime and wartime). No ship may undertake more than three expeditions per day.

LRU Descriptions (LRU). Two records describe the characteristics of each LRU. Among the data on each pair of records are quantity per aircraft, demand rates, NRTS and condemnation rates at each echelon, repair time at each echelon, and resupply times. As before, fields pertaining to cirfs and depots remain blank.

The first record begins with the LRU name and level of repair indicator (a value of 1 allows repair at any echelon). The *quantity per aircraft* (QPA) and minimum QPA indicate the number of the LRU on the aircraft and the number that must be in place before the aircraft can fly. When the minimum QPA is less than the QPA, full mission capability is no longer contingent upon a full complement of components. Thus, a spaceship could still be FMC with 6 of its 18 maneuvering thrusters broken.

The NRTS/condemnation/failed SRU indicator remains blank, meaning that a failed part must first enter the repair process before a decision regarding its future disposition can be made. Since there are no higher echelons to which an LRU may be NRTSed and no SRUs subject to failure, condemnation is the only possible alternative to successful repair.

Typically LRU demand (or removal) rates (here for offshore bases) are given in terms of demands per component flying hour. If instead, an LRU experiences demands in proportion to sorties flown, its demands per sortie indicator is set to 1, as with the docking mechanism, which is used once per expedition regardless of expedition duration.

Base-level repair is described next by three parameters: repair time, NRTS rate, and condemnation rate. The MPC requires four days to repair a failed telephone (or to conclude that repair is impossible). The NRTS rate is inapplicable because the MPC is unsupported by higher echelons. The condemnation rate of 0.99 means that 99 percent of *all* (not just the nonNRTSed) telephones removed from Z-455s are condemned; replacements are procured individually from some unspecified outside source of supply. When higher echelons are unavailable, the model treats NRTS actions as if they were condemnations. Thus, similar results would be obtained with a NRTS rate of 0.99 and no condemnation rate.

The second record repeats the LRU name and has an artifice ('X') in column 17 to distinguish this record from the first of the LRU pair (to facilitate manipulation of the LRU record group). Use of the artifice has no bearing on the model's performance.

The peacetime and wartime resupply times are defined as the procurement lead time for an LRU at its highest echelon of repair. The telephone's resupply times are set to 21 days each for this steady-state appraisal. Also included is the LRU's unit cost ($750,000 for a telephone) even though it is superfluous for capability assessments. Finally, the no-cannibalization indicator determines whether the model may consolidate an LRU's backorders on the fewest possible aircraft. Except for the telephone, each LRU of the Z-455 is assumed to be freely interchangeable among different ships. The telephone cannot be cannibalized, so its backorders will be randomly distributed across the fleet for purposes of computing NFMC ships.

Stock Levels (STK). These records specify both prepositioned stock levels and intra-scenario changes in stock levels at various locations. This optional record group is for those components with positive stock levels at some point in the scenario; if the stock level of each component remains zero throughout, no records are required. The absence of an explicit start time for these levels implies that these levels are always in effect.

Output from Programs PART, PIPE, and MOD

Under normal circumstances, programs PART, PIPE, and MOD generate no output reports. Instead they produce short messages listing their own names and the names of the data sets upon which they operated (ECHO and REPORT precede their reports with such messages as well). Unless internal error messages are triggered (or, in the case of PIPE, special options are selected), this is the full extent of the output from these programs.

Output from Program ECHO

Dyna-METRIC precedes its more substantive output reports with an exhaustive summary, or echo, of the input data set. Selecting option 13 (full echo suppression) or option 14 (component data echo suppression) turns off this feature.

Administrative Data, Option Selection, and Flying Program Echo. The information included in the TOP record group is summarized here (Fig. 3), along with the user-selected options from the OPT record group and the flying program information in the ACFT, FLHR, SRTS, and TURN record groups.

It is only here that the existence of duplicate bases is explicitly acknowledged. Elsewhere, any mention of the prototype carries with it the implication that its duplicates have the same behavior and attributes. In aggregate or *worldwide* reports, an accounting is made of the duplicate bases by a simple multiplication of the prototype's associated values.

Resupply and Repair Characteristics Echo. This echo (Fig. 4) displays the resupply and repair data from the description records for each location in the scenario. *Supplier* is the external source of supply, which is not explicitly represented. Notice that the sustained

```
SCENARIO INCLUDES ONE BASE,   WITH BASE ADMINISTRATIVE DELAY SET TO  0.000;
                  NO CIRFS;
                  NO DEPOTS.

ANALYSIS REQUESTED FOR TIMES     1

NUMBER OF BASES OF EACH TYPE --
 MPC
   1

FIXED REPAIR AND TRANSPORTATION TIMES.
TRANSPORTATION CUTOFFS APPLY ONLY TO FORWARD TRANSPORTATION.

 OPTION
 SELECTED  MEANING
    8      LIST PROBLEM LRUS -- UP TO  5 LRUS.
   11      CALCULATE PERFORMANCE AT  15 PERCENT NFMC BASED ON INPUT OR
           PREVIOUS STOCK.

NO BASES, CIRFS, OR DEPOTS CANNIBALIZE SRUS AND SUBSRUS.

FLYING PROGRAM FOR BASE -- MPC (OFFSHORE)
   DAYS  AIRCRAFT   SORTIES   FH/SRT   MAX. SORTIES    M1   M2   M3   M4   M5
PEACE        20        2.00     4.00        3.00
   1         20        2.00     4.00        3.00        X    X    X    X    X
```

Fig. 3—Administrative Data, Option Selection, and Flying Program Echo, Peacetime Analysis

demand start time (left blank on the BASE record) is adjusted to the last time of analysis plus one (day 2) so that sustained demand rates never go into effect at the MPC.

Component Data Echo. Component-related data from the LRU and STK record groups are echoed in six separate tables.

The first two tables (Fig. 5) contain references to cirf-served bases, cirfs, and depots. The MPC is considered a *lone base* (unsupported by a cirf) that operates without assistance from a depot, so references to those other locations may be ignored. Furthermore, quantity per application and replacement percentage are defined only in terms of indentured subcomponents. In Dyna-METRIC, an LRU's quantity per application is called its quantity per aircraft and is treated in a subsequent table.

The third table (also Fig. 5) draws largely from input record groups not included in the example. All LRUs are assumed by default to be of the RR variety. The terms *test equipment* and *test stand* refer to the model's constrained repair module, which is not being used at present. A variance-to-mean ratio of 1.0 implies that the size of a pipeline has a Poisson probability distribution. Wartime demand rate multipliers are used to express component demand rates during wartime as a multiple of their peacetime rates. Since these multipliers default to a value of 1.0, demand rates remain unchanged after the peacetime to wartime transition. Sustained demand rates default uniformly to zero; these are unimportant because their start time was set such that they never go into effect (during the period of interest). Finally, each LRU is essential to the execution of all wartime missions.

```
PEACETIME FORWARD PIPELINE --
'N' MEANS DOES NOT START TO EMPTY UNTIL CONNECTION HAS BEEN
    RESTABLISHED BETWEEN LOCATION AND HIGHER ECHELON
'Y' MEANS CONTINUES TO EMPTY.

NAME SUPPLIER
MPC       N

TIMES WHEN FORWARD TRANSPORTATION IS FIRST ESTABLISHED --
NAME SUPPLIER
MPC       0.0

TIMES WHEN FORWARD TRANSPORTATION IS FIRST CUTOFF --
NAME SUPPLIER
MPC       0.0

DURATION OF FORWARD TRANSPORTATION CUTOFF --
NAME SUPPLIER
MPC       0.0

        SUSTAINED   TIME    SRU AND SUBSRU  0-RR     1-RRR
         DEMAND  PEACETIME      REPAIR      REPAIR  REPAIR
LOCATION  START  REPARABLES     START       START   START
  NAME    TIME     ARRIVE       TIME        TIME    TIME
  MPC      2.0       1.0          1.0         1.0     1.0
```

Fig. 4—Resupply and Repair Characteristics Echo, Peacetime Analysis

TABLE 1

OVERALL INFORMATION AND INDENTURE RELATIONSHIPS
(AN ASTERISK AFTER THE DEMAND RATES INDICATES ALL DEMAND RATES FOR THAT COMPONENT ARE PER SORTIE RATHER THAN PER FLYING HOUR.)

PART NAME	NUMBER		CAN TEST AT CIRF?	NRTS OR CONDEMN	COST	---DEMANDS PER-- FLYING HOUR ONSHORE	OFFSHORE	LEVEL OF REPAIR	RESUPPLY (DAYS) PEACE	WAR	QUANTITY PER APPLICATION	REPLACEMENT PERCENTAGE DEPOT	CIRF	BASE
NAV COMPUTER	L	1	NO	AFTER TEST	20000.	0.00000	0.00050	BASE	21.0	21.0	----	-----	-----	-----
MANEUVER THRUS	L	2	NO	AFTER TEST	8000.	0.00000	0.00007	BASE	14.0	14.0	----	-----	-----	-----
FUEL CELL	L	3	NO	AFTER TEST	2500.	0.00000	0.00180	BASE	14.0	14.0	----	-----	-----	-----
DOCKING MECH	L	4	NO	AFTER TEST	5000.	0.00000	0.00006*	BASE	14.0	14.0	----	-----	-----	-----
TARGET COMPUTER	L	5	NO	AFTER TEST	30000.	0.00000	0.00040	BASE	28.0	28.0	----	-----	-----	-----
ASTEROID BLASTR	L	6	NO	AFTER TEST	9000.	0.00000	0.00175	BASE	28.0	28.0	----	-----	-----	-----
PART COLLECTOR	L	7	NO	AFTER TEST	6000.	0.00000	0.00030	BASE	21.0	21.0	----	-----	-----	-----
SENSOR	L	8	NO	AFTER TEST	15000.	0.00000	0.00020	BASE	35.0	35.0	----	-----	-----	-----
CONTROL PANEL	L	9	NO	AFTER TEST	10000.	0.00000	0.00055	BASE	21.0	21.0	----	-----	-----	-----
CELL MOB PHONE	L	10	NO	AFTER TEST	750000.	0.00000	0.00410	BASE	21.0	21.0	----	-----	-----	-----

TABLE 2

BASE, CIRF, AND DEPOT RELATED LRU, SRU AND SUBSRU DATA

PART NAME	NUMBER		------LONE BASE------ REPAIR TIME (DAYS)	NRTS RATE	CONDEMN RATE	--CIRF-SERVED BASE--- REPAIR TIME (DAYS)	NRTS RATE	CONDEMN RATE	---------CIRF-------- REPAIR TIME (DAYS)	NRTS RATE	CONDEMN RATE	ASSIGNED DEPOT	--------DEPOT--------- REPAIR TIME (DAYS)	REPAIR LIMIT PER DAY	CONDEMN RATE
NAV COMPUTER	L	1	2.00	0.000	0.600	0.00	0.000	0.000	0.00	0.000	0.000	NONE	0.00	9999.00	0.000
MANEUVER THRUS	L	2	0.50	0.000	0.300	0.00	0.000	0.000	0.00	0.000	0.000	NONE	0.00	9999.00	0.000
FUEL CELL	L	3	4.00	0.000	0.950	0.00	0.000	0.000	0.00	0.000	0.000	NONE	0.00	9999.00	0.000
DOCKING MECH	L	4	0.50	0.000	0.100	0.00	0.000	0.000	0.00	0.000	0.000	NONE	0.00	9999.00	0.000
TARGET COMPUTER	L	5	2.00	0.000	0.650	0.00	0.000	0.000	0.00	0.000	0.000	NONE	0.00	9999.00	0.000
ASTEROID BLASTR	L	6	2.00	0.000	0.750	0.00	0.000	0.000	0.00	0.000	0.000	NONE	0.00	9999.00	0.000
PART COLLECTOR	L	7	2.00	0.000	0.200	0.00	0.000	0.000	0.00	0.000	0.000	NONE	0.00	9999.00	0.000
SENSOR	L	8	2.00	0.000	0.950	0.00	0.000	0.000	0.00	0.000	0.000	NONE	0.00	9999.00	0.000
CONTROL PANEL	L	9	2.00	0.000	0.700	0.00	0.000	0.000	0.00	0.000	0.000	NONE	0.00	9999.00	0.000
CELL MOB PHONE	L	10	4.00	0.000	0.990	0.00	0.000	0.000	0.00	0.000	0.000	NONE	0.00	9999.00	0.000

TABLE 3

SPECIFIC LRU INFORMATION

LRU NAME	NUMBER		MAINT. TYPE	ASSIGNED TEST EQUIPMENT	PROB PMC TEST STAND WILL TEST	PIPELINE VARIANCE TO MEAN RATIO	CANNIBALIZE	WARTIME DEMAND RATE MULTIPLIER ONSHORE	OFFSHORE	SUSTAINED DEMAND RATE ONSHORE	OFFSHORE	ESSENTIAL IN MISSIONS M1	M2	M3	M4	M5
NAV COMPUTER	L	1	RR	NONE	1.0000	1.000	YES	1.000	1.000	0.000	0.000	X	X	X	X	X
MANEUVER THRUS	L	2	RR	NONE	1.0000	1.000	YES	1.000	1.000	0.000	0.000	X	X	X	X	X
FUEL CELL	L	3	RR	NONE	1.0000	1.000	YES	1.000	1.000	0.000	0.000	X	X	X	X	X
DOCKING MECH	L	4	RR	NONE	1.0000	1.000	YES	1.000	1.000	0.000	0.000	X	X	X	X	X
TARGET COMPUTER	L	5	RR	NONE	1.0000	1.000	YES	1.000	1.000	0.000	0.000	X	X	X	X	X
ASTEROID BLASTR	L	6	RR	NONE	1.0000	1.000	YES	1.000	1.000	0.000	0.000	X	X	X	X	X
PART COLLECTOR	L	7	RR	NONE	1.0000	1.000	YES	1.000	1.000	0.000	0.000	X	X	X	X	X
SENSOR	L	8	RR	NONE	1.0000	1.000	YES	1.000	1.000	0.000	0.000	X	X	X	X	X
CONTROL PANEL	L	9	RR	NONE	1.0000	1.000	YES	1.000	1.000	0.000	0.000	X	X	X	X	X
CELL MOB PHONE	L	10	RR	NONE	1.0000	1.000	NO	1.000	1.000	0.000	0.000	X	X	X	X	X

Fig. 5—Tables 1-3 of Component Data Echo, Peacetime Analysis

The fourth table (Fig. 6) lists the QPA and minimum QPA of each component. The fifth table, based upon a record group not yet encountered, lists the fraction of aircraft at each base on which a given LRU is installed. The default value of an LRU's application fraction is 1.0; thus, all ships at the MPC contain a complement of each of these 10 LRUs. The sixth table reports initial levels of component stock.

Component Count Echo. After the contents of the input data set are echoed, Dyna-METRIC provides a count of different classes of components (Fig. 7). Instead of processing many components all at once, the PART program partitions a large data set into smaller groups of components, accounting for families of indentured subcomponents.[1] This peacetime analysis is sufficiently small in scale that there was no need to partition the input data set.

TABLE 4

QUANTITIES PER AIRCRAFT

PART NAME	NUMBER	BASE	QP ACFT	QP APPL	MIN QPA
NAV COMPUTER	L 1	MPC	2	---	1
MANEUVER THRUS	L 2	MPC	18	---	12
FUEL CELL	L 3	MPC	10	---	10
DOCKING MECH	L 4	MPC	2	---	2
TARGET COMPUTER	L 5	MPC	1	---	1
ASTEROID BLASTR	L 6	MPC	2	---	2
PART COLLECTOR	L 7	MPC	2	---	1
SENSOR	L 8	MPC	36	---	36
CONTROL PANEL	L 9	MPC	2	---	2
CELL MOB PHONE	L 10	MPC	2	---	2

TABLE 5

APPLICATION FRACTIONS

LRU NAME	NUMBER		
NAV COMPUTER	L 1	MPC	1.00
MANEUVER THRUS	L 2	MPC	1.00
FUEL CELL	L 3	MPC	1.00
DOCKING MECH	L 4	MPC	1.00
TARGET COMPUTER	L 5	MPC	1.00
ASTEROID BLASTR	L 6	MPC	1.00
PART COLLECTOR	L 7	MPC	1.00
SENSOR	L 8	MPC	1.00
CONTROL PANEL	L 9	MPC	1.00
CELL MOB PHONE	L 10	MPC	1.00

TABLE 6

INITIAL STOCK LEVELS

PART NAME	NUMBER		
NAV COMPUTER	L 1	MPC	5
MANEUVER THRUS	L 2	MPC	20
FUEL CELL	L 3	MPC	15
DOCKING MECH	L 4	MPC	5
TARGET COMPUTER	L 5	MPC	5
ASTEROID BLASTR	L 6	MPC	9
PART COLLECTOR	L 7	MPC	5
SENSOR	L 8	MPC	30
CONTROL PANEL	L 9	MPC	5
CELL MOB PHONE	L 10	MPC	20

Fig. 6—Tables 4-6 of Component Data Echo, Peacetime Analysis

[1]The size of these groups is limited by several compilation parameters, described in App. B.

```
TOTAL LRU COUNT =  10
TOTAL SRU COUNT =   0
TOTAL SUBSRU COUNT =   0

MAXIMUM LRUS IN A GROUP =  10
MAXIMUM SRUS IN A GROUP =   0
MAXIMUM SUBSRUS IN A GROUP =   0

MAXIMUM SRUS TO A LRU =   0
MAXIMUM LRUS TO A SRU =   0
MAXIMUM SUBSRUS TO A SRU =   0
MAXIMUM SRUS TO A SUBSRU =   0
```

Fig. 7—Component Count, Peacetime Analysis

Output from Program REPORT

Performance measures such as those requested with option 11 are computed with respect to the status of the system at the end of a given time of analysis. REPORT also gives the output from option 8, the list of LRUs with unacceptably high probabilities of incapacitating at least some specified percentage of the entire fleet. Figure 8 shows the day 1 output for both options.

Performance Report. The option 11 report has two sections: one assuming *full cannibalization* and one assuming *partial cannibalization*. The partial cannibalization portion is based upon the actual cannibalization indicators specified in the input data set. The full cannibalization portion ignores all such specifications and describes the performance that can be achieved if every component may be cannibalized. Thus, it provides a good indication of the benefits that might accrue from pursuing a policy of full cannibalization.

The target number of NFMC aircraft at the base is determined by the first parameter for option 11 (set to 15 percent of the MPC's fleet of 20 ships, or 3 ships). This target NFMC figure serves as the threshold for determining whether an LRU is entered on the problem LRUs list. All other quantities being equal, the number of problem LRUs should increase or decrease as the target NFMC figure decreases or increases.

The first two performance measures in each section of the report are probabilities. The first gives the probability that less than 15 percent of the MPC's ships are NFMC at the end of the day—a very small number regardless of cannibalization assumption. The second is the probability of achieving the current day's sortie generation goal given the status at the end of the day. Like all of the statistics in this report, PROB.ACHIEVE SORTIES is better under full cannibalization than under partial cannibalization because of the influence of the telephone, which was considered unsuitable for cannibalization.

Next the report lists the number of FMC aircraft observed at the end of the day, with the option 11 confidence level. Under conditions of full cannibalization, the MPC is expected to have at least 10 FMC Z-455s (with .90 probability) at the end of day 1; this drops to 7 FMC ships when we account for the MPC's inability to cannibalize telephones.

The expected number of NFMC aircraft at a base is specified both in absolute terms (as E(NFMC)) and as a proportion of the base's total allocation of aircraft (as EXP. % NFMC).

The expected number of sorties and expected sorties per (FMC) aircraft are predictive values based upon the expected status of the system at the end of the current day. Given the condition of the MPC's fleet at the end of day 1, these two quantities estimate both the number of sorties and the turn rate that could be achieved under the input restrictions imposed on day 1 activity. These quantities are somewhat related: As E(SORTIES) falls ever shorter of the requested number of sorties, EXP. SORTIES/ACFT approaches the limiting turn rate given in the TURN record group; alternatively, as E(SORTIES) approaches the requested number of sorties, EXP. SORTIES/ACFT declines, reflecting a slackened demand for the services of FMC aircraft.

The last entry in the report gives the expected total number of backorders among all components. This is a traditional measure of system performance and is not as useful as the more integrated measures already discussed.

Problem LRUs List. The problem LRUs list appears in three separate pieces:

- a tabular listing of the LRUs and their fleetwide effects (Fig. 8);
- a graphic display of this same information (Fig. 9);
- a detailed report of each LRU's pipeline contents at the end of the current day (Fig. 10).

Component impact is an approximation of the expected number of aircraft rendered NFMC by shortages of a particular LRU. It is computed by dividing the LRU's expected number of backorders by its QPA (thus implicitly assuming its ability to be cannibalized). Target component impact, again, comes from option 11.

Any LRU incapacitating more aircraft than its target component impact with probability greater than (1 − the confidence level given as the second option 11 parameter) is included on the problem LRUs list. Therefore, although ASTEROID BLASTER has a component impact of two, it appears on the list because the probability that it actually incapacitates more than three ships is greater than .10 (1 − .90).

The third piece of the problem LRUs list is the detailed summary of each problem LRU's pipeline contents (Fig. 10). There are no incoming *retrograde* pipelines associated with the MPC because bases are the lowest echelon in the model. The *administrative delay* pipeline is zero to reflect instantaneous administrative processing. The expected number of a component *in test/repair* at the end of the day includes units that eventually will be condemned but must first complete the repair process (as determined by the NRTS/condemnation/failed SRU switch on the LRU record). The *awaiting parts* (AWP) pipeline, relevant only in the context of indentured subcomponents, remains zero. Parts *on order* come from the unnamed external source of supply. Summing these five quantities yields the total number of the item in pipelines at the end of day 1.

The expected number of backorders is the expected positive difference between total pipeline contents and stock level. Dividing this expected number of backorders by the LRU's QPA gives the item's component impact (e.g., 6.270 for the telephone assuming cannibalization).

PERFORMANCE BASED ON STOCK ON HAND ON DAY 1.

			FULL CANNIBALIZATION								PARTIAL CANNIBALIZATION							
BASE	TARG. NFMC	TOTAL ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC-90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC-90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	TOTAL BACK ORDERS
MPC	3	20	0.031	0.503	10	6.853	0.343	36.77	2.392	0.000	0.044	7	10.278	0.514	29.03	2.905	62.33

POTENTIAL PROBLEM PARTS REPORT

WORLD-WIDE AIRCRAFT STATUS - DAY 1
(* - INDICATES AN APPLICATION FRACTION NOT EQUAL TO ONE.)

				WORLD-WIDE					
LRU NAME	NUMBER	WUC	TEQ	COMP. IMPACT	MINIMUM COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG COMP. IMPACT
CELL MOB PHONE	L 10			6.27	6.27	3.00	MPC	6.27	3.00
FUEL CELL	L 3			3.48	3.48	3.00	MPC	3.48	3.00
ASTEROID BLASTR	L 6			2.04	2.04	3.00	MPC	2.04	3.00

Fig. 8—Performance Report and Problem LRUs List, Peacetime Analysis

```
WORLD-WIDE COMPONENT SUPPORT IMPACT - DAY     1

O = WORLD-WIDE COMPONENT IMPACT WITH CURRENT STOCK DISTRIBUTION
* = WORLD-WIDE COMPONENT IMPACT WITH STOCK REDISTRIBUTED

        10.-

         9.-

         8.-

         7.-
              ***
         6.-  ***
              ***
              ***
         5.-  ***
              ***
              ***
         4.-  ***
              ***
              ***  ***
         3.-  ***  ***
              ***  ***
              ***  ***
         2.-  ***  ***  ***
              ***  ***  ***
              ***  ***  ***
         1.-  ***  ***  ***
              ***  ***  ***
              ***  ***  ***
        -0.-  ***  ***  ***
               A    B    C
    LRU NAME             NUMBER    WUC          LRU NAME            NUMBER    WUC
A= CELL MOB PHONE        L  10              B= FUEL CELL            L   3
C= ASTEROID BLASTR       L   6
```

Fig. 9—Bar Chart of Problem LRUs, Peacetime Analysis

WARTIME DYNAMIC EXAMPLE

The second example incorporates three important features that appear in many Dyna-METRIC exercises: multiple echelons, indentured subcomponents, and scenario dynamics.

Suppose the Really Awesome New Doodads Corporation (which owns and operates the Z-455s) prepares to exploit the popularity of the "Pet Asteroid" by quadrupling the size of its fleet. Besides more spaceships and bases, a more viable logistics support structure is needed. In particular, there are now three MPCs with a depot on Earth's moon supporting them.

The original MPC (renamed MPC1) now has no repair capability, so its parts will use the repair facilities at MPC2. To model this, MPC2 is collocated with a cirf (called CRF1) that will serve the 60 spaceships at both bases. Rather than condemning components it cannot repair, MPC1 sends them to the cirf while simultaneously requisitioning replacements from the cirf. MPC3 accommodates 20 Z-455s and operates in a remote and highly radioactive section of the asteroid belt in order to collect glow-in-the-dark Pet Asteroids. Because of its great distance from CRF1, MPC3 is supported directly by the depot. The support relationships among DEPO, CRF1, and the three MPCs are depicted in Fig. 11.

Not content merely to enhance the quality of Z-455 logistics support, management also incorporated several technological improvements into shipboard components. These were

```
DETAILED PIPELINE SEGMENT REPORT FOR PROBLEM LRUS, SRUS, AND SUBSRUS

                                          ---------------------------WORLD-WIDE---------------------------
                     NUMBER    WUC     COMPONENT IMPACT  MINIMUM COMPONENT IMPACT  TARGET COMPONENT IMPACT
LRU CELL MOB PHONE    L  10                 6.270                  6.270                    3.000

LOC.  RETROGRADE  ADMINISTRATIVE  IN TEST/REPAIR  AWAITING PARTS  ON ORDER (NOT RECEIVED)    TOTAL   STOCK BACKORDERS    BO/QPA
MPC   ----------       0.000           5.248           0.000               27.276             32.524   20      12.540       6.270
                                          ---------------------------WORLD-WIDE---------------------------
                     NUMBER    WUC     COMPONENT IMPACT  MINIMUM COMPONENT IMPACT  TARGET COMPONENT IMPACT
LRU FUEL CELL         L   3                 3.482                  3.482                    3.000

LOC.  RETROGRADE  ADMINISTRATIVE  IN TEST/REPAIR  AWAITING PARTS  ON ORDER (NOT RECEIVED)    TOTAL   STOCK BACKORDERS    BO/QPA
MPC   ----------       0.000          11.520           0.000               38.304             49.824   15      34.824       3.482
                                          ---------------------------WORLD-WIDE---------------------------
                     NUMBER    WUC     COMPONENT IMPACT  MINIMUM COMPONENT IMPACT  TARGET COMPONENT IMPACT
LRU ASTEROID BLASTR   L   6                 2.044                  2.044                    3.000

LOC.  RETROGRADE  ADMINISTRATIVE  IN TEST/REPAIR  AWAITING PARTS  ON ORDER (NOT RECEIVED)    TOTAL   STOCK BACKORDERS    BO/QPA
MPC   ----------       0.000           1.120           0.000               11.760             12.880    9       4.088       2.044
```

Fig. 10—Problem LRU Pipeline Contents, Peacetime Analysis

```
            DEPO
            / \
           /   \
          /     \
       CRF1      \
        /\        \
       /  \        \
      /    \        \
   MPC1    MPC2     MPC3
```

Fig. 11—Z-455 Logistics Support Structure

confined to weaponry and communications and took the form of indentured subcomponents. The asteroid blaster was augmented by a single SRU—a revolutionary pulsed beam generator designed to produce asteroidal debris of especially pleasing shape and texture. The telephone benefitted from the addition of three SRUs, two of which contained subSRUs:

1. automatic dialing module
   — memory unit
   — redialing circuit
2. automatic answering system
   — message recorder
   — call screening device
3. mute switch

Despite the scheduled increase in fleet size, management considers it likely that at current workload levels, production will be insufficient to meet the demand for Pet Asteroids during the approaching holiday shopping season. Thus activity at MPC1 will surge to a higher level than its peacetime norm. Furthermore, as MPC2 and MPC3 commence operations, they too will move quickly toward a high level of activity. At the end of the surge, there will be diminished activity at all MPCs. Management wishes to predict how well this system will perform under such unprecedented and rigorous conditions.

Input Data Set

The extension of the example requires modification of the input data set. Previously constructed record groups need to be updated, and new record groups need to be added. Figure 12 shows the scenario record groups in the new data set.

Administrative Data (TOP). The first three records require few changes. The switch for transportation and repair time probability density function now indicates that transportation and repair times are to be considered exponentially distributed. Also, the cirf and depot experience administrative delays of 0.5 days each; this means components NRTSed to higher echelons from the MPCs must first spend a (constant) half day in an administrative pipeline before they may enter into the repair process.

Option Selection Records (OPT). Options 11 and 8 remain in effect, and option 23 (subcomponent supplement to the problem LRUs list) is added. It will report the two most troublesome SRUs indentured to each LRU on the problem LRUs list, and also the two most

```
columns
         1         2         3         4         5         6         7         8
12345678901234567890123456789012345678901234567890123456789012345678901234567890

WARTIME DYNAMIC ANALYSIS
 1       0.5  0.5  VERSION 4.4
   1  15  30  45  60
OPT
     11 15  .90
      8  5
     23  2  .15
DEPT
DEPO
CIRF
CRF1                                                5    5    5
BASE
MPC1CRF1  2.0  2.0
MPC2CRF1  0.0  0.0
MPC3                                                8    8    8                1
TRNS
CRF1 DEPO   2.0   2.0
MPC3 DEPO   3.5   3.5
ACFT
MPC1  20
MPC2  20  10  40
MPC3  20
FLHR
MPC1 4.0   1 5.0
MPC2 5.0
MPC3 5.0
SRTS
MPC1 2.0   1 3.0  30 2.0
MPC2 0.0   5 3.0  10 2.7   25 1.5
MPC3 0.0   8 2.0  50 1.0
TURN
MPC1 3.0
MPC2 3.0
MPC3 3.0
```

Fig. 12—Scenario Record Groups, Wartime Analysis

troublesome subSRUs indentured to each of these SRUs, as long as the parent component's AWP pipeline is 0.15 of its total pipeline.

Depot Descriptions (DEPT). These records describe the resupply and repair characteristics of each depot. This depot (DEPO) resembles the unsupported MPC in the first example. Resupply from sources external to the Really Awesome New Doodads Corporation is always available and never cut off. All types of repair are available at DEPO throughout both peacetime and wartime. The depot is unable to cannibalize SRUs among LRUs (and subSRUs among SRUs), so the SRU cannibalization switch remains blank.

In instances in which a depot (or, more commonly, a cirf masquerading as a depot for modeling purposes) is deployed in wartime to a location different from its peacetime location, the reparable arrival time specifies the day upon which peacetime reparable components temporarily abandoned during the deployment arrive at the new site and become available for repair and subsequent use. DEPO remains firmly on the moon at all times, so reparables arrive on day 1 by default.

Cirf Descriptions (CIRF). These records have the same format as the depot description records. Since CRF1 is supported by DEPO with respect to each Z-455 component, it has no direct access to outside suppliers. Resupply data are therefore immaterial and may be omitted. Repair facilities at CRF1 are available beginning on day 5 of wartime (with peacetime repair again implicit) for RR and RRR LRUs as well as for subcomponents. Like DEPO, CRF1 is unable to cannibalize subcomponents.

Base Descriptions (BASE). The BASE record group now reflects not only the addition of two new MPCs, but also the existence of a cirf, a depot, and indentured subcomponents. As in the case of CRF1, data regarding external sources of resupply are irrelevant because each MPC is fully supported by a higher echelon.

Records for cirf-served bases (MPC1 and MPC2) describe the resupply linkage between base and cirf. One-way transportation time between MPC1 and the cirf is two days, between MPC2 and the cirf zero days (the latter two are collocated). Resupply from the cirf is always available and never cut off. Although MPC2 does not exist (has no flying activity) before day 5, the model nonetheless considers cirf-to-base resupply to be available during peacetime (a harmless assumption). The cirf availability switch is irrelevant for both bases: for MPC1 because cirf resupply starts on day 0, and for MPC2 because it does not exist in peacetime. Repair of RR and RRR LRUs, and now also for subcomponents, technically starts on day 1 (with peacetime repair inherent), even though neither base has repair (only the cirf does). Base-level repair for MPC1 and MPC2 can be turned off later in the LRU record group. Finally, the MPCs cannot cannibalize subcomponents.

Aside from being a lone base, MPC3 is distinguished from its companions in two ways: It has repair facilities and is designated an onshore base. Repair begins on day 8, its first day of operations. Before this time, MPC3 does not exist in any real sense, even though Dyna-METRIC assumes that it has repair facilities available during peacetime. The absence of ships and activity assures that this "phantom" MPC cannot generate a repair workload before its intended time. MPC3 is classified as an onshore base to point to its deployment in a highly radioactive environment; its components may be assigned a different set of demand rates (in the LRU record group) rather than share a common set with the other MPCs.

Depot Transportation (TRNS). These records describe the transportation links between each depot and any lower echelons to which it provides direct logistics support. Neither MPC1 nor MPC2 appears here because they are fully supported by CRF1 rather than DEPO. Dyna-METRIC interprets their absence as an indication that each is connected to DEPO by an instantaneous and ever-present mode of transportation. Although inaccurate, this is acceptable because the specified support structure prevents resources from being shipped along these routes.

The data contained on a TRNS record are similar to those on the BASE record describing a base's supporting cirf. Here the transportation time between the cirf and depot is two days one way, between MPC3 and the depot 3.5 days one way. All assumptions of transportation availability are the same as those for resupply—always available and never cut off. Once again, the model harmlessly assumes that MPC3 exists—and has transportation links to DEPO—during peacetime.

Aircraft Levels (ACFT). Each MPC is assigned 20 spaceships in peacetime, even though there is no flying activity at MPC2 and MPC3. That level holds throughout the wartime scenario unless it is changed, as on day 10 when the fleet at MPC2 increases to 40 ships.

Flying Hours per Sortie (FLHR), Sortie Rates (SRTS), Maximum Sortie Rates (TURN). Each MPC experiences varying levels of activity. For example, MPC3 begins operations on day 8 of the wartime scenario, when the 20 Z-455s stationed there fly an average of two expeditions per ship per day. Each expedition is scheduled to last 5 hours. Activity drops on day 50, so the Z-455s average only one daily expedition per ship. Throughout the scenario, turnaround considerations limit the maximum number of daily expeditions to 3 per ship.

LRU Descriptions (LRU). Although the number of LRUs remains constant, each record must be modified to account for the presence of higher echelons and an onshore base. The revised LRU record group appears in Fig. 13.

Three fields control the mix of repair capability among echelons. Depot reparability is confirmed by naming DEPO as the responsible depot. Repair is allowed at all echelons by setting the level of repair indicator to 1. Cirf reparability is confirmed by setting the cirf reparability switch to 1 as well. Alternatively, settings of "DEPO," "1," and "0" would imply that the LRU could be repaired at the depot and bases but not at the cirf.

QPA and minimum QPA remain the same. If, however, an LRU's quantity per aircraft differed by base, the base-specific QPAs would be entered in the QPA record group (described later). Entries in the QPA record group override quantities in the LRU record group.

```
columns
         1         2         3         4         5         6         7
123456789012345678901234567890123456789012345678901234567890123456789012345

LRU
NAV COMPUTER       DEPO 1 1  2  1    .00062 .00050  2.0  .60          0.0 1.00
NAV COMPUTER     X 1.5  .35         1.8     0.02   21.  42.      20000
MANEUVER THRUS     DEPO 1 1 18 12    .00007 .00007  0.5  .30          0.0 1.00
MANEUVER THRUS   X 0.5  .05         0.9            14.  28.       8000
FUEL CELL          DEPO 1 1 10 10    .00200 .00180  4.0  .95          0.0 1.00
FUEL CELL        X 2.5  .45         4.0     0.02   14.  28.       2500
DOCKING MECH       DEPO 1 1  2 21    .00006 .00006  0.5  .10          0.0 1.00
DOCKING MECH     X 0.5  .00         0.5            14.  28.       5000
TARGET COMPUTER    DEPO 1 1  1  1    .00046 .00040  2.0  .65          0.0 1.00
TARGET COMPUTER  X 1.5  .30         1.6     0.02   28.  56.      30000
ASTEROID BLASTR    DEPO 1 1  2  2    .00195 .00175  2.0  .75          0.0 1.00
ASTEROID BLASTR  X 1.0  .35         2.0            28.  56.       9000
PART COLLECTOR     DEPO 1 1  2  1    .00038 .00030  2.0  .20          0.0 1.00
PART COLLECTOR   X 1.8  .05         2.1            21.  42.       6000
SENSOR             DEPO 1 1 36 36    .00200 .00020  2.0  .95          0.0 1.00
SENSOR           X 1.6  .50         2.0     0.07   35.  56.      15000
CONTROL PANEL      DEPO 1 1  2  2    .00061 .00055  2.0  .70          0.0 1.00
CONTROL PANEL    X 1.8  .25         2.2            21.  42.      10000
CELL MOB PHONE     DEPO 1 1  2  2    .00410 .00410  4.0  .99          0.0 1.00
CELL MOB PHONE   X 3.5  .60         4.5      .15   21.  42.     750000          1
```

Fig. 13—LRU Record Group, Wartime Analysis

As before, the NRTS/condemnation/failed SRU indicator is left blank, thereby requiring all LRUs to pass through the repair process before being NRTSed or condemned, and before any failed SRUs may be found. For this purpose, lack of repair[2] at an echelon is not considered equivalent to a NRTS rate of 1.0; an LRU that cannot be repaired at the base level is sent immediately to the next higher echelon upon its removal at an MPC, whereas an LRU that can be repaired at base level—but has a NRTS rate of 1.0—must still pass through the MPC repair cycle.

The inclusion of onshore demand rates reflects the addition of MPC3 to the scenario. Allowing different demand rates is especially important in an evaluation of the base-by-base performance of such components as the sensor, which has an onshore demand rate ten times as high as its offshore demand rate.

The updated records specify the LRU's repair time,[3] NRTS rate, and condemnation rate at four types of locations: lone bases, cirf-served bases, cirfs, and depots. The lone base numbers, applicable to parts at MPC3, remain as before in the first example. The cirf-served base numbers trigger parts at MPC1 and MPC2 to skip base repair (repair time of zero) and automatically be NRTSed to the cirf (NRTS rate of 1.00). The cirf numbers (on the second card) are similar to the lone base numbers. In all three cases, parts that cannot be repaired locally are never condemned (condemnation rates are left blank), but are NRTSed to the next higher echelon. Depot numbers include repair time and condemnation rate but not NRTS rate (because there is no higher echelon to which an irreparable part may be NRTSed).

The daily repair limit at the depot queues units awaiting repair if the number of the LRU's arrivals at the depot should exceed it. Because DEPO is so overstaffed and overequipped that any sort of capacity shortage is inconceivable, the depot repair limit remains blank (here, blank or 0 denotes unlimited capacity while "−1" denotes zero capacity).

Finally, the time it takes the highest echelon repairing the LRU (DEPO) to procure a replacement in wartime is uniformly twice that of peacetime.

SRU Descriptions (SRU). These records contain essentially the same information about SRUs that the LRU description records contain about LRUs. However, they diverge in one important respect. Recall that each LRU's ability to be cannibalized is declared individually. SRUs, however, are declared as a group (along with subSRUs) to be subject or not subject to cannibalization. Unlike that of LRUs, subcomponent cannibalization is site-specific and is controlled by a switch on the BASE, CIRF, and DEPT records. The subcomponent record groups (SRU, SSRU, and INDT) are shown in Fig. 14.

For each SRU, the field naming the depot responsible for repair, the level of repair indicator, and the cirf reparability switch may be chosen to yield any desired mix of repair capability among the various echelons. Observe that the pulsed beam generator is reparable at all locations, but the three SRUs indentured to the telephone are not base reparable.

Quantity per aircraft has the same meaning for SRUs as it does for LRUs. Minimum QPA, however, is not specified but inferred from the minimum QPA of the SRU's parent LRU.

Like LRU demand rates, SRU demand rates (both onshore and offshore) are usually given in terms of demands per component flying hour. However, if an SRU's parent LRU is identified as experiencing demands in proportion to sorties flown, then its own demand rates are also interpreted as demands per sortie. Note that an SRU's demand rates may never

---

[2]*Lack of repair* refers to precluding repair at an echelon through the component's level of repair indicator. If a part is base reparable, but repair is not yet available (as specified in the BASE record group), the parts *remain* at the base awaiting repair.

[3]Repair time includes time awaiting maintenance and in work (and time awaiting parts if there are no SRUs).

```
columns
        1         2         3         4         5         6         7
123456789012345678901234567890123456789012345678901234567890123456789012

SRU
PULSED BEAM GEN  DEPO 1 1  2      .00082 .00070  2.0  .70       0.0 1.00
PULSED BEAM GEN X 1.2  .30       1.4       .02  90.  90.     4000
AUTO DIAL MOD    DEPO 2 1  2      .00101 .00087
AUTO DIAL MOD   X 4.0  .70       4.4       .05  90.  90.   150000
AUTO ANSWER SYS  DEPO 2 1  2      .00091 .00081
AUTO ANSWER SYS X 3.0  .50       3.9       .02  90.  90.   100000
MUTE SWITCH      DEPO 2 1  2      .00150 .00142
MUTE SWITCH     X 2.0  .25       2.0            60.  60.    60000
SSRU
MEMORY UNIT      DEPO 2 1  4
MEMORY UNIT     X 3.0  .65       3.5       .02  90.  90.    40000
REDIAL CIRCUIT   DEPO 2 1  2
REDIAL CIRCUIT  X 2.5  .45       3.5       .02  90.  90.    50000
MESSAGE REC      DEPO 2 1  2
MESSAGE REC     X 1.0  .30       1.4       .01  90.  90.    15000
CALL SCREEN DEV  DEPO 2 1  2
CALL SCREEN DEV X 1.5  .35       1.5       .02  90.  90.    20000
INDT
ASTEROID BLASTR 1L
PULSED BEAM GEN  S  1 .4167
CELL MOB PHONE  1L
AUTO DIAL MOD    S  1 .2350
AUTO ANSWER SYS  S  1 .2128
MUTE SWITCH      S  1 .3731
AUTO DIAL MOD   1S
MEMORY UNIT      B  2 .7122 .8608
REDIAL CIRCUIT   B  1 .3450 .3886
AUTO ANSWER SYS 1S
MESSAGE REC      B  1 .4383 .3941
CALL SCREEN DEV  B  1 .6820 .6658
```

Fig. 14—Subcomponent Record Groups, Wartime Analysis

exceed the corresponding demand rates of its parent LRU. Also, SRU demands at the depot are partly determined by the depot replacement fraction (in the INDT record group, discussed below).

Component repair characteristics are described in the same manner for SRUs as they are for LRUs. Because their level of repair indicators preclude base-level repair, the three SRUs indentured to the telephone have only cirf- and depot-related repair data. An SRU's repair data elements are not restricted in any sense by the corresponding data elements of its parent LRU. Thus, the repair time, NRTS rate, and condemnation rate of an SRU bear no particular relation to those of its parent LRU.

SubSRU Descriptions (SSRU). These records are nearly identical to the SRU records, except for the absence of demand rates for subSRUs. Implied subSRU demand rates may be found in the indenture relationship (INDT) records.

Indenture Relationships (INDT). These records describe the configuration of an aircraft in terms of the relationships among its various components. LRUs that have no indentured SRUs are not included within the INDT record group.

Components may be listed in either of two ways: as the parent of some number of subcomponents, or as a subcomponent indentured to another (higher) component. LRUs may only be parents, while subSRUs may only be indentured subcomponents. SRUs may appear in both roles; but in any event, they must always be listed as being indentured to an LRU.

Structurally, an INDT group consists of a series of parent-children component groupings. Each parent is followed immediately by all components directly indentured to it ("grandchildren" components are not included). Two indicators identify each part as parent/child and as LRU/SRU/subSRU. For example, CELL MOB PHONE is the parent LRU to three SRUs: AUTO DIAL MOD, AUTO ANSWER SYS, and MUTE SWITCH. AUTO DIAL MOD is in turn the parent SRU to two subSRUs: MEMORY UNIT and REDIAL CIRCUIT.

Indentured components are accompanied by additional data. The quantity per higher assembly specifies the number of the subcomponent on its parent component (rather than on the entire aircraft). MEMORY UNIT, for instance, has a quantity of two per AUTO DIAL MOD, whereas its quantity per aircraft is four (two MEMORY UNITs on each of two AUTO DIAL MODs per ship).

The depot replacement fraction of an indentured subcomponent may be viewed as the probability that it is removed from its parent after its parent is NRTSed to the depot. (Note that under such circumstances, the subcomponent itself is not considered to have been formally NRTSed.)

The cirf and base replacement fractions apply only to subSRUs. In a manner analogous to that of the depot replacement fraction, these quantities denote the respective probabilities that a subSRU is removed from its parent SRU after both arrive together (and attached) at a cirf or base. As before, only the SRU is presumed to have been NRTSed. Since none of our parent SRUs are subject to base-level repair, the base replacement fraction is irrelevant and left blank.

Stock Levels (STK). Although the starting stock levels at MPC1 do not change, the STK record group now accommodates locations and time-phased deployment of component stocks (Fig. 15). Stock levels may change over time by including multiple records for the part and location. The replacement switch indicates whether the new level replaces (blank or 0), increases (1), or reduces (2) the current level. The new levels go into effect on the day given in columns 77–80. For example, the fuel cell has an initial stock level of 15 at MPC1 that changes to 25 on day 5 (at that location only); the mute switch starts with an initial five at MPC1 but gains three more on day 5.

Output from Program ECHO

The echo output is longer and encompasses greater detail than before, but is essentially the same. Below we point out only the differences between the two echoes.

Administrative Data, Option Selection, and Transportation Times Echo. In Fig. 16, the requested times of analysis are superseded by equal or greater times because of *time scaling*, a feature tied to the compilation parameter DMTIME. This parameter sets the chronologically latest internal time of analysis (set to 30 here). When the latest requested time of analysis is larger than DMTIME, the model compresses the requested time by an integer scaling factor (the latest requested time divided by DMTIME, rounded up to the nearest integer).

```
columns
         1         2         3         4         5         6         7         8
12345678901234567890123456789012345678901234567890123456789012345678901234567890

STK
NAV COMPUTER     MPC1    5 MPC2    8 MPC3    6 DEPO    1
MANEUVER THRUS   MPC1   20 MPC2   30 MPC3   20
FUEL CELL        MPC1   15 MPC2   40 MPC3   15 DEPO    7
DOCKING MECH     MPC1    5 MPC2   10 MPC3    5
TARGET COMPUTER  MPC1    5 MPC2    8 MPC3    5
ASTEROID BLASTR  MPC1    9 MPC2   12 MPC3   10 DEPO    4
PART COLLECTOR   MPC1    5 MPC2   10 MPC3    6
SENSOR           MPC1   30 MPC2   56 MPC3   62 DEPO    2
CONTROL PANEL    MPC1    5 MPC2    9 MPC3    6
CELL MOB PHONE   MPC1   20 MPC2   35 MPC3   22 DEPO    2
PULSED BEAM GEN  MPC1    2 MPC2    3 MPC3    2
AUTO DIAL MOD    MPC1    3 MPC2    4 MPC3    3
AUTO ANSWER SYS  MPC1    1 MPC2    2 MPC3    1
MUTE SWITCH      MPC1    5 MPC2    6 MPC3    6
MEMORY UNIT      CRF1    4 DEPO    1
REDIAL CIRCUIT   CRF1    3
MESSAGE REC      CRF1    3 DEPO    1
CALL SCREEN DEV  CRF1    3
NAV COMPUTER     MPC1    8                                                     5
MANEUVER THRUS   MPC1   25                                                     5
FUEL CELL        MPC1   25                                                     5
TARGET COMPUTER  MPC1    7                                                     5
ASTEROID BLASTR  MPC1   18                                                     5
PART COLLECTOR   MPC1    6                                                     5
SENSOR           MPC1   32                                                     5
CONTROL PANEL    MPC1    8                                                     5
CELL MOB PHONE   MPC1   24                                                     5
PULSED BEAM GEN 1MPC1    3                                                     5
AUTO ANSWER SYS 1MPC1    2 CRF1    2                                           5
MUTE SWITCH     1MPC1    3                                                     5
```

Fig. 15—STK Record Group, Wartime Analysis

In this case, the time scaling factor is 2 (60/30). Except for causing slight deviations among requested times of analysis, time scaling imposes no inconvenience upon the user. Such quantities as repair and resupply times and start times for various activities and resources are treated appropriately and should not be independently adjusted.

The summary of base- and cirf-to-depot transportation times shows that MPC1 and MPC2 are connected to DEPO by instantaneous transportation links. As discussed earlier, these links can never be used since (with respect to all components) both MPCs are served directly by CRF1.

Flying Program Echo. Aside from the inclusion of two additional MPCs and the incorporation of a dynamic program of activity, this echo is no different from its counterpart in the first example and so is omitted here.

```
SCENARIO INCLUDES  3 BASES,   WITH BASE ADMINISTRATIVE DELAY SET TO  0.000;
                  ONE CIRF,   WITH CIRF ADMINISTRATIVE DELAY SET TO  0.500;
                  ONE DEPOT, WITH DEPOT ADMINISTRATIVE DELAY SET TO  0.500.

ANALYSIS REQUESTED FOR TIMES      1    15    30    45    60
DUE TO TIME SCALING, ANALYSIS DONE FOR TIMES     2    16    30    46    60
TIME SCALE FACTOR IS   2.

NUMBER OF BASES OF EACH TYPE --
 MPC1 MPC2 MPC3
   1    1    1

RANDOMIZED REPAIR AND TRANSPORTATION TIMES.
TRANSPORTATION CUTOFFS APPLY ONLY TO FORWARD TRANSPORTATION.

 OPTION
 SELECTED  MEANING
    8      LIST PROBLEM LRUS -- UP TO  5 LRUS.
   11      CALCULATE PERFORMANCE AT  15 PERCENT NFMC BASED ON INPUT OR
           PREVIOUS STOCK.
   23      INCLUDE IN THE PROBLEM LRUS REPORT DATA
           ABOUT THE STATUS OF THE WORST   2 SRUS
           PER PROBLEM LRU, AND THE WORST   2 SUBSRUS PER PROBLEM SRU.

TRANSPORTATION TIMES BETWEEN LOCATIONS --
BASE MPC1 IS SERVED BY CIRF CRF1.
     TRANSPORTATION FROM BASE TO CIRF IS  2.00 DAYS.
     TRANSPORTATION FROM CIRF TO BASE IS  2.00 DAYS.
BASE MPC2 IS SERVED BY CIRF CRF1.
     TRANSPORTATION FROM BASE TO CIRF IS  0.00 DAYS.
     TRANSPORTATION FROM CIRF TO BASE IS  0.00 DAYS.
BASE MPC3 IS NOT SERVED BY A CIRF.

BASE OR CIRF TO DEPOT TRANSPORTATION TIMES, IN DAYS --
 BASE OR CIRF   DEPO
 CRF1           2.00
 MPC1           0.00
 MPC2           0.00
 MPC3           3.50

DEPOT TO BASE OR CIRF TRANSPORTATION TIMES, IN DAYS --
 BASE OR CIRF   DEPO
 CRF1           2.00
 MPC1           0.00
 MPC2           0.00
 MPC3           3.50

NO BASES, CIRFS, OR DEPOTS CANNIBALIZE SRUS AND SUBSRUS.
```

Fig. 16—Administrative Data Echo, Wartime Analysis

Resupply and Repair Characteristics Echo. The matrix-like method for displaying transportation links (Fig. 17) occasionally indicates the availability of nonexistent connections (e.g., transportation between SUPPLIER and MPC1 is established by default on day 0). Because MPC1 relies upon CRF1 for resupply of all Z-455 components, this implied link to SUPPLIER can never be used and may be ignored.

Component Data Echo. The first two tables (Fig. 18) show the components listed by family, with LRUs, SRUs, and subSRUs differentiated by lettered prefixes (L, S, or SS) in their assigned numbers, and by progressively larger indentations from the left-hand margin. Again, data restricted in use to subcomponents are quantity per application and replacement fractions. The cirf and base replacement fractions are additionally restricted in use to

```
PEACETIME FORWARD PIPELINE --
'N' MEANS DOES NOT START TO EMPTY UNTIL CONNECTION HAS BEEN
    RESTABLISHED BETWEEN LOCATION AND HIGHER ECHELON
'Y' MEANS CONTINUES TO EMPTY.

                    DEPOT
                    NAMES
NAME SUPPLIER CIRF  DEPO
MPC1     N      N     N
MPC2     N      N     N
MPC3     N            N
CRF1     N            N
DEPO     N

TIMES WHEN FORWARD TRANSPORTATION IS FIRST ESTABLISHED --
                    DEPOT
                    NAMES
NAME SUPPLIER  CIRF  DEPO
MPC1     0.0    0.0    0.0
MPC2     0.0    0.0    0.0
MPC3     0.0           0.0
CRF1     0.0           0.0
DEPO     0.0

TIMES WHEN FORWARD TRANSPORTATION IS FIRST CUTOFF --
                    DEPOT
                    NAMES
NAME SUPPLIER  CIRF  DEPO
MPC1     0.0    0.0    0.0
MPC2     0.0    0.0    0.0
MPC3     0.0           0.0
CRF1     0.0           0.0
DEPO     0.0

DURATION OF FORWARD TRANSPORTATION CUTOFF --
                    DEPOT
                    NAMES
NAME SUPPLIER  CIRF  DEPO
MPC1     0.0    0.0    0.0
MPC2     0.0    0.0    0.0
MPC3     0.0           0.0
CRF1     0.0           0.0
DEPO     0.0

          SUSTAINED    TIME     SRU AND SUBSRU  O-RR    I-RRR
           DEMAND   PEACETIME      REPAIR      REPAIR   REPAIR
LOCATION   START    REPARABLES     START       START    START
  NAME     TIME      ARRIVE        TIME        TIME     TIME
  MPC1     61.0       1.0           1.0         1.0      1.0
  MPC2     61.0       1.0           1.0         1.0      1.0
  MPC3     61.0       1.0           8.0         8.0      8.0
  CRF1    ------      1.0           5.0         5.0      5.0
  DEPO    ------      1.0           1.0         1.0      1.0
```

Fig. 17—Resupply and Repair Characteristics Echo, Wartime Analysis

subSRUs alone; recall that these act as surrogates for subSRU demand rates, which otherwise go unspecified.

The remaining four tables (and component frequency echo) are omitted here as they appear much the same as before, only expanded to allow for data pertaining to other echelons.

Output from Program REPORT

REPORT produces much output covering a broad span of time that helps in understanding the fleet's capabilities. This information can be useful in suggesting strategies (that may be assessed further with the model) to compensate for potential shortcomings. Here, we only extract the output for a single time of analysis (the end of day 30) to help explain the contents of the various reports.

Performance Report. This report (Fig. 19) again shows the effect of full cannibalization. If the telephone were cannibalized, the total expected number of NFMC ships drops from 28.5 (36 percent of the entire fleet of 80 ships) to 18.8 (about 24 percent). The expected number of sorties varies by only 9, but note the tremendous variation in the probability of achieving those sorties at each MPC: from 4 percent to 100 percent. The fact that results at MPC1 look much like the day 1 results from the peacetime analysis is only coincidence.

Problem LRUs List. This report reveals three LRUs to be the likely causes of problems. The worldwide component impact and target component impact are aggregates of each respective impact across all bases. The worldwide minimum component impact, though, may be defined as the worldwide component impact given an assumption of instantaneous lateral supply among all bases (or, equivalently, an optimal redistribution of existing base-level stocks). For LRUs whose stock levels are evenly distributed among bases (CELL MOB PHONE and FUEL CELL), the availability of instantaneous lateral supply is of scant benefit. But for LRUs that are plentiful at some bases and scarce at others (SENSOR), such a capability could offer substantial rewards. In general, a large difference between the worldwide component impact and the worldwide minimum component impact suggests a maldistribution of stock at the base level.

Selection of an LRU for the problem LRUs list is determined solely by its status at individual bases; even if it is a problem at just one of many bases, it is included on the list (SENSOR, for instance, appears to cause hardship only at MPC3). Within the list, overall ranking is established on the basis of worldwide component impact. The graphic display of the problem LRUs list is similar to that given in Fig. 9 and is omitted here.

The detailed pipeline contents are swelled by the selection of option 23 (subcomponent supplement to the problem LRUs list). *Troublesome* subcomponents are reported when the parent component has a significant AWP pipeline at some location (an AWP pipeline segment is significant if the fraction of the total pipeline it represents exceeds the fraction given as option 23's second parameter—0.15 in this case). Subcomponents may be included only because of the requirement that a certain number be listed for each parent component. However, some troublesome subcomponents may not be identified because of the limit imposed by option 23's first parameter. Figure 20 shows the pipeline segments of the problem parts associated with the telephone.

As before, retrograde pipelines at the MPCs have no meaning and are automatically ignored. MPC administrative pipelines for subcomponents are also undefined, since SRUs and subSRUs can never "arrive" at a base repair shop (whereas an LRU can arrive from the flight line). SubSRU AWP pipelines are ignored because subSRUs represent the lowest level of indenture and have no subcomponents to wait for.

TABLE 1

OVERALL INFORMATION AND INDENTURE RELATIONSHIPS
(AN ASTERISK AFTER THE DEMAND RATES INDICATES THAT DEMAND RATES FOR THAT COMPONENT ARE PER SORTIE RATHER THAN PER FLYING HOUR.)

PART NAME	NUMBER	CAN TEST AT CIRF?	NRTS OR CONDEMN	COST	---DEMANDS PER-- FLYING HOUR ONSHORE	OFFSHORE	LEVEL OF REPAIR	RESUPPLY (DAYS) PEACE	WAR	QUANTITY PER APPLICATION	REPLACEMENT PERCENTAGE DEPOT	CIRF	BASE
NAV COMPUTER	L 1	YES	AFTER TEST	20000.	0.00062	0.00050	BASE	21.0	42.0	----	-----	-----	-----
MANEUVER THRUS	L 2	YES	AFTER TEST	8000.	0.00007	0.00007	BASE	14.0	28.0	----	-----	-----	-----
FUEL CELL	L 3	YES	AFTER TEST	2500.	0.00200	0.00180	BASE	14.0	28.0	----	-----	-----	-----
DOCKING MECH	L 4	YES	AFTER TEST	5000.	0.00006	0.00006*	BASE	14.0	28.0	----	-----	-----	-----
TARGET COMPUTER	L 5	YES	AFTER TEST	30000.	0.00046	0.00040	BASE	28.0	56.0	----	-----	-----	-----
ASTEROID BLASTR	L 6	YES	AFTER TEST	9000.	0.00195	0.00175	BASE	28.0	56.0	----	-----	-----	-----
PULSED BEAM GEN	S 1	YES	AFTER TEST	4000.	0.00082	0.00070	BASE	90.0	90.0	1	0.417	-----	-----
PART COLLECTOR	L 7	YES	AFTER TEST	6000.	0.00038	0.00030	BASE	21.0	42.0	----	-----	-----	-----
SENSOR	L 8	YES	AFTER TEST	15000.	0.00200	0.00020	BASE	35.0	56.0	----	-----	-----	-----
CONTROL PANEL	L 9	YES	AFTER TEST	10000.	0.00061	0.00055	BASE	21.0	42.0	----	-----	-----	-----
CELL MOB PHONE	L 10	YES	AFTER TEST	750000.	0.00410	0.00410	BASE	21.0	42.0	----	-----	-----	-----
AUTO DIAL MOD	S 2	YES	AFTER TEST	150000.	0.00101	0.00087	CIRF	90.0	90.0	1	0.235	-----	-----
MEMORY UNIT	SS 1	YES	AFTER TEST	40000.	-------	-------	CIRF	90.0	90.0	2	0.712	0.861	0.000
REDIAL CIRCUIT	SS 2	YES	AFTER TEST	50000.	-------	-------	CIRF	90.0	90.0	1	0.345	0.389	0.000
AUTO ANSWER SYS	S 3	YES	AFTER TEST	100000.	0.00091	0.00081	CIRF	90.0	90.0	1	0.213	-----	-----
MESSAGE REC	SS 3	YES	AFTER TEST	15000.	-------	-------	CIRF	90.0	90.0	1	0.438	0.394	0.000
CALL SCREEN DEV	SS 4	YES	AFTER TEST	20000.	-------	-------	CIRF	90.0	90.0	1	0.682	0.666	0.000
MUTE SWITCH	S 4	YES	AFTER TEST	60000.	0.00150	0.00142	CIRF	60.0	60.0	1	0.373	-----	-----

TABLE 2

BASE, CIRF, AND DEPOT RELATED LRU, SRU AND SUBSRU DATA

PART NAME	NUMBER	------LONE BASE------ REPAIR TIME (DAYS)	NRTS RATE	CONDEMN RATE	--CIRF-SERVED BASE--- REPAIR TIME (DAYS)	NRTS RATE	CONDEMN RATE	---------CIRF-------- REPAIR TIME (DAYS)	NRTS RATE	CONDEMN RATE	ASSIGNED DEPOT	--------DEPOT--------- REPAIR TIME (DAYS)	REPAIR LIMIT PER DAY	CONDEMN RATE
NAV COMPUTER	L 1	2.00	0.600	0.000	0.00	1.000	0.000	1.50	0.350	0.000	DEPO	1.80	9999.00	0.020
MANEUVER THRUS	L 2	0.50	0.300	0.000	0.00	1.000	0.000	0.50	0.050	0.000	DEPO	0.90	9999.00	0.000
FUEL CELL	L 3	4.00	0.950	0.000	0.00	1.000	0.000	2.50	0.450	0.000	DEPO	4.00	9999.00	0.020
DOCKING MECH	L 4	0.50	0.100	0.000	0.00	1.000	0.000	0.50	0.000	0.000	DEPO	0.50	9999.00	0.000
TARGET COMPUTER	L 5	2.00	0.650	0.000	0.00	1.000	0.000	1.50	0.300	0.000	DEPO	1.60	9999.00	0.020
ASTEROID BLASTR	L 6	2.00	0.750	0.000	0.00	1.000	0.000	1.00	0.350	0.000	DEPO	2.00	9999.00	0.000
PULSED BEAM GEN	S 1	2.00	0.700	0.000	0.00	1.000	0.000	1.20	0.300	0.000	DEPO	1.40	9999.00	0.020
PART COLLECTOR	L 7	2.00	0.200	0.000	0.00	1.000	0.000	1.80	0.050	0.000	DEPO	2.10	9999.00	0.000
SENSOR	L 8	2.00	0.950	0.000	0.00	1.000	0.000	1.60	0.500	0.000	DEPO	2.00	9999.00	0.070
CONTROL PANEL	L 9	2.00	0.700	0.000	0.00	1.000	0.000	1.80	0.250	0.000	DEPO	2.20	9999.00	0.000
CELL MOB PHONE	L 10	4.00	0.990	0.000	0.00	1.000	0.000	3.50	0.600	0.000	DEPO	4.50	9999.00	0.150
AUTO DIAL MOD	S 2	0.00	0.000	0.000	0.00	0.000	0.000	4.00	0.700	0.000	DEPO	4.40	9999.00	0.050
MEMORY UNIT	SS 1	0.00	0.000	0.000	0.00	0.000	0.000	3.00	0.350	0.000	DEPO	3.50	9999.00	0.020
REDIAL CIRCUIT	SS 2	0.00	0.000	0.000	0.00	0.000	0.000	2.50	0.450	0.000	DEPO	3.50	9999.00	0.020
AUTO ANSWER SYS	S 3	0.00	0.000	0.000	0.00	0.000	0.000	3.00	0.500	0.000	DEPO	3.90	9999.00	0.020
MESSAGE REC	SS 3	0.00	0.000	0.000	0.00	0.000	0.000	1.00	0.300	0.000	DEPO	1.40	9999.00	0.010
CALL SCREEN DEV	SS 4	0.00	0.000	0.000	0.00	0.000	0.000	1.50	0.350	0.000	DEPO	1.50	9999.00	0.020
MUTE SWITCH	S 4	0.00	0.000	0.000	0.00	0.000	0.000	2.00	0.250	0.000	DEPO	2.00	9999.00	0.000

Fig. 18—Tables 1 and 2 of Component Data Echo, Wartime Analysis

PERFORMANCE BASED ON STOCK ON HAND ON DAY 30.

BASE	TARG. NFMC	TOTAL ACFT	FULL CANNIBALIZATION: PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	PARTIAL CANNIBALIZATION: PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	TOTAL BACK ORDERS
MPC1	3	20	0.118	0.505	10	6.727	0.336	36.53	2.421	0.001	0.044	7	10.282	0.514	29.00	2.906	38.27
MPC2	6	40	0.448	1.000	28	7.123	0.178	60.00	1.844	0.020	0.997	24	12.237	0.306	59.98	2.171	25.77
MPC3	3	20	0.069	0.876	13	4.958	0.248	39.51	2.315	0.017	0.648	12	6.053	0.303	38.61	2.200	144.40
TOTAL		80				18.808	0.235	136.03					28.573	0.357	127.60		208.45

POTENTIAL PROBLEM PARTS REPORT

WORLD-WIDE AIRCRAFT STATUS - DAY 30
(* - INDICATES AN APPLICATION FRACTION NOT EQUAL TO ONE.)

LRU NAME	NUMBER	WUC	TEQ	WORLD-WIDE: COMP. IMPACT	MINIMUM COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT
CELL MOB PHONE	L 10			16.40	16.23	12.00	MPC1	6.31	3.00	MPC2	6.81	6.00	MPC3	3.28	3.00
FUEL CELL	L 3			7.49	7.48	12.00	MPC1	2.56	3.00	MPC2	1.21	6.00	MPC3	3.72	3.00
SENSOR	L 8			2.80	1.42	12.00	MPC1	0.00	3.00	MPC2	0.00	6.00	MPC3	2.80	3.00

Fig. 19—Performance Report and Problem LRUs List, Wartime Analysis

DETAILED PIPELINE SEGMENT REPORT FOR PROBLEM LRUS, SRUS, AND SUBSRUS

	NUMBER	WUC	COMPONENT IMPACT	MINIMUM COMPONENT IMPACT	TARGET COMPONENT IMPACT
			---WORLD-WIDE---		
LRU CELL MOB PHONE	L	10	16.401	16.233	12.000

LOC.	RETROGRADE	ADMINISTRATIVE	IN TEST/REPAIR	AWAITING PARTS	ON ORDER (NOT RECEIVED)	TOTAL	STOCK	BACKORDERS	BO/QPA
DEPO	12.450	2.567	23.614	12.066	11.712	62.409	2	60.409	
CRF1	4.402	2.385	18.501	7.060	48.953	81.301	0	81.301	
MPC1	----------	0.000	0.000	0.000	36.597	36.597	24	12.622	6.311
MPC2	----------	0.000	0.000	0.000	48.575	48.575	35	13.623	6.812
MPC3	----------	0.000	6.539	0.000	21.753	28.292	22	6.557	3.279

SRU AUTO DIAL MOD S 2

LOC.	RETROGRADE	ADMINISTRATIVE	IN TEST/REPAIR	AWAITING PARTS	ON ORDER (NOT RECEIVED)	TOTAL	STOCK	BACKORDERS	BO/QPA
DEPO	0.756	0.190	5.965	3.918	1.640	12.469	0	12.469	12.469
CRF1	0.000	0.000	2.027	0.026	4.068	6.122	0	6.122	6.122
MPC1	----------	----------	0.000	0.000	0.000	0.000	3	0.000	0.000
MPC2	----------	----------	0.000	0.000	0.000	0.000	4	0.000	0.000
MPC3	----------	----------	0.000	0.000	0.047	0.047	3	0.000	0.000

SUBSRU MEMORY UNIT SS 1

LOC.	RETROGRADE	ADMINISTRATIVE	IN TEST/REPAIR	AWAITING PARTS	ON ORDER (NOT RECEIVED)	TOTAL	STOCK	BACKORDERS	BO/QPA
DEPO	0.361	0.090	6.417	----------	0.609	7.478	1	6.478	3.239
CRF1	0.000	0.000	0.825	----------	0.937	1.762	4	0.046	0.023
MPC1	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000
MPC2	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000
MPC3	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000

SUBSRU REDIAL CIRCUIT SS 2

LOC.	RETROGRADE	ADMINISTRATIVE	IN TEST/REPAIR	AWAITING PARTS	ON ORDER (NOT RECEIVED)	TOTAL	STOCK	BACKORDERS	BO/QPA
DEPO	0.057	0.014	1.504	----------	0.140	1.715	0	1.715	1.715
CRF1	0.000	0.000	0.155	----------	0.161	0.316	3	0.000	0.000
MPC1	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000
MPC2	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000
MPC3	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000

SRU AUTO ANSWER SYS S 3

LOC.	RETROGRADE	ADMINISTRATIVE	IN TEST/REPAIR	AWAITING PARTS	ON ORDER (NOT RECEIVED)	TOTAL	STOCK	BACKORDERS	BO/QPA
DEPO	0.504	0.127	4.520	1.128	0.556	6.834	0	6.834	6.834
CRF1	0.000	0.000	1.392	0.001	1.932	3.325	2	1.517	1.517
MPC1	----------	----------	0.000	0.000	0.000	0.000	3	0.000	0.000
MPC2	----------	----------	0.000	0.000	0.000	0.000	2	0.000	0.000
MPC3	----------	----------	0.000	0.000	0.032	0.032	1	0.001	0.001

SUBSRU MESSAGE REC SS 3

LOC.	RETROGRADE	ADMINISTRATIVE	IN TEST/REPAIR	AWAITING PARTS	ON ORDER (NOT RECEIVED)	TOTAL	STOCK	BACKORDERS	BO/QPA
DEPO	0.059	0.015	0.721	----------	0.083	0.878	1	0.294	0.294
CRF1	0.000	0.000	0.095	----------	0.075	0.170	3	0.000	0.000
MPC1	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000
MPC2	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000
MPC3	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000

SUBSRU CALL SCREEN DEV SS 4

LOC.	RETROGRADE	ADMINISTRATIVE	IN TEST/REPAIR	AWAITING PARTS	ON ORDER (NOT RECEIVED)	TOTAL	STOCK	BACKORDERS	BO/QPA
DEPO	0.118	0.030	1.218	----------	0.263	1.629	0	1.629	1.629
CRF1	0.000	0.000	0.243	----------	0.234	0.477	3	0.002	0.002
MPC1	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000
MPC2	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000
MPC3	----------	----------	0.000	----------	0.000	0.000	0	0.000	0.000

Fig. 20—Telephone Pipeline Contents, Wartime Analysis

Again, the total pipeline contents for a particular component at a particular location are the sum of its individual pipeline segments. However, for components for which NRTS/condemnation decisions and failed subcomponent discoveries are made before the initiation of repair, the total pipeline is the sum of the retrograde, administrative, and on-order segments, plus the maximum of the in-repair and AWP segments. Expected total backorders are calculated for subcomponents just as they are for LRUs.

An LRU's component impact (BO/QPA) is not specified for cirfs and depots because these are not operating locations. For a subcomponent, BO/QPA may be taken to mean the ratio of its expected total backorders to its quantity per application; this corresponds to its expected effect upon its parent component (assuming, of course, that the subcomponent may be cannibalized). The component impact of subcomponents is specified both at the base level and for cirfs and depots.

CONSTRAINED REPAIR

An unstated assumption in the first two examples is that repair facilities at each location are *unconstrained*: always sufficient in capacity to accommodate whatever loads may be placed upon them. Such an assumption is quite adequate in many situations, especially when *repair cycle times* are used in place of repair times in the LRU/SRU/SSRU records.[1] But under nonstationary conditions (such as a wartime surge) in which flying activity, resupply availability, or repair resource levels change, an assumption of unconstrained repair can lead to inaccuracies in projected performance.

Recall that depot-level repair of individual components can be constrained in a rather simplistic fashion by specifying appropriate values for depot repair limit. Dyna-METRIC also has a constrained repair module that represents constrained repair on a more extensive and sophisticated basis. For LRUs only, this module:

- assigns each LRU to a particular type of repair resource,
- adjusts the number of servers of each type over time, and
- accounts for variable server productivity and degradation of server capabilities.

Queues for LRUs awaiting maintenance are determined explicitly as a function of server availability, workload, and other factors; hence, hands-on processing times (rather than repair cycle times) should be specified in the LRU record group.

In this example, the Really Awesome New Doodads Corporation wishes to consider the inherent limitations of their repair facilities. Specifically, they want to observe the effects of finite numbers of servers, restrictions in the amount of time that an individual server may dedicate to actual repair work, and incapacitation of servers. Given the demanding nature of the wartime scenario, they expect a decline in projected performance; they only hope that the new results fall short of an outright refutation of their earlier assumption of unconstrained repair capacity.

Throughout the Z-455 logistics support structure, component repair is performed by three classes of semi-intelligent, programmable test-and-repair androids (TRAs):

[1]A component's repair cycle time is the sum of its actual hands-on processing time, time awaiting maintenance, and time awaiting parts (if it has no explicit indentured subcomponents).

1. TRAAs repair MANEUVER THRUS, DOCKING MECH, ASTEROID BLASTR, PART COLLECTOR, and CONTROL PANEL.
2. TRABs repair NAV COMPUTER, TARGET COMPUTER, and SENSOR.
3. TRACs repair FUEL CELL and CELL MOB PHONE.

Although TRAs are complex and consist of many intricate and exotic components, they are also unusually rugged and reliable. In the rare event of a failure, spare parts may be obtained from an Earth-based supplier.

INPUT DATA SET

The constrained repair module calls for three additional record groups (called TEST, TBED, and TPRT) for *each* repair resource. These are the only record groups that may appear more than once in an input data set. Figure 21 shows the record groups for the three types of TRA.

Constrained Repair Availability (TEST). The TEST record group must appear first; it names the resource and specifies server availability as a function of the number of servers at a location. An isolated server is presumed to divide its operating time between two pursuits: test and repair of LRUs, and self-test. Its availability is defined as the fraction of its operating time that it may spend in test and repair of LRUs. In actual practice with automated test equipment, at a location with many test stands, self-test is frequently replaced by cross-testing on other stands. Typically, this is a much more efficient process; hence, the availability of individual servers often increases as the total number of servers increases. If cross-testing is inapplicable to the repair resource (e.g., human workers), availability may be set to a constant value throughout. TRAAs benefit greatly by operating together; each TRAA in a group of three or more has an availability of 0.80, whereas a lone TRAA has an availability of just 0.55.

Server Levels (TBED). Each TBED record group is associated only with the immediately preceding TEST group. For each location, the TBED record specifies the failure and resupply parameters of the named repair resource, and also gives the time-varying number of servers of that type. Note that even though MPC1 and MPC2 do not have repair, they must be assigned one server for their LRUs to spend zero time on before being NRTSed to CRF1 for repair (because the LRUs were coded NRTS after test).

At DEPO and MPC3, TRAAs have backorder rates (expressed in terms of component backorders per day of active operation), perhaps as a consequence of their frequent exposure to radioactive contaminants on the Z-455 LRUs they repair. A server with failed components is considered partially mission capable and able to test and repair only some specified subset of the LRUs that a fully capable server can handle.[5] Regardless of the number of servers present, Dyna-METRIC never allows more than one server to become PMC at each location. This rule is founded upon two assumptions:

1. All server components may be cannibalized (note that this is unsuited to modeling "backorders" among human workers).
2. The likelihood that a particular component is responsible for incapacitating more than one server is vanishingly small (this has its roots in the extreme complexity of typical automated test equipment).

[5]The subset of reparable LRUs may be designated in the VTM record group, described below under "Treatment of Unusual Data Elements." Without such a designation, a PMC server is taken by default to have the same capabilities as a fully capable server.

```
columns
         1         2         3         4         5         6         7
1234567890123456789012345678901234567890123456789012345678901234567890

TEST
TRAA             .55  .60  .80
TBED
DEPO .0002    7       3
CRF1         14       2  5  5
MPC1                  1
MPC2                  1
MPC3 .0007   14       2
TPRT
MANEUVER THRUS
DOCKING MECH
ASTEROID BLASTR
PART COLLECTOR
CONTROL PANEL
TEST
TRAB             .50  .90
TBED
DEPO .0004    7       3
CRF1         14       2  5  5
MPC1                  1
MPC2                  1
MPC3 .0012   14       2
TPRT
NAV COMPUTER
TARGET COMPUTER
SENSOR
TEST
TRAC             .63  .75  .85
TBED
DEPO .0005    7       3
CRF1         14       2  5  6
MPC1                  1
MPC2                  1
MPC3 .0011   14       2
TPRT
FUEL CELL
CELL MOB PHONE
```

Fig. 21—TEST, TBED and TPRT Record Groups

Resupply of components needed to restore the full capabilities of PMC servers is handled separately from resupply of aircraft components. Each location is considered directly linked to an external supplier of server components, thereby skirting the normal protocol among echelons. Once ordered, server components are received after a delay equal to the resupply time (7 days for DEPO and 14 days for CRF1 and MPC3). As with aircraft components, resupply may be cut off for any span of time; however, resupply start time is automatically taken to be day 1 of the wartime scenario, with peacetime resupply again an implicit assumption. The early availability of resupply is inconsistent with the intra-scenario deployment times of MPC2 and MPC3, but this is of small concern because of the absence of activity before deployment.

Server levels are specified in the same manner as are aircraft levels in the ACFT record group. CRF1, for example, has two TRAAs in peacetime and through day 4 of wartime. On day 5, it receives another three TRAAs, bringing its total allocation to five; this quantity is in effect for the remainder of the scenario.

LRU Repair Assignments (TPRT). TPRT records list the LRUs assigned to the constrained repair resource named in the preceding TEST record. This group applies strictly to the TEST group that it follows.

Because the use of the constrained repair module is the only feature that distinguishes the current exercise from the second example, most of the previously existing portions of the input data set may be carried over without modification. The exception is the LRU record group, where repair times now correspond to actual hands-on processing times. The new repair times are arbitrarily 10 percent of the original repair cycle times.

Program PART

To keep processing problems to a minimum we run the partitioner program to group LRUs with their associated repair resources. The only effect this has on any output report is the reordering of parts. For example, LRU 1 is now MANEUVER THRUS because it was the first LRU listed under the first named repair resource (TRAA).

Output from Program ECHO

ECHO includes an additional report detailing the constrained repair resources. Information from the TEST and TBED record groups are echoed here, whereas information from the TPRT record groups are incorporated into the third ECHO table. Figure 22 presents the new echo report for the TRAAs. Note that there are six *alphas* (availability factors) even though only three were specified. The last specified factor is repeated until DMCHANGE number of factors is given.

Output from Program REPORT

The substitution of constrained repair for Dyna-METRIC's normal representation of repair can have dramatic effects on projected performance, particularly in dynamic scenarios. If constrained repair resources are plentiful, and if the differences between LRU repair cycle times and hands-on processing times are large, performance can improve substantially. Alternatively, if repair resources are tightly constrained and activity levels soar, performance can suffer. Our constrained repair parameters yield a significant drop in most performance measures on day 30 (Fig. 23).

The component impacts of the two driving LRUs are considerably higher than before, and so cause reductions of 7 percent (under full cannibalization) and 23 percent (under partial cannibalization) in the expected number of expeditions that may be undertaken. The inadequacy of TRAC server levels may be seen in a comparison of the current *in test/repair* pipeline segments for the telephone with those from the second example (Fig. 20).

```
DETAILED INFORMATION--  TEST STANDS

CONSTRAINED REPAIR RESOURCE TRAA
 ALPHA -- 0.550 0.600 0.800 0.800 0.800 0.800
   NUMBER OF STANDS --
      B.O.  RSPLY
      RATE  TIME
MPC1 0.0000  0.
  1 STANDS UNTIL DAY  61

MPC2 0.0000  0.
  1 STANDS UNTIL DAY  61

MPC3 0.0007 14.
  2 STANDS UNTIL DAY  61

CRF1 0.0000 14.
  2 STANDS UNTIL DAY   5
  5 STANDS UNTIL DAY  61

DEPO 0.0002  7.
  3 STANDS UNTIL DAY  61
```

Fig. 22—Echo for TRAA Constrained Repair Resource

TREATMENT OF UNUSUAL DATA ELEMENTS

For many Dyna-METRIC users, the features just discussed are sufficient to conduct meaningful capability assessments. What remains to be examined are the features that address important but often unusual issues. Described below are the input record groups for:

- aircraft attrition (ATTR)
- site-specific application fractions (APPL)
- site-specific QPAs (QPA)
- changing mission requirements (MESL)
- alternative maintenance concepts (ILM)
- non-Poisson removal processes (VTM).

Each group is, in turn, incorporated into the second example (dynamic analysis) to highlight its effect through comparison with the previous results.

Attrition Rates (ATTR)

ATTR records describe the aircraft attrition rates as a function of time at each operating base. Peacetime attrition is disallowed; otherwise, over the infinite span of peacetime, a nonzero attrition rate would yield the absurd result of infinite attrition.

PERFORMANCE BASED ON STOCK ON HAND ON DAY 30.

			FULL CANNIBALIZATION							PARTIAL CANNIBALIZATION							
BASE	TARG. NFMC	TOTAL ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	TOTAL BACK ORDERS
MPC1	3	20	0.025	0.197	7	9.310	0.465	31.04	2.712	0.000	0.001	4	13.576	0.679	19.27	2.999	43.54
MPC2	6	40	0.032	0.958	21	13.560	0.339	59.71	2.228	0.000	0.326	14	21.912	0.548	52.70	2.620	41.68
MPC3	3	20	0.001	0.410	9	7.537	0.377	35.45	2.478	0.000	0.017	6	11.224	0.561	26.28	2.962	181.50
TOTAL		80				30.406	0.380	126.21					46.712	0.584	98.25		266.72

POTENTIAL PROBLEM PARTS REPORT
WORLD-WIDE AIRCRAFT STATUS - DAY 30

(* - INDICATES AN APPLICATION FRACTION NOT EQUAL TO ONE.)

				WORLD-WIDE												
LRU NAME	NUMBER	WUC	TEQ	COMP. IMPACT	MINIMUM COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT
CELL MOB PHONE	L 10		TRAC	29.35	29.35	12.00	MPC1	9.03	3.00	MPC2	13.31	6.00	MPC3	7.01	3.00
FUEL CELL	L 9		TRAC	7.85	7.85	12.00	MPC1	2.55	3.00	MPC2	1.51	6.00	MPC3	3.80	3.00
SENSOR	L 8		TRAB	3.59	1.86	12.00	MPC1	0.00	3.00	MPC2	0.00	6.00	MPC3	3.59	3.00

DETAILED PIPELINE SEGMENT REPORT FOR PROBLEM LRUS, SRUS, AND SUBSRUS

			WORLD-WIDE		
	NUMBER	WUC	COMPONENT IMPACT	MINIMUM COMPONENT IMPACT	TARGET COMPONENT IMPACT
LRU CELL MOB PHONE	L 10		29.353	29.347	12.000

LOC.	RETROGRADE	ADMINISTRATIVE	IN TEST/REPAIR	AWAITING PARTS	ON ORDER (NOT RECEIVED)	TOTAL	STOCK	BACKORDERS	BO/QPA
DEPO	10.128	2.177	59.646	2.415	4.212	78.577	2	76.577	
CRF1	4.402	2.385	21.195	7.448	66.159	101.589	0	101.589	
MPC1	----------	0.000	0.000	0.000	42.066	42.066	24	18.068	9.034
MPC2	----------	0.000	0.000	0.000	61.617	61.617	35	26.617	13.309
MPC3	----------	0.000	24.080	0.000	11.931	36.011	22	14.021	7.011

Fig. 23—Output from REPORT, Constrained Repair

Suppose the Really Awesome New Doodads Corporation has run afoul of a militant group of alien environmentalists who are outraged by what they perceive to be a commercially motivated desecration of the asteroid belt. Although the environmentalists have been tolerant of MPC1's modest peacetime activities, management expects them to take violent action in response to the deployment of the other two MPCs. Only MPC3's fleet is considered to be safe from loss (the aliens are apparently allergic to radiation and have no recourse to long-range weapons).

The ATTR record group (Fig. 24) shows that at MPC1, alien incursions are expected to destroy an average of one Z-455 per 1000 expeditions over the first four days of wartime. These attacks are expected to double in intensity for the next ten days as MPC2 and MPC3 commence operations. Meanwhile, MPC2 is expected to lose three ships per 1000 expeditions from the time it starts operations (day 5) through day 14. At both locations, management thinks the aliens will have exhausted their energies by day 15 and will abandon the fight in order to pursue worthier causes.

Dyna-METRIC accounts for attrition losses on the basis of expected value. Losses are reflected in ECHO's summary of flying program requirements and restrictions at each base (Fig. 25). Note that the level of activity per ship undergoes no adjustment in order to compensate for attrition; thus, attrition losses invariably imply a corresponding reduction in a base's total required program of activity.

Although aircraft attrition losses should always result in a net reduction in total sortie generation, their effect is occasionally mitigated by improvements in the situation of the remainder of the system. In the case of MPC1, the logistics resources that originally were intended to support 20 ships are required to support only 18 after a period of 11 days. Furthermore, computed daily component demands eventually fall to 90 percent (18 ships out of 20) of their nonattrition levels, thereby yielding smaller pipeline quantities and fewer expected backorders. Indeed, in exceptionally rigorous scenarios, these "benefits" of attrition are often shown (fallaciously) to outweigh the effects of the losses themselves. This is illustrated in Fig. 26 in the day 30 performance report.

A comparison of this report with its counterpart for the second example (Fig. 19) reveals that under both cannibalization policies, the expected number and percentage of NFMC ships at MPC1 and MPC2 decrease. Expected sorties show marginal reductions at all MPCs under full cannibalization, but rise at MPC1 under partial cannibalization. We are thus confronted by a paradox: The total number of degraded ships improves but their overall performance decreases after incurring attrition losses. This unusual and inaccurate result confirms the

```
columns
         1         2         3         4         5         6         7
1234567890123456789012345678901234567890123456789012345678901234567890

ATTR
MPC1 .001    5 .002   15 .000
MPC2 .003   15 .000
```

Fig. 24—ATTR Record Group

FLYING PROGRAM FOR BASE -- MPC1(OFFSHORE)

DAYS		AIRCRAFT	SORTIES	FH/SRT	MAX. SORTIES	M1	M2	M3	M4	M5
PEACE		20	2.00	4.00	3.00					
1		20	3.00	5.00	3.00	X	X	X	X	X
2-	11	19	3.00	5.00	3.00	X	X	X	X	X
12-	29	18	3.00	5.00	3.00	X	X	X	X	X
30-	60	18	2.00	5.00	3.00	X	X	X	X	X

FLYING PROGRAM FOR BASE -- MPC2(OFFSHORE)

DAYS		AIRCRAFT	SORTIES	FH/SRT	MAX. SORTIES	M1	M2	M3	M4	M5
PEACE		20	0.00	5.00	3.00					
1-	4	20	0.00	5.00	3.00	X	X	X	X	X
5		20	3.00	5.00	3.00	X	X	X	X	X
6-	9	19	3.00	5.00	3.00	X	X	X	X	X
10		39	2.70	5.00	3.00	X	X	X	X	X
11-	13	38	2.70	5.00	3.00	X	X	X	X	X
14-	24	37	2.70	5.00	3.00	X	X	X	X	X
25-	60	37	1.50	5.00	3.00	X	X	X	X	X

Fig. 25—Flying Program Echo, Attrition

hazards of using Dyna-METRIC to analyze overly demanding scenarios without first activating option 20 (described later).

Although its ships suffer no attrition losses, MPC3 experiences slight reductions in projected performance. On the problem LRUs list, CELL MOB PHONE has a slightly larger component impact at MPC3 with the introduction of attrition at the other MPCs. This outcome may be attributed to the allocation to MPC3 of a larger share of the component backorders at DEPO. The fact that DEPO's expected telephone backorders decrease as a result of attrition losses may explain the insignificant size of the increase in component impact at MPC3.

Application Fractions (APPL)

APPL records specify the fraction of aircraft at each base that contain a particular LRU. Such fractions hold for both peacetime and wartime and need only be specified for those LRUs with fractions other than 1.0 at some base. Suppose that only part of the overall Z-455 fleet is outfitted with fuel cells and targeting computers. Figure 27 shows, for example, that only 90 percent of the ships at MPC1 (or 18 ships) have these LRUs. The absence of records for other LRUs implies that each has an application fraction of 1.0 at all MPCs. The application fractions are summarized in Table 5 of the output from ECHO.

Since an application fraction of less than 1.0 decreases the number of units in service (and the number of units subject to failure), we would expect reduced pipelines for FUEL CELL and TARGET COMPUTER and improved performance for the fleet as a whole (compared with the results in Fig. 19).

Figure 28 gives the new performance report and problem LRUs list for day 30. Such performance measures as E(NFMC) and E(SORTIES) show only marginal improvements, if any.

PERFORMANCE BASED ON STOCK ON HAND ON DAY 30.

			FULL CANNIBALIZATION							PARTIAL CANNIBALIZATION							
BASE	TARG. NFMC	TOTAL ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC-90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC-90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	TOTAL BACK ORDERS
MPC1	3	18	0.309	0.728	9	5.155	0.286	34.04	2.330	0.018	0.288	8	7.700	0.428	30.22	2.529	29.80
MPC2	6	37	0.643	1.000	27	5.515	0.149	55.50	1.780	0.136	0.999	24	9.389	0.254	55.50	2.026	18.80
MPC3	3	20	0.068	0.872	13	4.974	0.249	39.49	2.314	0.016	0.637	12	6.100	0.305	38.55	2.205	144.70
TOTAL		75				15.645	0.209	129.03					23.190	0.309	124.26		193.30

POTENTIAL PROBLEM PARTS REPORT

WORLD-WIDE AIRCRAFT STATUS - DAY 30
(* - INDICATES AN APPLICATION FRACTION NOT EQUAL TO ONE.)

				WORLD-WIDE												
LRU NAME	NUMBER	WUC	TEQ	COMP. IMPACT	MINIMUM COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT
CELL MOB PHONE	L 10			13.09	12.83	11.25	MPC1	4.61	2.70	MPC2	5.17	5.55	MPC3	3.31	3.00
FUEL CELL	L 3			6.62	6.58	11.25	MPC1	2.06	2.70	MPC2	0.85	5.55	MPC3	3.72	3.00
SENSOR	L 8			2.80	1.34	11.25	MPC1	0.00	2.70	MPC2	0.00	5.55	MPC3	2.80	3.00

Fig. 26—Output from REPORT, Attrition

```
columns
              1         2         3         4         5         6         7
1234567890123456789012345678901234567890123456789012345678901234567890

APPL
FUEL CELL        MPC1   .90 MPC2   .85 MPC3   .75
TARGET COMPUTER  MPC1   .90 MPC2   .85 MPC3   .90
```

Fig. 27—APPL Record Group

The absence of more substantial improvements may be attributed largely to the overwhelming influence of the telephone. However, the problem LRUs list shows a pronounced difference in the effect of FUEL CELL—its worldwide component impact is cut by 35 percent.

As before, TARGET COMPUTER does not appear on the problem LRUs list.[6] Its contribution to degraded performance can be only negligible at best when compared with those of the telephone and the fuel cell.

Quantity per Aircraft (QPA)

The QPA record group accommodates site-specific component quantities per aircraft (and for LRUs, minimum quantities per aircraft). There is no mechanism for accommodating site-specific quantities per LRU or SRU. QPA records are entered for a subcomponent only to maintain a state of compatibility with its parent. Values on QPA records override those on component description records; however, if a component has no QPA record, or if a base is absent from a QPA record, the missing values are taken by default to be the quantities given in the appropriate component description record. QPA data are summarized in Table 4 of the output from ECHO.

The LRU QPA affects the computation of performance measures in two ways. As the quantity increases or decreases, so do the component's flying hours, hence demands, pipeline contents, and expected number of backorders. The reverse happens to the component impact associated with a particular number of backorders (recall that component impact is given by the ratio of the LRU's backorders to its QPA). Clearly, these two properties are offsetting to some extent; however, given a constant level of stock, the net effect of a higher QPA is generally a higher component impact.

Miscellaneous Record Groups

The remaining three record groups (VTM, ILM, and MESL) are closely tied to one another. By themselves, the maintenance types defined in the ILM group are inoperative unless VTM records are used to assign each LRU to a particular type of repair. Similarly, the mission support profiles contained in the MESL records must be used in conjunction with the LRU mission essentiality switches specified in the VTM group. Below we describe the record groups individually but consider the joint effect of VTM records with ILM and MESL records.

[6]Detailed pipeline reports for all parts may be obtained with option 15 (described below in "Additional Capability Assessment Options").

PERFORMANCE BASED ON STOCK ON HAND ON DAY 30.

			FULL CANNIBALIZATION							PARTIAL CANNIBALIZATION							
BASE	TARG. NFMC	TOTAL ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC-90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC-90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	TOTAL BACK ORDERS
MPC1	3	20	0.146	0.505	10	6.668	0.333	36.53	2.413	0.001	0.044	7	10.282	0.514	29.00	2.906	32.95
MPC2	6	40	0.448	1.000	28	7.083	0.177	60.00	1.843	0.020	0.997	24	12.237	0.306	59.98	2.171	18.50
MPC3	3	20	0.335	0.877	13	4.454	0.223	39.51	2.314	0.081	0.649	12	5.888	0.294	38.61	2.219	130.91
TOTAL		80				18.205	0.228	136.03					28.408	0.355	127.60		182.36

POTENTIAL PROBLEM PARTS REPORT

WORLD-WIDE AIRCRAFT STATUS - DAY 30
(* - INDICATES AN APPLICATION FRACTION NOT EQUAL TO ONE.)

				WORLD-WIDE											
LRU NAME	NUMBER	WUC	TEQ	COMP. IMPACT	MINIMUM COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT
CELL MOB PHONE	L 10			16.40	16.23	12.00	MPC1	6.31	3.00	MPC2	6.81	6.00	MPC3	3.28	3.00
FUEL CELL	L 3			4.88	4.77	10.05*	MPC1	2.03	2.70*	MPC2	0.49	5.10*	MPC3	2.37	2.25*
SENSOR	L 8			2.80	1.42	12.00	MPC1	0.00	3.00	MPC2	0.00	6.00	MPC3	2.80	3.00

Fig. 28—Output from REPORT, Application Fractions

Variance-to-Mean Data (VTM). The VTM record group serves a multitude of purposes, including the specification of each LRU's:

- maintenance type (in conjunction with the ILM record group);
- wartime demand rates (in conjunction with the LRU record group);
- pipeline variance-to-mean ratio;
- reparability on PMC servers (in conjunction with the constrained repair module);
- mission essentiality (in conjunction with the MESL record group);
- sustained demand rates.

Figure 29 shows the VTM record group for the LRUs of the Z-455. The target computer, for instance, utilizes type 1 maintenance (or repair) resources. Its demand rates in wartime (at both onshore and offshore bases) are identical to its peacetime demand rates (since the two demand rate multipliers are set to 1.0). The variance-to-mean ratio for the number of target computers in pipelines is 1.0; consequently, Dyna-METRIC considers target computer pipeline values to be drawn from a Poisson distribution.[7] The probability that a failed computer can be repaired by a PMC server (when the constrained repair module is being used) is 1.0. The target computer is essential only to the execution of the second mission. Finally, the sustained demand rates are left blank.

Note that VTM records could be omitted for MANEUVER THRUS and the other LRUs that are set to default values: maintenance type 0, wartime demand rate multipliers of 1.0, VTM ratio of 1.0, 100 percent probability of repair by a PMC server, essential to the execution of all missions, and no sustained demand rates. The contents of the VTM record group are summarized in Table 3 of th. output from ECHO.

Maintenance Deployment (ILM). The ILM record group offers a convenient method for partitioning LRUs into different classes for which repair availability may differ. It consists

```
columns
         1         2         3         4         5         6         7
1234567890123456789012345678901234567890123456789012345678901234567890

VTM
NAV COMPUTER     1  1.0  1.0  1.0  1.0 11000
MANEUVER THRUS   0  1.0  1.0  1.0  1.0 11000
FUEL CELL        2  1.0  1.0  1.0  1.0 11000
DOCKING MECH     0  1.0  1.0  1.0  1.0 11000
TARGET COMPUTER  1  1.0  1.0  1.0  1.0 01000
ASTEROID BLASTR  0  1.0  1.0  1.0  1.0 11000
PART COLLECTOR   0  1.0  1.0  1.0  1.0 10000
SENSOR           1  1.0  1.0  1.0  1.0 11000
CONTROL PANEL    0  1.0  1.0  1.0  1.0 11000
CELL MOB PHONE   2  1.0  1.0  1.0  1.0 01000
```

Fig. 29—VTM Record Group

[7]A variance-to-mean ratio less than 1.0 corresponds to a binomial distribution, and a variance-to-mean ratio greater than 1.0 corresponds to a negative binomial distribution.

of a single record that names the LRU maintenance (repair) types, followed by a series of records that give the maintenance type's start time at each location during wartime (as usual, peacetime availability at all locations is assumed automatically). The BASE, CIRF, and DEPT records provide for two maintenance types, RR and RRR.[8] The ILM records override any values set there and are mostly used to define more than two maintenance types. Start times are summarized in the echo of resupply and repair characteristics.

Figure 30 shows the ILM records and the VTM records from before (with VTM records that hold default values omitted and all LRUs considered essential for both missions). Besides RR (type 0) maintenance, there are two Special Maintenance Types. An LRU assigned to maintenance type 0 on its VTM record may enter repair immediately at DEPO, on day 5 at CRF1, and on day 8 at MPC3. LRUs assigned to type 2 (SMT2), however, are irreparable at all locations for four weeks (until day 29).

LRUs that are temporarily irreparable at a particular location are treated in a different manner than those specified as being always irreparable at certain echelons (by using the level of repair indicator in the LRU record, for instance). Temporarily irreparable LRUs are accumulated in local repair pipelines until repair of the appropriate type becomes available.

The output for day 30 (Fig. 31) shows that performance is severely degraded at MPC1 and MPC2 because of shortages of fuel cells and telephones (the type 2 LRUs for which repair was delayed until day 29). As a consequence of that delay, the *in test/repair* pipeline segment at MPC1 appears to contain more telephones than were ever allocated to MPC1 (40 installed on ships and 24 spares). This is because Dyna-METRIC computes component failures on the basis of the required flying program whether or not it is achievable. Option 20 (described later) dispels this paradox by adjusting the flying program.

```
columns
         1         2         3         4         5         6         7
1234567890123456789012345678901234567890123456789012345678901234567890

ILM
      RR    SMT1  SMT2
DEPO   1     1     29
CRF1   5     5     29
MPC3   8    12     29
       .
       .
       .
VTM
NAV COMPUTER       1   1.0   1.0   1.0   1.0 11000
FUEL CELL          2   1.0   1.0   1.0   1.0 11000
TARGET COMPUTER    1   1.0   1.0   1.0   1.0 11000
SENSOR             1   1.0   1.0   1.0   1.0 11000
CELL MOB PHONE     2   1.0   1.0   1.0   1.0 11000
```

Fig. 30—ILM Record Group

[8]Although these maintenance types frequently have names that imply a particular mode of handling (for instance, RR means Remove and Replace), Dyna-METRIC endows none of them with any unique characteristics.

PERFORMANCE BASED ON STOCK ON HAND ON DAY 30.

			FULL CANNIBALIZATION							PARTIAL CANNIBALIZATION							
BASE	TARG. NFMC	TOTAL ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	TOTAL BACK ORDERS
MPC1	3	20	0.000	0.000	0	19.956	0.998	0.13	3.000	0.000	0.000	0	19.993	1.000	0.02	3.000	181.34
MPC2	6	40	0.000	0.028	4	29.458	0.736	31.52	2.960	0.000	0.000	1	36.616	0.915	10.15	3.000	184.74
MPC3	3	20	0.000	0.020	8	9.260	0.463	32.18	2.942	0.000	0.000	5	12.226	0.611	23.32	2.999	192.80
TOTAL		80				58.674	0.733	63.83					68.836	0.860	33.49		558.89

POTENTIAL PROBLEM PARTS REPORT

WORLD-WIDE AIRCRAFT STATUS - DAY 30
(* - INDICATES AN APPLICATION FRACTION NOT EQUAL TO ONE.)

				WORLD-WIDE											
LRU NAME	NUMBER	WUC	TEQ	COMP. IMPACT	MINIMUM COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT
CELL MOB PHONE	L 10			65.05	65.05	12.00	MPC1	28.02	3.00	MPC2	29.24	6.00	MPC3	7.79	3.00
FUEL CELL	L 3			32.63	32.63	12.00	MPC1	12.53	3.00	MPC2	12.63	6.00	MPC3	7.47	3.00
SENSOR	L 8			2.85	1.48	12.00	MPC1	0.00	3.00	MPC2	0.00	6.00	MPC3	2.84	3.00

DETAILED PIPELINE SEGMENT REPORT FOR PROBLEM LRUS, SRUS, AND SUBSRUS

			WORLD-WIDE		
	NUMBER	WUC	COMPONENT IMPACT	MINIMUM COMPONENT IMPACT	TARGET COMPONENT IMPACT
LRU CELL MOB PHONE	L 10		65.048	65.046	12.000

LOC.	RETROGRADE	ADMINISTRATIVE	IN TEST/REPAIR	AWAITING PARTS	ON ORDER (NOT RECEIVED)	TOTAL	STOCK	BACKORDERS	BO/QPA
DEPO	37.688	4.760	15.062	2.264	1.907	61.682	2	59.682	
CRF1	4.402	2.385	96.510	12.331	46.952	162.580	0	162.580	
MPC1	----------	0.000	0.000	0.000	80.033	80.033	24	56.033	28.017
MPC2	----------	0.000	0.000	0.000	93.480	93.480	35	58.480	29.240
MPC3	----------	0.000	23.470	0.000	14.107	37.578	22	15.583	7.791

Fig. 31—Output from REPORT, Maintenance Deployment

Mission Requirements (MESL). This record group specifies changing mission support profiles at operating bases. Although it gives some indication of mission activity at each base, neither it nor any other group contains an explicit mechanism for specifying how many missions of each type are flown on a daily basis. The primary function of the MESL group (in conjunction with the VTM record group) is to identify and discount nonessential LRUs before the computation of performance reports, problem LRUs lists, and the like. ECHO summarizes each base's mission support profile as part of its flying program echo.

Suppose that Z-455s may undertake two missions: asteroidal debris collection (M1) and alien attack suppression (M2). Figure 32 shows the MESL records and the original VTM records (with the maintenance type in column 18 ignored). In peacetime, MPC1 supports only M1; starting day 1 (the outbreak of hostilities), it supports M2 as well.[9] By day 30, M2 becomes unnecessary and MPC1 reverts to supporting just M1. MPC2 has a similar mission support profile except for the days on which they go into effect. MPC3, though, supports only M1 (since its Z-455s are essentially immune to alien attack).

When taken together, these MESL and VTM groups specify that, beginning on day 30 at MPC1 (and on day 25 at MPC2), both the targeting computer and the telephone become entirely superfluous to Z-455 operations; furthermore, they are always superfluous to operations at MPC3. The particle collector is nonessential to the second mission but is always essential to the first mission at all MPCs; consequently, it is never entirely superfluous at any location.

```
columns
         1         2         3         4         5         6         7
1234567890123456789012345678901234567890123456789012345678901234567890

MESL
MPC1 10000   111000 3010000
MPC2 11000 2510000
MPC3 10000
         .
         .
         .
VTM
NAV COMPUTER        1.0  1.0  1.0      11000
MANEUVER THRUS      1.0  1.0  1.0      11000
FUEL CELL           1.0  1.0  1.0      11000
DOCKING MECH        1.0  1.0  1.0      11000
TARGET COMPUTER     1.0  1.0  1.0      01000
ASTEROID BLASTR     1.0  1.0  1.0      11000
PART COLLECTOR      1.0  1.0  1.0      10000
SENSOR              1.0  1.0  1.0      11000
CONTROL PANEL       1.0  1.0  1.0      11000
CELL MOB PHONE      1.0  1.0  1.0      01000
```

Fig. 32—MESL Record Group

[9]Our compilation of the model provides space for five missions. See the discussion of the DMMISSNS parameter in App. B for details.

Although an LRU may be nonessential, it is nevertheless considered to be always onboard the aircraft. It undergoes the same removal, repair, and resupply processes that it would if it were mission essential. In the constrained repair module, nonessential LRUs still compete for repair resources on an equivalent footing with essential LRUs. The only effect of the LRU's nonessentiality at a particular base is that it makes no contribution to degradation of performance at that base, regardless of its component impact.

REPORT's output for day 30 (Fig. 33), compared with that of the second example (Fig. 19), shows how sortie generation increases dramatically at MPC1 and MPC2, where the telephone is normally the chief culprit for degraded performance. At MPC3, however, the telephone is not usually a high driver; hence, performance gains are only marginal. As expected, the statistics for full cannibalization and partial cannibalization become identical because of the discounting of the telephone (the only LRU not subject to cannibalization).

Although it is no longer a true problem LRU, the telephone continues to be listed in the problem LRUs list; furthermore, Dyna-METRIC continues to account for its failures, repair, and resupply. This is done only for purposes of information, however, and should not be construed as an indication of the telephone's ongoing contribution to degraded performance.

ADDITIONAL CAPABILITY ASSESSMENT OPTIONS

Just as it has unusual data elements, Dyna-METRIC also has a number of capability assessment options that are more specialized than such workhorses as option 11 and option 8. Discussed below are:

- detailed parts disposition (option 15)
- achievable sorties on PMC and FMC aircraft (option 20)
- constrained depot workload report (option 1)
- daily demands report (option 18)
- depot workload report (option 7).

Detailed Parts Disposition (Option 15)

Option 15 produces a report similar in content to the detailed pipeline segments section of the problem LRUs list (option 8). Rather than listing the location-by-location pipelines for each part, it lists part-by-part pipelines for each location. Option 15 has no associated parameters.

The detailed parts disposition at DEPO and MPC1 for day 30 of the second example is given in Fig. 34. (The similar segments for CRF1 and the other MPCs are omitted.) This report is somewhat sparser in detail than its counterpart in the problem LRUs list. No mention is made of component impacts. Nevertheless, the pipeline segments themselves are identical and may be of use to the analyst who is curious about the status of components not contained in the list of problem LRUs.

Achievable Sorties for PMC and FMC Aircraft (Option 20)

Under normal circumstances, Dyna-METRIC computes component removals strictly on the basis of required flying activity; whether the required program is actually achievable is academic. This technique is adequate when applied to undemanding scenarios in which projected sortie generation either meets or comes close to meeting stipulated requirements. It

PERFORMANCE BASED ON STOCK ON HAND ON DAY 30.

BASE	TARG. NFMC	TOTAL ACFT	FULL CANNIBALIZATION: PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	PARTIAL CANNIBALIZATION: PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	TOTAL BACK ORDERS
MPC1	3	20	0.756	1.000	16	3.015	0.151	40.00	2.358	0.756	1.000	16	3.015	0.151	40.00	2.358	38.27
MPC2	6	40	1.000	1.000	37	1.655	0.041	60.00	1.565	1.000	1.000	37	1.655	0.041	60.00	1.565	25.77
MPC3	3	20	0.132	0.999	15	4.221	0.211	40.00	2.441	0.132	0.999	15	4.221	0.211	40.00	2.441	144.40
TOTAL		80				8.890	0.111	140.00					8.890	0.111	140.00		208.45

POTENTIAL PROBLEM PARTS REPORT

WORLD-WIDE AIRCRAFT STATUS - DAY 30
(* - INDICATES AN APPLICATION FRACTION NOT EQUAL TO ONE.)

LRU NAME	NUMBER	WUC	TEQ	WORLD-WIDE: COMP. IMPACT	MINIMUM COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT
CELL MOB PHONE	L 10			16.40	16.23	12.00	MPC1	6.31	3	MPC2	6.81	6.00	MPC3	3.28	3.00
FUEL CELL	L 3			7.49	7.48	12.00	MPC1	2.56		MPC2	1.21	6.00	MPC3	3.72	3.00
SENSOR	L 8			2.80	1.42	12.00	MPC1	0.00		MPC2	0.00	6.00	MPC3	2.80	3.00

Fig. 33—Output from REPORT, Mission Essential

DETAILED PIPELINE SEGMENT REPORT AT DEPOT DEPO ON DAY 30--

PART NAME	NUMBER	RETRO.	ADMIN.	IN TEST	AWP	ORDERED	TOTAL	STOCK	BACKORDERS
NAV COMPUTER	L 1	0.97	0.19	0.72	0.00	0.14	2.02	1	1.15
MANEUVER THRUS	L 2	0.34	0.06	0.11	0.00	0.00	0.51	0	0.51
FUEL CELL	L 3	24.01	4.70	38.82	0.00	2.52	70.05	7	63.05
DOCKING MECH	L 4	0.00	0.00	0.00	0.00	0.00	0.00	0	0.00
TARGET COMPUTER	L 5	0.36	0.07	0.23	0.00	0.06	0.72	0	0.72
ASTEROID BLASTR	L 6	3.58	0.69	2.89	1.73	0.00	8.91	4	4.94
PART COLLECTOR	L 7	0.15	0.03	0.11	0.00	0.00	0.28	0	0.28
SENSOR	L 8	52.38	8.02	32.32	0.00	16.80	109.53	2	107.53
CONTROL PANEL	L 9	0.96	0.18	0.82	0.00	0.00	1.95	0	1.95
CELL MOB PHONE	L 10	12.45	2.57	23.61	12.07	11.71	62.41	2	60.41
PULSED BEAM GEN	S 1	0.56	0.12	1.23	0.00	0.47	2.38	0	2.38
AUTO DIAL MOD	S 2	0.76	0.19	5.96	3.92	1.64	12.47	0	12.47
AUTO ANSWER SYS	S 3	0.50	0.13	4.52	1.13	0.56	6.83	0	6.83
MUTE SWITCH	S 4	0.44	0.11	3.76	0.00	0.00	4.32	0	4.32
MEMORY UNIT	SS 1	0.36	0.09	6.42	0.00	0.61	7.48	1	6.48
REDIAL CIRCUIT	SS 2	0.06	0.01	1.50	0.00	0.14	1.71	0	1.71
MESSAGE REC	SS 3	0.06	0.02	0.72	0.00	0.08	0.88	1	0.29
CALL SCREEN DEV	SS 4	0.12	0.03	1.22	0.00	0.26	1.63	0	1.63

DETAILED PIPELINE SEGMENT REPORT AT BASE MPC1 ON DAY 30--

PART NAME	NUMBER		ADMIN.	IN TEST	AWP	ORDERED	TOTAL	STOCK	BACKORDERS
NAV COMPUTER	L 1	-------	0.00	0.00	0.00	1.96	1.96	8	0.00
MANEUVER THRUS	L 2	-------	0.00	0.00	0.00	1.52	1.52	25	0.00
FUEL CELL	L 3	-------	0.00	0.00	0.00	50.64	50.64	25	25.64
DOCKING MECH	L 4	-------	0.00	0.00	0.00	0.03	0.03	5	0.00
TARGET COMPUTER	L 5	-------	0.00	0.00	0.00	0.83	0.83	7	0.00
ASTEROID BLASTR	L 6	-------	0.00	0.00	0.00	7.25	7.25	18	0.00
PART COLLECTOR	L 7	-------	0.00	0.00	0.00	0.98	0.98	6	0.00
SENSOR	L 8	-------	0.00	0.00	0.00	19.37	19.37	32	0.01
CONTROL PANEL	L 9	-------	0.00	0.00	0.00	2.29	2.29	8	0.00
CELL MOB PHONE	L 10	-------	0.00	0.00	0.00	36.60	36.60	24	12.62
PULSED BEAM GEN	S 1	-------	0.00	0.00	0.00	0.00	0.00	5	0.00
AUTO DIAL MOD	S 2	-------	0.00	0.00	0.00	0.00	0.00	3	0.00
AUTO ANSWER SYS	S 3	-------	0.00	0.00	0.00	0.00	0.00	3	0.00
MUTE SWITCH	S 4	-------	0.00	0.00	0.00	0.00	0.00	8	0.00
MEMORY UNIT	SS 1	-------	0.00	0.00	0.00	0.00	0.00	0	0.00
REDIAL CIRCUIT	SS 2	-------	0.00	0.00	0.00	0.00	0.00	0	0.00
MESSAGE REC	SS 3	-------	0.00	0.00	0.00	0.00	0.00	0	0.00
CALL SCREEN DEV	SS 4	-------	0.00	0.00	0.00	0.00	0.00	0	0.00

⋮

Fig. 34—Detailed Parts Disposition, Option 15

is a must if requirements for spare components are to be computed. But for capability assessments in rigorous scenarios, its shortcomings soon become apparent. Suppose, for example, that MPC1 is required to generate 50 expeditions on a particular day, but that some combination of circumstances restricts its projected output to only 10 expeditions. Among these limiting circumstances is the pattern of component failures that would arise were all 50 expeditions actually achievable. It is unfair to impose such a drain upon MPC1's resources when it is inconceivable that so many failures could ever really occur. Furthermore, it is quite possible that if failures are generated on the basis of a more realistic number of expeditions, MPC1's projected output could rise above 10 expeditions.

Option 20 addresses this problem on an LRU-by-LRU basis at each time of analysis. In a separate computation for each LRU, option 20 iteratively reduces each base's original total sortie generation goal until, from the standpoint of that LRU alone, the adjusted goal can be achieved with some specified level of confidence. Upon the completion of this procedure for all LRUs, the user may be reasonably certain that excessive LRU failures no longer constitute a valid excuse for degraded performance.

Option 20 has two parameters. The sortie rate (on the SRTS record) is successively multiplied by the percent in the first parameter until the confidence level in the second parameter is achieved. In general, both the accuracy and execution time of the search increase as these parameters increase. In this example, unachievable sorties will be scaled back to 95 percent of the unachievable level on each step until a confidence of 0.80 is achieved.

Suppose MPC1 has a sortie rate of 50 on day 20, but the confidence of actually achieving this (with respect to the asteroid blaster) is only 0.25. Option 20 reduces the sortie rate to 95 percent of 50, or 47.5. If the confidence of achieving this new program (again, with respect to the blaster) is only 0.50, option 20 again reduces MPC1's sortie rate to 95 percent of its current value (95 percent of 47.5, or 45.125) until the confidence of achieving it equals or exceeds 0.80. This final value then becomes the actual "required" sortie rate on the basis of which blaster failures are computed. This procedure is repeated for each LRU at each time of analysis.

After completing its iterative search for the greatest number of sorties achievable at each base with respect to each LRU at each time of analysis, option 20 reports its findings to File 7 in PIPE.[10] For each LRU that cannot withstand the rigors of the user-specified program of activity, option 20 provides a new SRTS record group. Our example yields two such overtaxed LRUs, listed in Fig. 35 with their newly computed SRTS groups.

These records exhibit several noteworthy characteristics. The days on which sortie rates change at the various MPCs correspond exactly to those specified in the SRTS group (without time scaling). Moreover, user-specified rates act as upper bounds on the values computed by option 20. Thus, option 20 may reduce sortie rate goals, but it never increases them above their original levels.

```
OPTION 20 RESULTS:     4 LRUS HAD THEIR FLYING PROGRAMS SCALED BACK.
SEE OUTPUT FILE 7 FOR THE LRUS AND THEIR ASSOCIATED SORTIE RATES.

 FUEL CELL          L   3
       BASED ON INPUT AND PURCHASED STOCK FOR TIMES     2
MPC11.71    12.57  302.00  610.00   00.00   00.00   0

 CELL MOB PHONE     L  10
       BASED ON INPUT AND PURCHASED STOCK FOR TIMES     2
MPC11.63    12.44  302.00  610.00   00.00   00.00   0

 FUEL CELL          L   3
       BASED ON INPUT AND PURCHASED STOCK FOR TIMES    16    30    46    60
MPC11.71    12.57  302.00  610.00   00.00   00.00   0

 CELL MOB PHONE     L  10
       BASED ON INPUT AND PURCHASED STOCK FOR TIMES    16    30    46    60
MPC11.63    12.44  302.00  610.00   00.00   00.00   0
MPC20.00    52.85  102.56  251.50  610.00   00.00   0
MPC30.00    81.80  501.00  610.00   00.00   00.00   0
```

Fig. 35—Adjusted Flying Program, Option 20

[10] Option 20 has different file requirements than a regular Dyna-METRIC run; its files are described in App. A.

The "four" LRUs listed are really two LRUs with more than one associated SRTS group: a group for day 2 and another for days 16, 30, 46, and 60. These multiple listings are due to the time-varying stock levels. The first group gives the highest sortie rates achievable through day 2 (ignore all entries for times beyond day 2); the second group gives the highest sortie rates achievable through day 16, through day 30, and so forth. The number of sortie rates is equal to DMANALYS, the number of times of analysis set when the model was compiled (in this case, six).

The model-computed SRTS records are often incomplete in the sense that there are some MPCs for which a record is missing for some time of analysis. For example, the day 2 SRTS group for FUEL CELL contains only a record for MPC1. The highest sortie rate that can go into effect at MPC1 on day 1 and still be achievable (from the standpoint of the fuel cell) with a confidence level of 0.80 on day 2 is 2.57 expeditions per Z-455. The absence of records for MPC2 and MPC3 implies that the requirements established for those bases on day 2 are achievable without reduction (in both cases, this is automatic since operations begin after day 2).

Because sortie rate reductions are effected on an LRU-by-LRU basis, it is frequently the case that different rates are specified at the same time for the same base (compare the peacetime rates at MPC1). In part, this is a compromise between reality and Dyna-METRIC's somewhat conservative assumption that only FMC aircraft may fly sorties. With option 20 in effect, PMC aircraft may be thought to fly sorties as well, thereby accounting for any perceived variations in total flying hours among different LRUs. In short, option 20 generates LRU failures based upon both FMC and PMC sorties, although, as usual, only FMC sorties are counted when overall performance is measured.

The performance report (option 11) for day 30 (Fig. 36), compared with the earlier results without option 20 (Fig. 19), shows improvement in performance. Total worldwide sortie generation increases 2 percent under full cannibalization and 9 percent under partial cannibalization. The expected number of NFMC ships drops in both cases, while the number of FMC ships with 90 percent confidence rises (especially for MPC1). Note that PROB. ACHIEVE SORTIES continues to refer to the original, user-specified program of activity and not the model-computed sortie rates.

The problem LRUs list (option 8) also shows the effects of the relaxed method for computing component failures. While all three of the LRUs originally on the list continue to appear, their component impacts are diminished considerably. Although the presence of any entries at all might suggest that the model-computed program is still not achievable, such a conclusion is unwarranted. While these LRUs can be expected to incapacitate more than the target number of ships at one or more MPC, the new expedition generation rates with respect to each are sufficiently slack that these losses represent no degradation in required performance.

In addition to presenting a more realistic depiction of achievable sortie rates, the activation of option 20 prevents the occurrence of paradoxes of the sort observed earlier with the ATTR and ILM record groups.

Constrained Depot Workload (Option 1)

Option 1 constrains the depot's daily wartime workload to a percentage of its daily peacetime workload. Using this option without a peacetime flying program means that the depot completes no repairs during peacetime, implying a wartime limit of zero, which effectively cuts off depot repair.

PERFORMANCE BASED ON STOCK ON HAND ON DAY 30.

			FULL CANNIBALIZATION							PARTIAL CANNIBALIZATION							
BASE	TARG. NFMC	TOTAL ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC-90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC-90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	TOTAL BACK ORDERS
MPC1	3	20	0.543	0.889	13	3.878	0.194	39.56	2.258	0.162	0.717	12	5.461	0.273	38.98	2.187	24.77
MPC2	6	40	0.586	1.000	30	6.007	0.150	60.00	1.781	0.081	1.000	26	10.281	0.257	60.00	2.035	22.96
MPC3	3	20	0.099	0.959	14	4.504	0.225	39.87	2.372	0.064	0.937	14	4.733	0.237	39.85	2.284	141.86
TOTAL		80				14.389	0.180	139.43					20.475	0.256	138.82		189.60

POTENTIAL PROBLEM PARTS REPORT

WORLD-WIDE AIRCRAFT STATUS - DAY 30
(* - INDICATES AN APPLICATION FRACTION NOT EQUAL TO ONE.)

				WORLD-WIDE												
LRU NAME	NUMBER	WUC	TEQ	COMP. IMPACT	MINIMUM COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT	BASE	COMP. IMPACT	TARG. COMP. IMPACT
CELL MOB PHONE	L 10			10.75	10.18	12.00	MPC1	3.06	3.00	MPC2	5.64	6.00	MPC3	2.04	3.00
FUEL CELL	L 3			6.74	6.73	12.00	MPC1	1.86	3.00	MPC2	1.17	6.00	MPC3	3.71	3.00
SENSOR	L 8			2.80	1.42	12.00	MPC1	0.00	3.00	MPC2	0.00	6.00	MPC3	2.80	3.00

Fig. 36—Output from REPORT, Option 20

The first parameter contains the percentage of peacetime daily completions that can finish depot repair on each day of war; the second parameter is not used. In this example, DEPO can repair four times as many LRUs during wartime, so the first parameter is 400 (percent). Thus, if DEPO repaired four fuel cells per day during peacetime, it can repair at most sixteen per day during wartime.

Option 1 alone does not produce output reports. It must be run in conjunction with option 11 (for a performance report) and/or option 18 (for the computed maximum daily *throughput*, or number of completions).

The day 30 performance report (Fig. 37) shows that putting a cap on depot repair completions degrades performance. When there are no limits on depot repair (Fig. 19), about 19 spaceships are expected to be degraded (under the assumption of full cannibalization); but when depot repair throughput is limited to 400 percent that of peacetime, 26.5 spaceships are expected to be degraded.

Daily Demands Report (Option 18)

The expected daily demands for repair and supply at each location are reported when option 18 has been selected. This option has no parameters.

Demands for repair refers to the number of parts arriving at each location that require repair by that location's repair facility before NRTS/condemnation decisions are made. If such decisions are made before repair, the number of items actually entering repair is less than the reported daily demands for repair. (For example, if the NRTS rate is .3, the number entering repair is .7 times the daily demand rate for repair.)

Demands on supply refers to the expected number of requisitions received at that location on that day. At bases, the daily demands on base supply for LRUs is the number of daily removals of LRUs from aircraft. For cirfs and depots, the requisition arrives immediately when a base NRTSes (or condemns) a component, but the NRTSed component (and thus the demand for cirf or depot repair) does not arrive until a transportation time later. (Condemned components generate demands on supply at the higher echelon, but not demands for repair.)

Running options 1 and 18 with the second example produces the report shown in Fig. 38 (only the values associated with the first three LRUs are listed). Option 1 results are reported under the heading *Maximum Daily Thruput* for the depot; without option 1, depot throughput would be unlimited like that of the cirf.

Consider the status of the fuel cell on day 2. At MPC1, 5.4 fuel cells are expected to be removed and enter base level repair. Simultaneously, their replacements are requested from base supply. At MPC2 and MPC3, there are no expected removals because neither base is in operation yet. The expected number of requisitions received by CRF1 for the fuel cell is 5.4, while the depot expects no requisitions. The expected number of fuel cells that arrive at CRF1 for repair is 3.807, at the depot 0.819.

Depot Workload Report (Option 7)

Option 7 produces a report on minimum and maximum depot workload for each LRU, and purchases additional LRU depot stock to satisfy the requisitions that cannot be met by initial depot stock or by depot repairs. The confidence with which this is to be accomplished is specified in the second parameter (our example uses .80). The cost of each LRU is specified in the LRU record group.

PERFORMANCE BASED ON STOCK ON HAND ON DAY 30.

BASE	TARG. NFMC	TOTAL ACFT	FULL CANNIBALIZATION: PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	PARTIAL CANNIBALIZATION: PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	TOTAL BACK ORDERS
MPC1	3	20	0.023	0.303	8	8.304	0.415	33.53	2.594	0.000	0.004	5	12.345	0.617	22.96	2.990	50.37
MPC2	6	40	0.118	0.993	24	10.935	0.273	59.96	2.079	0.000	0.766	18	18.219	0.455	58.33	2.394	51.54
MPC3	3	20	0.000	0.188	11	7.270	0.364	37.81	2.439	0.000	0.061	9	8.509	0.425	34.35	2.818	286.09
TOTAL		80				26.510	0.331	131.30					39.072	0.488	115.63		388.00

Fig. 37—Performance Report, Option 1

DETAILED DEMAND RATE REPORT

PART NAME	NUMBER	LOC.	MAXIMUM DAILY THRUPUT	DEMAND RATE FOR	DAY		DAY		DAY		DAY		DAY	
NAV COMPUTER	L 1	MPC1	-------	REP./SUPP.	2	0.300	16	0.300	30	0.250	46	0.200	60	0.200
NAV COMPUTER	L 1	MPC2	-------	REP./SUPP.	2	0.000	16	0.540	30	0.300	46	0.300	60	0.300
NAV COMPUTER	L 1	MPC3	-------	REP./SUPP.	2	0.000	16	0.248	30	0.248	46	0.248	60	0.124
NAV COMPUTER	L 1	CRF1	-------	SUPPLY	2	0.300	16	0.840	30	0.550	46	0.500	60	0.500
NAV COMPUTER	L 1	CRF1	9999.99	REPAIR	2	0.212	16	0.840	30	0.582	46	0.500	60	0.500
NAV COMPUTER	L 1	DEPO	-------	SUPPLY	2	0.000	16	0.436	30	0.362	46	0.324	60	0.250
NAV COMPUTER	L 1	DEPO	0.22	REPAIR	2	0.035	16	0.390	30	0.383	46	0.324	60	0.259
MANEUVER THRUS	L 2	MPC1	-------	REP./SUPP.	2	0.378	16	0.378	30	0.315	46	0.252	60	0.252
MANEUVER THRUS	L 2	MPC2	-------	REP./SUPP.	2	0.000	16	0.680	30	0.378	46	0.378	60	0.378
MANEUVER THRUS	L 2	MPC3	-------	REP./SUPP.	2	0.000	16	0.252	30	0.252	46	0.252	60	0.126
MANEUVER THRUS	L 2	CRF1	-------	SUPPLY	2	0.378	16	1.058	30	0.693	46	0.630	60	0.630
MANEUVER THRUS	L 2	CRF1	9999.99	REPAIR	2	0.266	16	1.058	30	0.733	46	0.630	60	0.630
MANEUVER THRUS	L 2	DEPO	-------	SUPPLY	2	0.000	16	0.129	30	0.113	46	0.107	60	0.069
MANEUVER THRUS	L 2	DEPO	0.04	REPAIR	2	0.006	16	0.117	30	0.116	46	0.107	60	0.072
FUEL CELL	L 3	MPC1	-------	REP./SUPP.	2	5.400	16	5.400	30	4.500	46	3.600	60	3.600
FUEL CELL	L 3	MPC2	-------	REP./SUPP.	2	0.000	16	9.720	30	5.400	46	5.400	60	5.400
FUEL CELL	L 3	MPC3	-------	REP./SUPP.	2	0.000	16	4.000	30	4.000	46	4.000	60	2.000
FUEL CELL	L 3	CRF1	-------	SUPPLY	2	5.400	16	15.120	30	9.900	46	9.000	60	9.000
FUEL CELL	L 3	CRF1	9999.99	REPAIR	2	3.807	16	15.119	30	10.469	46	9.001	60	9.000
FUEL CELL	L 3	DEPO	-------	SUPPLY	2	0.000	16	9.795	30	8.961	46	7.859	60	6.113
FUEL CELL	L 3	DEPO	5.18	REPAIR	2	0.819	16	8.429	30	9.394	46	7.877	60	6.450

.
.
.

Fig. 38—Detailed Demand Rate Report, Options 1 and 18

Basically, the depot is required to ship an LRU to satisfy each requisition as received. During peacetime, LRUs are assumed to begin repair as soon as they arrive at the depot. During a wartime surge, the depot may not be able to satisfy requisitions by simply repairing LRUs. The requisitions arrive almost immediately, but the actual LRUs to be repaired may not arrive for several days. The depot must have stock reserves to last through the period when depot repairs cannot keep up with requisitions.

Option 7 considers only the initial stock levels (and any stock it purchases) and disregards deployed stock. ECHO prints a warning message if it encounters new stock levels. However, the model ignores these new values as if they were not in the STK record group and proceeds.

Program PIPE shows the value of purchased depot stock at each time of analysis (Fig. 39). The model determines the day on which cumulative requisitions most exceed initial depot stock plus the cumulative output from depot repair. Because requisitions must be satisfied as soon as they are received, sufficient depot LRUs are purchased to cover any shortfall with the given confidence. This stock is assumed to have been available from the start of the scenario.

The depot workload report covers only the days selected as times of analysis.[11] Figure 40 shows the portion of the report dealing with the first three LRUs. Each LRU initially had no depot stock. The model computed a requirement of three NAV COMPUTERs worth $60,000, four MANEUVER THRUSs worth $32,000 and three FUEL CELLs worth $8,000.

For NAV COMPUTER on day 30, there are .3621 expected requisitions and .3890 expected arrivals for repair. If, as in peacetime, the depot starts repairing each computer as soon as it arrives, .3552 would complete repair. To satisfy requisitions as soon as they are received, .3040 computers must start repair on day 30.

The right-hand part of the table gives the cumulative values from day 1 through each particular time of analysis. For example, *cumulative demands on supply* on day 30 (9.85 for NAV COMPUTER) is the sum of the daily demands from day 1 through day 30.

VALUE OF CIRF AND DEPOT LRUS AT END OF DAY --
(COSTS GIVEN IN THOUSANDS OF DOLLARS)

LOC.	DAY 2	DAY 16	DAY 30	DAY 46	DAY 60
CRF1	0.00	0.00	0.00	0.00	0.00
DEPO	6307.50	6307.50	6307.50	6307.50	6307.50
ALL	6307.50	6307.50	6307.50	6307.50	6307.50

VALUE OF ADDED CIRF AND DEPOT LRUS --
(COSTS GIVEN IN THOUSANDS OF DOLLARS)

LOC.	DAY 2	DAY 16	DAY 30	DAY 46	DAY 60
CRF1	0.00	0.00	0.00	0.00	0.00
DEPO	6307.50	0.00	0.00	0.00	0.00
ALL	6307.50	0.00	0.00	0.00	0.00

Fig. 39—Value of Stock Purchased, Option 7

[11]Option 7 has different file requirements than a regular Dyna-METRIC run. File 10 in PIPE contains the depot workload report associated with times of analysis; file 11 describes the same but for each internal (scaled) day. Appendix A explains the file requirements and job control language used with this option.

DEPOT WORKLOAD REPORT
(COSTS IN THOUSANDS OF DOLLARS)

PART NAME	NUMBER	LOC.	INIT. DEPOT STOCK	INC. DEPOT STOCK	COST INC. STOCK	DAY	DAILY DEMANDS ON SUPPLY	DAILY DEMANDS ON REPAIR	DAILY MAXIMUM REPAIR OUTPUT	DAILY MINIMUM INDUCTIONS	CUMULATIVE DEMANDS ON SUPPLY	CUMULATIVE DEMANDS ON REPAIR	CUMULATIVE MAXIMUM REPAIR OUTPUT	CUMULATIVE MINIMUM INDUCTIONS
NAV COMPUTER	L 1	DEPO	0	3	60.	2	0.0000	0.0405	0.0408	0.0328	0.00	0.08	0.08	0.07
						16	0.4361	0.3790	0.3135	0.3167	4.00	2.95	2.36	2.83
						30	0.3621	0.3890	0.3552	0.3040	9.85	8.83	7.54	7.08
						46	0.3239	0.3244	0.2912	0.3040	15.14	14.25	12.45	11.95
						60	0.2500	0.2611	0.2393	0.0250	19.01	18.38	16.21	15.21
MANEUVER THRUS	L 2	DEPO	0	4	32.	2	0.0000	0.0073	0.0034	0.0026	0.00	0.01	0.01	0.01
						16	0.1285	0.1144	0.0508	0.0340	1.22	0.84	0.37	0.34
						30	0.1128	0.1166	0.0527	0.0340	2.94	2.56	1.14	0.81
						46	0.1071	0.1072	0.0482	0.0340	4.67	4.31	1.93	1.36
						60	0.0693	0.0726	0.0328	0.0187	5.77	5.56	2.49	1.80
FUEL CELL	L 3	DEPO	0	3	8.	2	0.0000	0.9384	2.2778	0.9384	0.00	1.88	4.56	1.88
						16	9.7946	8.1493	12.3843	8.1493	83.95	61.36	90.30	61.36
						30	8.9611	9.5113	19.1086	9.5113	223.88	198.41	339.22	198.41
						46	7.8586	7.8841	15.8506	7.8840	353.30	331.17	611.36	331.17
						60	6.1126	6.5053	13.6497	6.5052	449.42	433.24	820.09	433.24
. . .														

Fig. 40--Depot Workload Report, Option 7

REQUIREMENTS COMPUTATIONS

When Dyna-METRIC is used to recommend spares levels, it is important to understand the order in which the computation proceeds. In particular, the model buys

- subSRUs before SRUs
- SRUs before LRUs
- depot stock before cirf stock
- cirf stock before base stock.

This way, the number of components in AWP status reflects the number of subcomponents purchased. Also, it reduces the number of components on order to reflect the spares purchased for the higher echelon. For example, the purchase of LRUs at the cirf and SRUs at the base will lower the requirement for LRUs at the base because fewer will be AWP or on order.

Input Data Set

Suppose that some junior managers doubted the company's decision to forgo purchasing additional spare components as a means of achieving greater levels of activity. They were sure that sensible levels of stock would augment performance. They went overboard, however, and decided to buy stock for all types of components at all locations.

Option Selection Records (OPT). Note the haphazard order of the options in Fig. 41. Any order will do because the computations proceed as previously mentioned.

Option 6 purchases SRU and subSRU spares. Subcomponents are typically less expensive than the parent LRU, so a stringent performance goal was selected—enough spares for each SRU and subSRU so that the probability of that part grounding no spaceships (assuming no stock of the parent LRU or SRU) is at least .90.

```
columns
         1         2         3         4         5         6         7
1234567890123456789012345678901234567890123456789012345678901234567890

OPT
         6  0   .90
         2 15   .90
         5 15   .90
         3 15   .90
         4 15   .90
        17      .85
         9
        12 15   .90
```

Fig. 41—OPT Record Group, Stock Buy

Options 2 and 5 purchase LRU spares for the depot and cirf, respectively. For each LRU, these options buy stock at DEPO and CRF1 to assure that no more than 15 percent of the fleet they serve is grounded (for that LRU) with at least a .90 probability.

Three options direct the model to purchase LRUs at each base. Option 3 works the same way as options 2 and 5—it buys stock to an individual aircraft goal, but for bases. Option 4 buys base LRU stock for an *overall* aircraft goal (explained below). Option 17 asks for sufficient base LRU spares so that, even without cannibalization, 85 percent of the expected number of spaceships are up and flying. Note that this second parameter has a different meaning than the other options' parameters: It is similar to the others' first parameter but reflects the percent of aircraft *not* grounded.

Option 4 is different from the other options because it buys base LRU spares for an *overall* aircraft goal. Though its parameters are the same as those for option 3, option 4 adds more base spares based on a marginal analysis that holds down total costs. Option 3 only ensures that the probability of a specific component grounding more than 15 percent of the spaceships is less than .10 (1 − .90). But if each of the ten LRUs has a .90 probability of not grounding too many Z-455s, the probability that *none* of those LRUs grounds too many ships is $.90^{10}$, or .35. Option 4, then, buys more spare base LRUs to guarantee (with .90 confidence) no more than 15 percent aircraft down.

Finally, option 9 saves the recommended stock levels (the sum of initial and purchased stock) in a formatted file for future use,[12] and option 12 reassesses capability. Option 12 produces essentially the same report as does option 11, but bases the performance on the recommended spares levels. (Option 11 bases its report on input stock as well as stock bought for an earlier time of analysis.)

Stock Levels (STK). The STK records associated with time-varying stock levels must be removed from the input data set. The model cannot purchase and deploy stock in the same run. It can, however, compute additional stock based on given initial stock levels, so it is not always necessary to remove every STK record (the initial stock level records may stay). In this example, the entire STK record group is omitted.

Output Reports

The ECHO program reports the selected options and their parameters. The PIPE program produces two reports about all but the base-level LRU spares. The REPORT program outputs the option 12 performance report associated with all spares and two reports about the base-level LRU spares.

The first PIPE report gives the costs of spares associated with options 2, 5, and 6, which is everything except LRUs at bases (Fig. 42). The top two tables show the value of original (day 0) and purchased spares at each time of analysis for LRUs at cirfs and depots, and SRUs and subSRUs at all locations. The bottom two tables show how much was spent on those spares to achieve the performance goal on each day of analysis. (The model does not necessarily achieve the performance goal on other days.) For example, LRU spares for DEPO were bought on days 16 and 30 to achieve the 85 percent available aircraft goal. The second report (Fig. 43) presents the spares levels in STK record format (same as File 3 of PIPE).

The first REPORT output, produced by option 12, is essentially the same as the performance report produced by option 11. The title is different, though, and indicates that the per-

[12]In the standard JCL, the base LRU stock levels are written to File 3 of REPORT, and all other stock levels are written to File 3 of PIPE. See App. A for details.

formance evaluation is based on the newly purchased stock levels. In particular, the report reflects all new stocks that were purchased for that time of analysis. Figure 44 shows that on day 30 performance is better than our goal of 15 percent or fewer aircraft NFMC with .90 confidence (e.g., 7.2 percent expected NFMC under the assumption of full cannibalization). This is because of the additional stock purchased to hold down the number of cannibalizations (option 17).

VALUE OF CIRF AND DEPOT LRUS AT END OF DAY --
(COSTS GIVEN IN THOUSANDS OF DOLLARS)

LOC.	DAY 2	DAY 16	DAY 30	DAY 46	DAY 60
CRF1	4500.00	21802.49	24052.49	24052.49	24052.49
DEPO	0.00	12000.00	27000.00	27000.00	27000.00
ALL	4500.00	33802.49	51052.49	51052.49	51052.49

VALUE OF SRUS AND SUBSRUS AT END OF DAY --
(COSTS GIVEN IN THOUSANDS OF DOLLARS)

LOC.	DAY 2	DAY 16	DAY 30	DAY 46	DAY 60
MPC1	0.00	0.00	0.00	0.00	0.00
MPC2	0.00	0.00	0.00	0.00	0.00
MPC3	0.00	4.00	4.00	4.00	4.00
CRF1	354.00	1567.00	1607.00	1607.00	1607.00
DEPO	1189.00	2221.00	3966.00	4016.00	4016.00
ALL	1543.00	3792.00	5577.00	5627.00	5627.00

VALUE OF ADDED CIRF AND DEPOT LRUS --
(COSTS GIVEN IN THOUSANDS OF DOLLARS)

LOC.	DAY 2	DAY 16	DAY 30	DAY 46	DAY 60
CRF1	4500.00	17302.49	2250.00	0.00	0.00
DEPO	0.00	12000.00	15000.00	0.00	0.00
ALL	4500.00	29302.49	17250.00	0.00	0.00

VALUE OF ADDED SRUS AND SUBSRUS --
(COSTS GIVEN IN THOUSANDS OF DOLLARS)

LOC.	DAY 2	DAY 16	DAY 30	DAY 46	DAY 60
MPC1	0.00	0.00	0.00	0.00	0.00
MPC2	0.00	0.00	0.00	0.00	0.00
MPC3	0.00	4.00	0.00	0.00	0.00
CRF1	354.00	1213.00	40.00	0.00	0.00
DEPO	1189.00	1032.00	1745.00	50.00	0.00
ALL	1543.00	2249.00	1785.00	50.00	0.00

FINAL STOCK LEVELS FOR SRUS AND SUBSRUS, AND FOR
CIRFS AND DEPOTS. BASE LRUS WILL BE WRITTEN BY REPORT.

Fig. 42—Value of Stock Purchased, Options 2, 5, and 6

NAV COMPUTER	DEPO	0				
NAV COMPUTER	CRF1	0				
MANEUVER THRUS	DEPO	0				
MANEUVER THRUS	CRF1	0				
FUEL CELL	DEPO	0				
FUEL CELL	CRF1	21				
DOCKING MECH	DEPO	0				
DOCKING MECH	CRF1	0				
TARGET COMPUTER	DEPO	0				
TARGET COMPUTER	CRF1	0				
ASTEROID BLASTR	DEPO	0				
ASTEROID BLASTR	CRF1	0				
PART COLLECTOR	DEPO	0				
PART COLLECTOR	CRF1	0				
SENSOR	DEPO	0				
SENSOR	CRF1	0				
CONTROL PANEL	DEPO	0				
CONTROL PANEL	CRF1	0				
CELL MOB PHONE	DEPO	36				
CELL MOB PHONE	CRF1	32				
PULSED BEAM GEN	DEPO	4				
PULSED BEAM GEN	CRF1	3				
PULSED BEAM GEN	MPC1	0	MPC2	0	MPC3	1
AUTO DIAL MOD	DEPO	13				
AUTO DIAL MOD	CRF1	5				
AUTO DIAL MOD	MPC1	0	MPC2	0	MPC3	0
AUTO ANSWER SYS	DEPO	9				
AUTO ANSWER SYS	CRF1	4				
AUTO ANSWER SYS	MPC1	0	MPC2	0	MPC3	0
MUTE SWITCH	DEPO	7				
MUTE SWITCH	CRF1	4				
MUTE SWITCH	MPC1	0	MPC2	0	MPC3	0
MEMORY UNIT	DEPO	11				
MEMORY UNIT	CRF1	3				
MEMORY UNIT	MPC1	0	MPC2	0	MPC3	0
REDIAL CIRCUIT	DEPO	4				
REDIAL CIRCUIT	CRF1	1				
REDIAL CIRCUIT	MPC1	0	MPC2	0	MPC3	0
MESSAGE REC	DEPO	2				
MESSAGE REC	CRF1	1				
MESSAGE REC	MPC1	0	MPC2	0	MPC3	0
CALL SCREEN DEV	DEPO	3				
CALL SCREEN DEV	CRF1	1				
CALL SCREEN DEV	MPC1	0	MPC2	0	MPC3	0

Fig. 43—Computed Spares Levels, Options 2, 5, and 6

PERFORMANCE BASED ON STOCK PURCHASED FOR 30 DAY SUPPORT

			FULL CANNIBALIZATION							PARTIAL CANNIBALIZATION							
BASE	TARG. NFMC	TOTAL ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	PROB. < 15% NFMC	PROB. ACHIEVE SORTIES	FMC- 90% CONF	E(NFMC)	EXP. % NFMC	E(SORTIES)	EXP. SORTIES /ACFT	TOTAL BACK ORDERS
MPC1	3	20	0.926	0.999	17	1.411	0.071	40.00	2.150	0.829	0.996	16	2.222	0.111	39.99	2.225	6.48
MPC2	6	40	0.996	1.000	35	2.974	0.074	60.00	1.622	0.734	1.000	32	5.250	0.131	60.00	1.733	5.94
MPC3	3	20	0.909	0.996	17	1.376	0.069	39.99	2.135	0.865	0.998	16	2.061	0.103	40.00	2.216	55.19
TOTAL		80				5.761	0.072	139.99					9.534	0.119	139.99		67.61

Fig. 44—Performance Report, Option 12

The second and third reports are just like those output by PIPE but for base-level LRU stocks. The value of the stocks and the final stock levels are reported in Fig. 45. Again, stock is not purchased for MPC2 and MPC3 until they start operations. As before, the stock levels (File 3 of REPORT) are in STK record format for use in a subsequent run.

MISCELLANEOUS OPTIONS

Four remaining options provide ways to manipulate input data for use with Dyna-METRIC. Each is described in detail in the input specifications (App. E) and file requirements writeup (App. A).

Briefly, option 10 initializes peacetime pipelines using previously saved or measured data. (If option 10 is not selected, the model initializes peacetime pipelines based on the peacetime flying program specified in the input data set.) Option 16 saves the pipelines at the end of the scenario for subsequent use with option 10. Option 21 allows the specification of LRU-dependent flying programs. Option 22 constructs a file of pipeline initialization data for use with option 10 (based on a file of user-supplied data).

VALUE OF BASE LRUS AT END OF DAY --
(COSTS GIVEN IN THOUSANDS OF DOLLARS)

BASE	DAY 2	DAY 16	DAY 30	DAY 46	DAY 60
MPC1	2304.50	3946.00	6946.00	6946.00	6946.00
MPC2	0.00	852.50	852.50	852.50	852.50
MPC3	0.00	8246.50	14447.00	16164.50	16164.50
ALL	2304.50	13045.00	22245.50	23963.00	23963.00

VALUE OF ADDED BASE LRUS --
(COSTS GIVEN IN THOUSANDS OF DOLLARS)

BASE	DAY 2	DAY 16	DAY 30	DAY 46	DAY 60
MPC1	2304.50	1641.50	3000.00	0.00	0.00
MPC2	0.00	852.50	0.00	0.00	0.00
MPC3	0.00	8246.50	6200.50	1717.50	0.00
ALL	2304.50	10740.50	9200.50	1717.50	0.00

NAV COMPUTER	MPC1	5	MPC2	3	MPC3	4
MANEUVER THRUS	MPC1	5	MPC2	3	MPC3	2
FUEL CELL	MPC1	54	MPC2	63	MPC3	75
DOCKING MECH	MPC1	1	MPC2	0	MPC3	0
TARGET COMPUTER	MPC1	3	MPC2	1	MPC3	1
ASTEROID BLASTR	MPC1	13	MPC2	12	MPC3	11
PART COLLECTOR	MPC1	4	MPC2	3	MPC3	2
SENSOR	MPC1	25	MPC2	27	MPC3	196
CONTROL PANEL	MPC1	6	MPC2	5	MPC3	5
CELL MOB PHONE	MPC1	8	MPC2	0	MPC3	17

Fig. 45--Base LRU Stock Reports, Options 3, 4, and 17

HELPFUL HINTS

Appendix F summarizes how to model cirfs, depots, and constrained repair, use VTM records, and specify missions and maintenance types. Appendix G is an alphabetical index to the data fields and the record groups where they are found.

V. MATHEMATICAL FOUNDATIONS

This section describes the mathematical foundations of Dyna-METRIC. The model computes expected pipeline contents and their probability distributions, then uses the probability distributions to predict performance and to recommend cost-effective stock levels.

COMPUTING PIPELINES

This discussion focuses on the integral that computes the expected contents of any pipeline segment. That integral is solved for the two repair-time probability distributions: exponential and deterministic. The discussion continues with descriptions of how LRUs enter the network of pipeline segments, how the model derives the expected contents of each pipeline segment, the effect of flying only achievable rather than requested sorties, and the effect of constrained repair on the pipeline computations.

General Pipeline Equations

The extended Palm's theorem (Hillestad and Carrillo, 1980; Crawford, 1981) is used to compute the expected contents of each pipeline segment. The theorem states that if parts arrive for service according to a nonhomogeneous Poisson process with intensity function *m(t)*, and if the service process is independent of the arrival process (for example, it does not take longer to service an LRU if there are many LRUs in the system), then the number of LRUs in the pipeline segment at time t has a Poisson distribution with mean

$$\lambda(t) = \int_{-\infty}^{t} m(s)P(s,t)\,ds \quad ,$$

where $P(s,t)$ is the probability that an LRU entering the system at time s has not completed service by time t.

In Dyna-METRIC, $m(s)$ is a step function. That is, there are a set of values, $M_0, M_1, M_2, M_3, \ldots$ such that $m(s)$ for s less than one equals M_0, $m(s)$ for s between one and two equals M_1, $m(s)$ for s between two and three equals M_2, and so on. The special structure of $m(s)$ helps reduce the extended Palm's integral to the general pipeline equations used throughout the model.

Suppose we want the expected pipeline contents at the end of day t. Day t is defined as time t through $t+1$, so the expected pipeline contents would be:

$$\lambda_t = \int_{-\infty}^{t+1} m(s)P(s,t+1)\,ds \quad .$$

Before proceeding, we need to know more about $P(s,t+1)$. Two special cases of this integral are considered: exponentially distributed service times and deterministic (or constant) service times. Let R be the expected service time. For deterministic service times,

$$P(s,t+1) = \begin{cases} 0 \text{ when } t+1-s \geq R \\ 1 \text{ when } t+1-s < R \end{cases} .$$

Thus

$$\lambda_t = \int_{t+1-R}^{t+1} m(s)\, ds \quad .$$

If $t+1-R < 1$, then

$$\lambda_t = M_0(R-t) + \sum_{n=1}^{t} M_n \quad .$$

Otherwise, there is some integer j with $j \leq t+1-R < j+1$, and the expected pipeline segment contents are given by

$$\lambda_t = M_j(j+R-t) + \sum_{n=j+1}^{t} M_n \quad .$$

Exponential service times are treated differently. At the start of day 1, the contents of the pipeline segment have a Poisson distribution with mean

$$\lambda_0 = \int_{-\infty}^{1} M_0\, P(s,1)\, ds = M_0 \int_{-\infty}^{1} e^{-(1-s)/R} ds = M_0\, R \quad .$$

Proceeding inductively, assume that the pipeline at the start of day t is Poisson with mean λ_{t-1}. We will derive λ_t as a function of λ_{t-1} and M_t and demonstrate that the pipeline at the end of day t (and start of day $t+1$) is Poisson with that mean.

Components in the pipeline segment at the end of day t were either there at the start of day t or arrived during the course of the day. What is p, the probability that a component in the pipeline at the start of day t is still in the pipeline at the end of day t? Recall that for the exponential distribution, the probability that the service time is greater than $T+s$, given that the service time is greater than T, is equal to the probability that the service time is greater than s. To compute p, the model does not need to know how long the components have been in service before day t. So $p = e^{-1/R}$.

If a component was in the pipeline segment at the start of day t, with probability p it is still in the pipeline segment at the end of day t. In fact, the number of components in the pipeline segment at the end of day t that were in the segment at the start of day t have a Poisson distribution with mean $p\lambda_{t-1}$.

What about components arriving during day t? The number of these still in the system at time $t+1$ are Poisson with mean $M_t R(1-e^{-1/R})$. Because summing two independent Poisson random variables results in a Poisson random variable with mean equal to the sum of the means, we have the number of components in the system at that end of day t Poisson distributed with mean

$$\lambda_t = \lambda_{t-1} e^{-1/R} + M_t R(1-e^{-1/R}) \quad .$$

Further, for both the exponential and the deterministic service times, the number of items completing service between times t and $t+1$ is Poisson with mean $\lambda_{t-1} + M_t - \lambda_t$.

To derive the equations describing expected pipeline segment contents and their probabilistic distributions, strong assumptions are made. In particular, it is assumed that demands for service are not correlated and that the service process does not change as a function of the number of demands. If demands arrive in clusters, or if the service process is sped up (or slowed down) as demands occur, the assumptions are violated and the probability distribution is no longer Poisson.

If real-world failure and repair processes were analyzed in detail, the assumptions would probably be violated. Most repair shops provide priority support to a component whose failure rate has increased substantially, thereby effectively reducing its service time. Moreover, failures are often correlated because a failure may trigger increased vigilance or even special inspections to detect emerging problems.

Fortunately, the Poisson distribution is quite robust; a substantial deviation from the assumptions is required to make the resultant distribution vary strongly from the assumed Poisson distribution. Thus, it has been accepted widely in numerous stockage computations models, including METRIC (Sherbrooke, 1968) and (implicitly) in the WRSK/BLSS computations.

For demand processes that sufficiently violate the assumptions, a distribution other than Poisson better approximates the contents of the total segment. Although Dyna-METRIC itself does not explicitly model failure processes other than the nonhomogeneous Poisson, the user may specify pipeline variance to mean ratios other than one (a VTMR of one corresponds to the Poisson distribution). In such cases, the model computes expected pipeline contents using the methodology described earlier, and then uses a negative binomial or binomial distribution with that mean to compute expected backorders and aircraft availability.

Why are the negative binomial and binomial distributions used to characterize pipeline quantities when the VTMR is greater than or less than one, respectively? From the start, observe that both of these distributions converge in distribution to the Poisson as their VTMRs approach one (Feller, 1968; Mood, Graybill and Boes, 1974). Thus, apart from any other justifications for using these distributions, they do have the nice property of letting the user do sensitivity analyses for VTMRs near one with distributions that are "almost Poisson."

At least two physical models give rise to negative binomial pipelines, the *parameter uncertainty model* and the *bulk arrival model*. The parameter uncertainty model has the following form. Suppose (in a steady-state world) that item demands are Poisson with parameter m, but m is unknown and randomly distributed according to a gamma law with known parameters. It has been shown that for deterministic repair times, the pipeline quantities are then negative binomial (Parzen, 1962). Furthermore, and of much greater importance, pipelines are negative binomial for general repair distributions so long as demands arrive according to the nonstationary Poisson process with mean value function $m(t) = m \times h(t)$ where $h(t)$ is nonstationary, but known, and m is constant, but unknown, with an uncertainty that is gamma distributed (Crawford, 1987). This models the situation where the flying program varies in a known manner, and the demand rate per flying hour is constant but unknown.

Given the flexibility of the gamma distribution, little modeling generality is lost by assuming m has that distribution. Also, the gamma distribution is the conjugate sampling distribution of the Poisson (DeGroot, 1970), which makes it particularly useful when the Poisson parameter is unknown and estimated from empirical data, as is done in standard Air Force data systems.

The bulk arrival model has the following form. Suppose events occur according to a non-stationary Poisson process with a known mean value function, and that at each event one or more demands occur with the number of demands being distributed logarithmically. Repair times can be arbitrarily distributed, but all items arriving in any particular bunch must have identical repair times (and hence leave repair simultaneously). From the dynamic Palm's theorem, the number of bunches in the system at any given time is Poisson distributed. From our assumptions, the number of items in each bunch is logarithmically distributed. Because a Poisson distribution compounded by a logarithmic distribution is negative binomial (Feller, 1968), the pipeline quantity has a negative binomial distribution. Deterministic repair times constitute a special case where all items (not just those in each bunch) share a common repair time.

If all items arriving in a given bunch do not have identical repair times, then, in general, the pipeline quantities will not have a negative binomial distribution. However, the pipeline distributions will still fall within the general class of compound Poisson distributions (Crawford, 1987).

As indicated above, either uncertainty about demand rates or bulk arrivals of demands can give rise to negative binomial pipelines. When a VTMR greater than one is empirically observed, the user must look at what is happening at the flight line and in the standard data systems to choose the appropriate physical interpretation of the situation; mathematically the choice is arbitrary.

Pipeline quantities that are binomially distributed can also be shown to arise in response to various physical situations. For example, consider a situation with deterministic repair times and where each sortie (or flying hour) gives rise to a broken part with probability p. Then the number of items in repair at time t has a binomial distribution with parameters S_t and p, where S_t is the number of sorties flown since one repair time before t, and p is as described above. The Poisson model is simply the limit of this model as the time periods become very short and the probabilities become very small (that is, as the VTMR goes to one, holding the expected number of failures constant).

LRU Removal Process

For each day, each base, and each LRU, the model computes the expected LRU failures. These are determined by the LRU demand rates and by each base's flying program. If the LRU demand rate is given in failures per flying hour, the expected number of daily demands is

Demand rate per flying hour of the LRU ×
Quantity per aircraft of the LRU ×
Application fraction of the LRU ×
Aircraft (after attrition) ×
Requested sorties per aircraft ×
Flying hours per sortie.

The number of aircraft at each base after attrition is determined in a straightforward manner. Let A_n be the number of aircraft on day n after attrition. Let B_n be the number of aircraft on day n if there were no attrition. By definition, $A_0 = B_0$ because there is no attrition in peacetime. A_n is computed inductively from A_{n-1} as follows:

$$A_n = B_n - B_{n-1} + A_{n-1} -$$

$$[A_{n-1} \times \text{ requested sorties per aircraft on day } n-1 \times$$

$$\text{Prob(an aircraft is lost during a sortie)}].$$

The computation of expected LRU demands and aircraft after attrition assume each base flies all of its requested sorties. When option 20 is selected, expected LRU demands are based on revised requested sortie rates, but attrition is still based on the original requested sortie rates.

Expected Pipeline Segment Contents

At this point, our discussion has shown how the model computes the expected contents of a pipeline segment, given the service process associated with the segment and the expected daily demands for that segment. Now we describe how the model derives the expected demands or arrivals for a pipeline segment based on the expected departures from a preceding pipeline segment. Included are details about the derivations of the expected number of AWP components (based on the pipelines and stock levels of the component's subcomponents) and the number of components ordered from a higher echelon but not received by a lower echelon (based on the pipelines and stock at the higher echelon). Also included is a discussion of how backorders at a higher echelon are allocated between bases and cirfs.

There are many steps in the pipeline computation. Table 2 summarizes those steps for the case where NRTS and condemnation decisions are made after repair. In the first step, LRUs arrive from the flight line for processing by the base administrative pipeline segment. The mathematics describing this segment are standard. In the second step, the LRUs that leave the base administrative pipeline segment enter the base repair pipeline segment, also described by the standard mathematics. In the last step, SRU backorders are used to complete the computation of the base LRU pipeline by determining the number of LRUs that are AWP. The expected base LRU pipelines are the foundation of most of the model's reports.

The table mentions five different types of processing: standard, depot repair, allocation, on-order, and AWP. Each will be dealt with in turn.

Standard Processing. Standard pipeline segments include the administrative and repair pipelines at the base, cirf, and depot for all parts, also the depot on-order pipeline segments. Each segment has as input the output of some earlier segment(s), with the exception of the base administrative pipeline for LRUs, which uses removals from the flight line. For instance, LRUs entering the base repair segment are those leaving the base administrative segment; SRUs entering cirf repair are those leaving the cirf administrative segment, plus those removed from LRUs leaving the cirf repair segment.

Let A_n be the expected demands on day n for the pipeline segment being analyzed. The expected daily contents of that segment are given by E_n (which is part of the total pipeline for that component), computed from the expected daily demands and service rate of the segment (administrative delay if this is an administrative segment, repair time if it is a repair segment, etc.) using the expected pipeline equations presented earlier. The expected departures from the segment are given by $D_n = E_{n-1} + A_n - E_n$, with $D_0 = A_0$.

Depot Repair Processing. This type of processing is used for all the depot repair segments. It differs from standard processing because of the optional daily cap on depot repairs. If there is no cap, then standard processing is used.

Let A_n be the expected demands on day n for depot repair for the part being analyzed. Let M be the maximum expected number of this part that can complete repair per wartime

Table 2

PIPELINE COMPUTATION STEPS

(NRTS/condemnation decisions and subcomponent removals occur after test)

Pipeline Segment	Input to Segment Output from Segment	Processing
LRU base admin	Flight line removals	standard
LRU base repair	LRU base admin	standard
LRU base-to-cirf	LRU base repair x NRTS	standard
LRU cirf admin	LRU base-to-cirf	standard
LRU cirf repair	LRU cirf admin	standard
LRU cirf-to-depot	LRU cirf repair x NRTS	standard
LRU depot admin	LRU cirf-to-depot	standard
LRU depot repair	LRU depot admin	depot repair
LRU depot on-order	(LRU depot repair + LRU cirf repair + LRU base repair) x condemn	standard
SRU base repair	LRU base repair x (1-(NRTS+condemn)) x replacement fraction	standard
SRU base-to-cirf	SRU base repair x NRTS	standard
SRU cirf admin	SRU base-to-cirf	standard
SRU cirf repair	SRU cirf admin + LRU cirf repair x (1-(NRTS+condemn)) x replace	standard
SRU cirf-to-depot	SRU cirf repair x NRTS	standard
SRU depot admin	SRU cirf-to-depot	standard
SRU depot repair	SRU depot admin + LRU depot repair x (1-(NRTS+condemn)) x replaces	depot repair
SRU depot on-order	(SRU depot repair + SRU cirf repair + SRU base repair) x condemn	standard
SubSRU base repair	SRU base repair x (1-(NRTS+condemn)) x replacement fraction	standard
SubSRU base-to-cirf	subSRU base repair x NRTS	standard
SubSRU cirf admin	subSRU base-to-cirf	standard
SubSRU cirf repair	subSRU cirf admin + SRU cirf repair x (1-(NRTS+condemn)) x replace	standard
SubSRU cirf-to-depot	subSRU cirf repair x NRTS	standard
SubSRU depot admin	subSRU cirf-to-depot	standard
SubSRU depot repair	subSRU depot admin + SRU depot repair x (1-(NRTS+condemn)) x replaces	depot repair
SubSRU depot on-order	(subSRU depot repair + subSRU cirf repair + subSRU base repair) x condemn	standard
Allocated subSRU depot backorders	subSRU depot pipeline and stock	allocation
SubSRU cirf on-order	allocated subSRU depot backorders	on-order
Allocated subSRU cirf backorders	subSRU cirf pipeline and stock	allocation
SubSRU base on-order	allocated subSRU cirf backorders	on-order
Allocated subSRU base backorders	subSRU base pipeline and stock	allocation
SRU depot AWP	allocated subSRU depot backorders	AWP
Allocated SRU depot backorders	SRU depot pipeline and stock	allocation
SRU cirf on-order	allocated SRU depot backorders	on-order
SRU cirf AWP	allocated subSRU cirf backorders	AWP
Allocated SRU cirf backorders	SRU cirf pipeline and stock	allocation
SRU base on-order	allocated SRU cirf backorders	on-order
SRU base AWP	allocated subSRU base backorders	AWP
Allocated SRU base backorders	SRU base pipeline and stock	allocation
LRU depot AWP	allocated SRU depot backorders	AWP
Allocated LRU depot backorders	LRU depot pipeline and stock	allocation
LRU cirf on-order	allocated LRU depot backorders	on-order
LRU cirf AWP	allocated SRU cirf backorders	AWP
Allocated LRU cirf backorders	LRU cirf pipeline and stock	allocation
LRU base on-order	allocated LRU cirf backorders	on-order
LRU base AWP	allocated SRU base backorders	AWP

day. (There is no cap during peacetime.) The model handles this cap as follows. First it computes F_n, the expected daily contents of this segment if there were no limit M. Second, it computes the expected daily departures from depot repair, d_n, again assuming there is no limit M, with $d_0 = A_0$, and $d_n = F_{n-1} + A_n - F_n$. Third, it computes D_n, the expected daily departures from depot repair, limited by M during wartime. Thus, $D_0 = A_0$ and $D_1 = \min(M, d_1)$.

For n larger than 1,

$$D_n = \min(M, \sum_{j=1}^{n} d_j - \sum_{j=1}^{n-1} D_j)$$

Finally, it computes the expected daily contents of the depot repair pipeline segment, $E_0 = F_0$ and $E_n = E_{n-1} + A_n - D_n$.

Allocation Processing. Standard processing and depot repair processing are sufficient for all the steps in Table 2 through the computation of the subSRU depot on-order segment. At this point, only the depot subSRU pipeline has been completely computed—the LRU and SRU pipelines still need their AWP segments, the base and cirf pipelines still need their on-order segments. With information about the depot subSRU pipeline and stock level, the depot subSRU backorders can be computed. Once those backorders have been *allocated* between AWP SRUs at the depot and requisitions for subSRUs from the cirfs, the model computes cirf subSRUs on-order and depot SRUs in AWP (using processes described below), thus completing the depot SRU pipeline and the cirf subSRU pipeline. The model works back in this manner from subSRUs to SRUs to LRUs, and from depots to cirfs to bases, using allocation, on-order, and AWP processing techniques, until the base LRU pipeline (the key pipeline) has been computed.

We will demonstrate allocation processing with subSRUs at the depot. Let E_n be the expected contents on day n of that portion of the subSRU's depot pipeline attributable to removals of the current SRU from the current LRU. This is just the sum of subSRUs in transit to the depot from cirfs, plus subSRUs in the depot administrative segment, plus subSRUs in the depot repair segment, plus subSRUs that have been ordered by the depot but not yet received.

Suppose the subSRU is indentured to more than one type of SRU, or that the current SRU is indentured to more than one type of LRU. Then the total subSRU depot pipeline may contain instances of the subSRU that are not attributable to removals of the current SRU from the current LRU. Let Q_L be the quantity per aircraft of the LRU being analyzed, Q_S be the quantity per LRU of the current SRU, and Q_B be the quantity per SRU of the current subSRU. Suppose Q_{BT} is the quantity per *aircraft* of the current subSRU. We approximate the total depot pipeline on day n for this subSRU by $ET_n = E_n Q_{BT}/(Q_B Q_S Q_L)$. For example, if only half the occurrences on the aircraft of the current subSRU are indentured to the current SRU indentured to the current LRU, the total pipeline is assumed to be twice as big as E_n.

If the depot lacks sufficient stock to cover the pipeline, there is a shortage somewhere in the system. Shortages of subSRUs may cause SRUs from which subSRUs have been removed to be delayed (AWP) pending receipt of a serviceable subSRU. Or a shortage of a subSRU may cause a delay in shipping a subSRU to a cirf that has submitted a requisition (cirfs submit requisitions whenever they NRTS or condemn a part; when a base condemns a part, it requisitions the cirf for a replacement, and the cirf in turn requisitions the depot).

Let BT_n be the total expected backorders of the subSRU. The equation for computing expected backorders is presented below in the discussion of pipeline distributions. Let B_n be the expected shortages attributable to the current SRU and current LRU. The expected backorders are allocated among the SRUs to which this subSRU is indentured, so that $B_n = BT_n Q_B Q_S Q_L / Q_{BT}$.

These allocated backorders, B_n, are in turn allocated to the cirfs and the depot. This is not done explicitly. Instead, we determine SR_n, the expected number of requisitions the depot satisfies on day n, and allocate those resources to satisfy the oldest unsatisfied requisitions. Shortages are thus translated into requisitions that have not been satisfied.

Let RC_{nc} be the expected number of requisitions for the subSRU received on day n by the depot from cirf c. This is the sum of the condemnations made on day n by cirf c's bases, plus the day n NRTS and condemnations that occurred at cirf c.

Let RD_n be the expected requirement for this subSRU generated by the repair of SRUs at the depot on day n. Let R_n be the total daily demands at the depot for serviceable subSRUs—the sum of RC_{nc} over all the cirfs, plus RD_n.

We are now in a position to compute SR_n, the number of requisitions the depot satisfies on day n. (Satisfying a requisition means only that the depot has shipped the component, not that the component has been received.) First, becuase of the steady-state nature of peacetime, $SR_0 = R_0$. For n greater than zero, $SR_n = B_{n-1} + R_n - B_n$.

Let SC_{nc} be the expected number of this subSRU that the depot ships to cirf c on day n. Let SD_n be the expected number of demands for subSRUs caused by subSRU removals from SRUs at the depot that are satisfied on day n. We set $SC_{0c} = RC_{0c}$ and $SD_0 = RD_0$. SR_n is divided between SC_{nc} and SD_n so that the oldest outstanding requisitions are satisfied first.

At the start of the conflict, there were B_0 unsatisfied requisitions. We assume these were allocated so that cirf c starts with $BC_{0c} = B_0 RC_{0c}/R_0$ unsatisfied requisitions and the depot starts with $U_{0i} = B_0 RD_0/R_0$, where i is the index of the subSRU. So for day n, the number of unsatisfied requisitions for cirf c is

$$BC_{nc} = BC_{n-1,c} + RC_{nc} - SC_{nc}$$

and for the depot is

$$U_{ni} = U_{n-1,i} + RD_n - SD_n \ .$$

U_{ni} is used below in the AWP processing example.

On-order Processing. The example of allocation processing computed the expected subSRU requisitions made by the cirf on the depot that are satisfied each day. Given that and the expected number of daily requisitions, we can compute the expected number of subSRUs ordered by each cirf that have not been received (they are either in transit or have not been sent). This example is typical of pipeline segments that are determined by on-order processing: From earlier steps, the expected number of shipments from the higher echelon was computed. That and the daily requistions imply the number of unreceived requisitions.

Consider in more detail the subSRU cirf on-order pipeline segment. RC_{nc} was defined above as the expected number of requisitions made on day n for this subSRU by cirf c. SC_{nc} was the expected number of this subSRU that the depot will ship to cirf c on day n, and BC_{nc} the number of unshipped requisitions. We want to compute E_{nc}, the expected number of subSRUs on day n that cirf c will have ordered but not received. E_{0c} is given by the expected number of subSRUs in transit plus the peacetime expected number of unsatisfied requisitions.

That is

$$E_{0c} = R_{0c} \times \text{(depot-to-cirf transportation time)} + BC_{0c} \ .$$

Standard processing determines the expected number of subSRUs in transit from the depot to the cirf (given the daily depot shipments), as well as D_{nc}, the expected number of subSRUs cirf c receives from the depot on day n. E_{nc} is then computed from $E_{nc} = E_{n-1,c} + RC_{nc} - D_{nc}$.

AWP Processing. The method for determining the number of AWP SRUs at the depot is typical of AWP processing, and is used as an example here. Let E_n be the expected number of SRUs in AWP at the depot on day n.

U_{ni}, computed in the example of allocation processing, is the expected unsatisfied demands for subSRU i at the depot on day n. We assume the shortfall of subSRU i for repairing SRUs at the depot has a Poisson distribution with mean U_{ni}.

There are two cases, depending on whether subSRUs are cannibalizable. Consider first the case where subSRUs are fully cannibalizable between identical SRUs awaiting parts. For each type of subSRU i, determine H_{ni}, the expected number of SRUs experiencing shortages of that subSRU. If the subSRU's quantity per SRU is 1, then $H_{ni} = U_{ni}$. If its quantity per SRU is Q_B,

$$H_{ni} = \sum_j j \ \text{Prob}[(j-1)Q_B < \text{number of shortages} \leq jQ_B] \ .$$

We then set $E_n = \max_i H_{ni}$.

In the second case, there is no cannibalization of subSRUs between identical SRUs awaiting parts. Temporarily assume all subSRUs have a quantity per SRU of 1. Further, assume the subSRUs are numbered so that subSRU 1 has the worst expected shortfall, subSRU 2 the second worst, and so on. Let E_{ni} be the expected AWP quantity if the SRU only had the first i subSRUs indentured to it. Let P_i be the probability that an SRU in service at the depot has at least one of those subSRUs removed. Then we approximate

$$\begin{aligned} P_0 &= 0; \\ P_{i+1} &= 1 - (1 - P_i) \times \text{Prob(subSRU } i{+}1 \text{ not removed from SRU)} \ ; \\ E_{n1} &= U_{n1} \ ; \text{ and} \\ E_{n,i+1} &= E_{ni} + (1 - P_i) \, U_{n,i+1} \ . \end{aligned}$$

If there are I types of subSRU, $E_n = E_{nI}$. For subSRUs with quantities per SRU greater than one, the model acts as if there were Q_B different types of identical subSRUs, each with an expected shortfall of U_{ni}/Q_B.

LOGISTICS SYSTEM CONSTRAINTS

This discussion covers the mathematics involved when the system is constrained. The first constraint deals with not being able to achieve the sorties asked for in the flying program. The second concerns constraints on repair resources.

Achievable Sorties for PMC and FMC Aircraft

Without option 20 selected, for purposes of computing expected daily component removals, Dyna-METRIC assumes that all requested sorties are flown. When the expected sorties cannot be flown, this assumption leads to an overestimation of expected pipeline contents and an underestimation in the performance report of FMC aircraft and sorties. With option 20 selected, *for each component* the model determines the number of PMC and FMC sorties that can be flown with that component operational (not flying sorties when the component is not operational) and computes the daily removals *of that component* based on those sorties rather than on the requested sorties.

This implies that aircraft continue to fly sorties and generate failures of, for instance, the radar, even if the engine is no longer available (although the engine itself is not flown once it is unavailable). Because the model does not know which components are essential and which are not, it leaves that determination up to the user. To facilitate this, the model writes out the achievable sorties for each component that could not achieve the requested sorties. The user should then skim through these sorties per aircraft records (they are formatted just like Dyna-METRIC input sortie level records, and may be substituted into the input data set as such) and determine the lowest sorties achievable record associated with a component whose failure would preclude even PMC sorties. That sorties per aircraft record should replace the sorties per aircraft record for that base in the Dyna-METRIC input data set, and the model should be rerun. We have used this approach rather than only flying FMC sorties because the latter approach would not include LRU removals from PMC sorties and hence would underestimate pipeline sizes.

Associated with option 20 are two parameters. The first, F, is the sortie reduction factor. The second, C, is the confidence level. The algorithm to compute the PMC and FMC sorties flown with a given component operational (and thus with a potential for being removed) proceeds as described below. For each component and each time of analysis, the following steps are performed.

Step 1: Compute the daily removals and daily pipeline contents based on the (current) requested flying program. On the first pass through this algorithm, this is the input requested flying program.

Step 2: If the requested sorties at base b can be achieved with k aircraft down, but not with $k+1$ down, determine p_b, the probability that there are $k \times QPA$ or fewer backorders for the current LRU at base b. (If the aircraft can still be flown with $q < QPA$ of this LRU working, then the allowable number of backorders is $k \times QPA + (AC - k) \times (QPA - q)$, where AC is the number of aircraft at base b, and p_b is the probability of having no more than $k \times QPA + (AC - k) \times (QPA - q)$ backorders.)

Step 3: For those bases with $p_b < C$, scale back the requested sorties as described below, and return to Step 1. If for all bases, $p_b \geq C$, proceed to the next time of analysis.

How are the sorties scaled back? Recall that the requested sorties are specified as a step function. That is, for days n_j through $n_{j+1} - 1$, the sortie rate is S_j. (Assume $n_0 = -\infty$.) Suppose the current day of analysis is N.

Case 1: N is the first day of analysis. Then, for all j such that n_j is not larger than N, S_j is reduced to $F \times S_j$. (F should be between zero and one.)

Case 2: There was a day of analysis N_1 before N. Then for all j such that n_j is larger than N_1 and $n_j \leq N$, and for the j such that N is between n_j and $n_{j+1} - 1$, S_j is reduced to $F \times S_j$.

Consider the following example:

S_0 is the sortie rate until day 3.
S_1 is the sortie rate from day 3 until day 5.
S_2 is the sortie rate from day 5 until day 20.

S_3 is the sortie rate from day 20 until the end of the conflict.

Days of analysis are 1, 6, 15, and 25.

If the expected backorders on day 1 are too high, S_0 is scaled back. If the expected backorders on day 6 are too high, S_1 and S_2 are scaled back. If the expected backorders on day 15 are too high, S_2 is scaled back. And if the expected backorders on day 25 are too high, S_3 is scaled back.

Constrained Priority Repair

The development of the expected pipeline equations assumed that there were either ample repair resources, so that components requiring repair could immediately commence the repair process, or that, for some initial period, there were no repair resources at all. The constrained repair submodel was designed for the analyses of cases between those two extremes, cases where there was some, but not ample, repair capacity. It was originally developed to analyze the queueing behavior of avionics components being tested on automated test equipment.

The submodel affects performance in two ways. First, it increases the expected pipeline contents of LRUs assigned to constrained priority repair by adding in the expected LRUs queued for repair. Second, it influences model dynamics by delaying the demands on higher echelons for repair and supply.

In standard queueing notation, this is a $D/D/K(t)$ system. That is, the arrival process for repair is assumed to be deterministic (not Poisson, as in the remainder of the model), repair times are deterministic (not exponential, even if exponential LRU repair times are selected in the main model), and there is a finite, time-dependent number of servers (also called test stands).

Notation

I – Number of types of LRU.

J – Number of types of repair resources.

T_i – Expected time required to test/repair LRU i.

D_{in} – Demands for repair of LRU i occurring on day n.

R_{ij} – Indicates whether LRU i can be tested on servers of type j. One means it can, zero means it cannot.

$K(j,n)$ – Number of servers of type j available on day n.

$A_j(k)$ – if k servers of type j are collocated, for each $A_j(k)$ days spent testing LRUs, $1-A_j(k)$ days must be used for self-test and testing other stands. (So on day n, at most $K(j,n)A_j[K(j,n)]$ time is available for testing LRUs on servers of type j.)

During peacetime, components are assumed to be scheduled for testing on a first come, first served basis. (Priority scheduling, discussed below, is used during wartime.) If there were only one type of LRU (LRU i) and one type of repair resource (type j), the expected number of LRUs in queue and in test during peacetime is given by the standard $M/M/K$ queueing equation (Gross and Harris, 1974),

$$D_{i0}T_i + \frac{\dfrac{(D_{i0}T_i)^{K(j,0)}D_{i0}/T_i}{(K(j,0)-1)!\,(K(j,0)/T_i - D_{i0})^2}}{\displaystyle\sum_{k=0}^{K(j,0)-1} \frac{D_{i0}^k T_i^k}{k!} + \frac{D_{i0}^{K(j,0)}\ T_i^{K(j,0)-1}}{(K(j,0)-1)!\,(K(j,0)/T_i - D_{i0})}} \quad . \tag{1}$$

To approximate queues in the more general case, the above equation is modified. The peacetime demand rate of LRUs that are repaired on multiple types of repair resources is allocated based on the number of applicable servers available. For example, if LRU i can be tested on types j_1 and j_2, and (during peacetime) there is one server of type j_1 and two stands of type j_2, the server of type j_1 is responsible for one-third of LRU i's peacetime demands while the servers of type j_2 are responsible for the remaining two-thirds. Let $D_{ij}(0)$ be the peacetime demand rate for LRU i that has been assigned to stand type j.

To compute the expected number of LRUs in queue or repair at server type j, the model also needs the total peacetime demand rate,

$$D_j(0) = \sum_{i=1}^{I} D_{ij}(0)$$

and the expected repair time for an LRU,

$$E_j[T] = \sum_{i=1}^{I} D_{ij}(0)T_i/D_j(0).$$

Substituting $D_j(0)$ for D_{i0} and $E_j[T]$ for T_i into Eq. (1) for expected LRUs in the system gives $Q_j(0)$, the expected number of LRUs in queue or test during peacetime for server type j. To determine the number of LRU i in queue or test at servers of type j, the model allocates $Q_j(0)$ based on demand rates. That is,

$$Q_{ij}(0) = Q_j(0)D_{ij}(0)/D_j(0) \quad .$$

Of the LRUs in queue or test, how does the model determine which were actually undergoing repair at the start of the conflict? Initially, the available servers are allocated among the LRUs they serve based on the number of LRUs of each type in the system, and on their test times. More specifically, N_{ij} is the number of servers of type j allocated to test LRU i, or

$$N_{ij} = \frac{Q_{ij}(0)T_i}{\displaystyle\sum_{x=1}^{I} Q_{xj}(0)T_x}$$

If $Q_{ij}(0)$ is less than or equal to N_{ij}, then all LRUs of type i in the system during peacetime that have been allocated to server type j are assumed to be in test. If $Q_{ij}(0) < N_{ij}$, then the unused allocation, $N_{ij} - Q_{ij}(0)$, is reallocated to the remaining LRUs. The algorithm proceeds as follows for each type of repair resource.

Step 1: For all LRUs such that $Q_{ij}(0) \leq N_{ij}$, let M_{ij} (the number of LRU i being tested on servers of type j at the start of the scenario) be set to $Q_{ij}(0)$. If all LRUs have been processed, the computation is done.

Step 2: Let I_M be the set of LRUs for which M_{ij} has not been defined. Redefine, for i in I_M,

$$N_{ij} = \frac{Q_{ij}(0)T_i}{\sum_{x \in I_M} Q_{xj}(0)T_x}\left[K(j,0) - \sum_{x \notin I_M} M_{xj}\right].$$

For all LRU i in I_M such that $Q_{ij}(0) < N_{ij}$, set M_{ij} to $Q_{ij}(0)$ and remove i from I_M. Repeat this step until no further deletions are made.

Step 3: For i in I_M, set

$$M_{ij} = \frac{Q_{ij}(0)T_i}{\sum_{x \in I_M} Q_{xj}(0)T_x}\left[K(j,0) - \sum_{x \notin I_M} M_{xj}\right].$$

The number of LRU i in queue (but not in test) at the start of the conflict is

$$Q_i(0) - \sum_{j=1}^{J} [Q_{ij}(0) - M_{ij}] .$$

If M_{ij} is the expected number of LRU i on servers of type j at the start of the conflict, how many LRUs of type i that were being tested during peacetime will complete repair between times n_1 and n_2, and how much test time will they use between those times? Answers to these questions give the effect on the wartime repair resource operations of the LRUs that were in test at the start of the scenario.

For LRU i, assume that the amount of time already spent in test at time 0 is uniformly distributed between 0 and T_i. We compute the expected time spent testing LRU i between times n_1 and n_2 by conditioning on T, the amount of repair time remaining before the LRU finishes testing. T is also uniformly distributed between 0 and T_i.

Let R be the time spent in repair between times n_1 and n_2. Further, assume $T_i \geq n_2 \geq n_1 > 0$. (If $n_2 > T_i$, substitute T_i for n_2 in the following.) Then

$$\begin{aligned} E(R) &= E(R \mid T > n_2)\,\mathrm{Prob}(T > n_2) + E(R \mid n_2 > T > n_1)\,\mathrm{Prob}(n_2 > T > n_1) + \\ &\quad E(R \mid n_1 > T)\,\mathrm{Prob}(n_1 > T) \\ &= M_{ij} \times [(n_2 - n_1)(T_i - ((n_2 + n_1)/2 - n_1)(n_2 - n_1)/T_i + n_1/T_i] \\ &= M_{ij} \times [1 - (n_1 + n_2)/2T_i]. \end{aligned}$$

Expected completed repairs between n_1 and n_2 is

$$M_{ij} \times \mathrm{Prob}(n_2 > T > n_1) = M_{ij}(n_2 - n_1)/T_i .$$

Two methods describe nonfunctioning repair resources. The first approach computes the amount of time servers spend in self-test and testing other servers, when they are not available to test any LRUs. The second approach describes PMC servers that may be able to test some types of LRU, but not others.

Time spent in self-test depends on the number of collocated servers. Typically, more servers imply less time per server needed for self-test. If k servers of type j are collocated, for each $A_j(k)$ spent testing LRUs, $1 - A_j(k)$ days are spent in self-test. This parameter is used to compute two different things, depending on context. The amount of time available on day n for testing LRUs is $K(j,n)\,A_j[K(j,n)]$. The amount of time a server is tied up because of a decision to test LRU i is $T_i/A_j[K(j,n)]$.

PMC servers are more complicated. The model assumes that a server is composed of many parts, some of which fail as the server is used. Failures that cannot be repaired immediately occur at the rate of λ_j per day spent in self-test or testing LRUs. After an expected (exponentially distributed) delay of μ_j, a replacement part is received. Resupply of replacement parts may not be available during a (user-specified) mid-scenario cutoff. The model also assumes servers are composed of parts whose individual failure rates are so low that it can, with very little error, assume there is at most one break of each type. This assumption allows the consolidation of all breaks onto one server of type j at each location.

The interesting question is, what is the probability at a given time that the server on which breaks are consolidated is able to test LRU type i? Let P_i be the probability that a failure of a server part does not render the server incapable of testing LRU i. Then the arrival rate of failures that disable the server for testing LRU i is $(1 - P_i)\lambda_j$ per day spent testing. The server is able to test LRU i if there are no unrepaired server failures that affect LRU i. The extended Palm's theorem describing pipeline quantities for LRUs can also be used to determine E(breaks), the expected number of unrepaired server failures affecting LRU i. Because this pipeline is Poisson, the probability that the server on which breaks are consolidated can test LRU i is the probability that no applicable breaks remain in the system, or $e^{-E(\text{breaks})}$.

If LRU i is being tested on the server where breaks are consolidated, and a break occurs that renders the server incapable of testing LRU i, the component is removed and set aside, and a component that can be tested is assigned to the server. The removed component waits for the server to be repaired before completing test, rather than being reassigned to another server.

The algorithm for scheduling components is fairly simple. Each day is analyzed in turn; all arrivals for the day are incorporated into the queue at the start of the day. (This is an optimistic assumption because it allows the maximum possible server utilization.) Let L_{in} be the number of LRU i that finish testing on day n. The number of LRUs in the system at the end of day n is

$$Q_i(0) + \sum_{j=1}^{J} M_{ij} + \sum_{t=1}^{n} (D_{it} - L_{it}) .$$

The three steps of the scheduling algorithm on day n are repeated while a queue or available test time remains, as follows.

Step 1: Determine which type of repair resource j has the most unassigned time.

Step 2: Determine which LRU i (with $R_{ij} = 1$) has the largest value for queue less stock.

Step 3: Determine which server k of type j has the most unassigned time. Determine n, the time when this LRU will complete test. Increment L_{in} and label the server as unavailable until time n.

If the number of servers decreases during the scenario, the model sets aside the LRUs that were in test on the servers being removed and does not resume their repair until their respective servers are returned.

We know of two major shortcomings with this approach. First, by assuming deterministic demands that arrive at the start of the day, and by using deterministic repair times, we have removed all variance from our queueing system. So long as demands for repair do not exceed repair capacity, queues do not build. As a result, expected queues are underestimated. Another problem involves the use of these queues once computed. Dyna-METRIC assumes that the pipelines of these LRUs are Poisson and independent. However, the use of a priority scheduling scheme has introduced dependence. Thus, when the distribution of NFMC aircraft is computed, the model is pessimistic and may ground too many aircraft.

PIPELINE DISTRIBUTIONS

Dyna-METRIC uses the pipeline variance to mean ratio to determine the pipeline distribution. A VTMR of one corresponds to a Poisson distribution, less than one to a binomial distribution, and greater than one to a negative binomial distribution. These distributions are used in turn to recommend stock levels for components, and to predict performance at each base for each time of analysis.

The following describes how pipeline distributions are derived from expected pipeline values and user-supplied VTMRs. It also shows how to compute expected backorders without having first to compute the pipeline distribution, saving computational resources in those cases where only expected backorder information is needed.

Notation

λ – Expected pipeline size of the component at the location being analyzed.

v – VTMR of the component's pipeline value.

p_j – Probability that there are j of the component in the pipeline.

$E(BO)$ – Expected backorders of the component.

S – Stock level of the component.

Poisson Distribution

The Poisson distribution is used for components with a pipeline VTMR of one. (Recall that SRUs and subSRUs have VTMRs of one, and that one is the default VTMR for LRUs.) In the case of the Poisson, we know that $p_j = e^{-\lambda}\lambda^j/j!$. For $j = 0$, $p_j = e^{-\lambda}$. For $j > 0$, $p_j = \lambda p_{j-1}/j$.

Binomial Distribution

The binomial distribution is used for LRUs with a pipeline VTMR of less than one. p_j for the binomial is

$$p_j = \frac{m!}{(m-j)!j!} r^j (1-r)^{m-j}$$

where $mr = \lambda$ and $(1-r) = v$. Solving for r and m yields $r = 1 - v$ and $m = \lambda/(1-v)$. The iterative definition, from which the model computes the terms of this distribution, is $p_0 = (1-r)^m$ and $p_j = r(m-j+1)p_{j-1}/[j(1-r)]$.

Negative Binomial Distribution

The negative binomial distribution is used for LRUs with a pipeline VTMR greater than one. p_j for the negative binomial is

$$p_j = \frac{(f + j - 1)!}{(f - 1)!j!} g^f (1 - g)^j$$

where $f(1 - g)/g = \lambda$ and $1/g = \upsilon$. Solving for f and g yields $g = 1/\upsilon$ and $f = \lambda/(\upsilon - 1)$.

For $j = 0, p_0 = g^f$. For $j > 0, p_j = p_{j-1}(f + j - 1)(1 - g)/j$. The terms of this distribution are computed with the iterative definition.

Expected Backorders

Expected backorders, or $E(BO)$, is defined to be $E[(\lambda - S)^+]$. Working from the definition, we have

$$E(BO) = \sum_{j=S}^{\infty} (j - S)p_j$$

$$= \sum_{j=0}^{\infty} jp_j - \sum_{j=0}^{\infty} Sp_j - \sum_{j=0}^{S-1} (j - S)p_j$$

$$= \lambda - S - \sum_{j=0}^{S-1} (j - S)p_j \ .$$

The primary advantage of using the final equality to compute expected backorders is that the number of terms of the pipeline distribution that are needed is not more than the stock level of the component, frequently a small value.

PROBABILITY DISTRIBUTIONS OF NFMC AIRCRAFT

Most information in the performance report for a given base is derived directly from the probability distribution of NFMC aircraft at that base. This discussion covers the computation of that probability distribution.

Notation

A – Number of aircraft at the base, after attrition.

I – Number of different types of LRU of which the aircraft are constructed.

Q_i – QPA of LRU i.

q_i – QPA of LRU i that must be operational if the aircraft is FMC. (By definition, q_i cannot be larger than Q_i.)

$p_i(j)$ – Probability that there are j of LRU i in the pipeline at the base. (See the previous discussion for how this probability is computed.)

S_i – [illegible] level of LRU i at the base.

b_{ij} – Probability that there are j backorders of LRU i at the base.

$P_{FC}(a)$ – Probability a or fewer aircraft are NFMC at this base, assuming a policy of full cannibalization.

$P_{PC}(a)$ – Probability a or fewer aircraft are NFMC at this base, assuming some types of LRU can be cannibalized and some types cannot.

Backorders

The probability that there are j backorders of LRU i is b_{ij}. If j is not zero, then for there to be j backorders, there would be $(j + S_i)$ of LRU i in the pipeline. If j is zero, there would be S_i or fewer of LRU i in the pipeline. That is:

$$b_{ij} = \begin{cases} \sum_{k=0}^{S_i} p_i(j), \text{ for } j = 0 \\ p_i\,(j + S_i), \text{ for } j > 0. \end{cases}$$

Full Cannibalization

$P_{FC}(a)$ is the probability that a or fewer aircraft are NFMC, assuming a policy of full cannibalization. Under such a policy, a or fewer aircraft are NFMC if and only if, for each LRU i, the number of aircraft held down by LRU i is no more than a. Let $P_i(a)$ be the probability that a or fewer aircraft are NFMC because of LRU i. Then

$$P_{FC}(a) = \prod_{i=1}^{I} P_i(a) \ .$$

But what is $P_i(a)$? For there to be at most a planes down because of LRU i, there must be at least $(A - a)$ aircraft with, for each aircraft, at least q_i of LRU i operational. That means that at least $(A - a)q_i$ of LRU i must be working; or, equivalently, at most $AQ_i - (A - a)q_i$ may be backordered. (When $Q_i = q_i$, this simplifies to aQ_i backorders.) So $P_i(a)$ is the probability of $AQ - (A - a)q_i$ or fewer backorders of LRU i, or

$$P_i(a) = \sum_{j=0}^{AQ_i - (A - a)q_i} b_{ij} \ .$$

Two major assumptions inherent to this approach should be emphasized. First, all the aircraft at a base look alike. This causes the treatment of LRUs with application fractions other than zero or one to be optimistic: All backorders can be consolidated onto the same aircraft. Consider the following situation: There are four aircraft at the base, made up of LRU A, LRU B, and LRU C. LRUs A and B have application fractions of .5, LRU C has an application fraction of 1. There are two backorders of LRU A, two of LRU B, and none of LRU C. Under our assumptions, the backorders of LRUs A and B would be consolidated onto the same aircraft, leaving two aircraft FMC. Suppose, however, that two of the aircraft were composed of LRUs A and C, and two of LRUs B and C. In this case, cannibalization should not help, and there should be no FMC aircraft at all. However, there could be one aircraft with LRUs A, B, and C, one with LRUs A and C, one with LRUs B and C, and one with LRU C only. In this case, there would be one aircraft FMC. Clearly, the application fraction alone does not define the aircraft structure with sufficient detail to determine legal cannibalizations. If consolidation of the backorders of low application fraction LRUs onto the same aircraft is leading to overly optimistic results, code the LRUs as not cannibalizable and use the performance measures based on a policy of partial cannibalization.

The second assumption is that cannibalization can be done instantly and without consuming resources. If this seems inappropriate, code the LRUs that are difficult or expensive to cannibalize as not cannibalizable, and again use the performance measures that are based on a policy of partial cannibalization.

Partial Cannibalization

$P_{PC}(a)$ is the probability that a or fewer aircraft are NFMC, assuming some types of LRU can be cannibalized and some types cannot. Assume, without loss of generality, that the first I_C types of LRU can be cannibalized, and the remaining $(I - I_C)$ types of LRU cannot. If $I_C = I$, then all LRUs are fully cannibalizable and $P_{PC}(a) = P_{FC}(a)$. So we will assume that I_C is less than I—that at least one type of LRU cannot be cannibalized.

Under this policy, a or fewer aircraft are NFMC if and only if, for each cannibalizable LRU i, a or fewer aircraft are down *and*, for the remaining $(I - I_C)$ noncannibalizable LRUs, considered together, a or fewer aircraft are down. This is because the backorders for cannibalizable LRUs can be consolidated onto the aircraft down for noncannibalizable LRUs. Let $P_i(a)$ be the probability that a or fewer aircraft are NFMC because of cannibalizable LRU i. Let $P_{NC}(a)$ be the probability that a or fewer aircraft are NFMC because of noncannibalizable LRUs. Then

$$P_{PC}(a) = P_{NC}(a) \prod_{i=1}^{I_C} P_i(a) \ .$$

when I_C is at least one (there is at least one type of cannibalizable LRU). $P_{PC}(a) = P_{NC}(a)$ when $I = 0$, the case where no LRUs can be cannibalized.

Next $P_{NC}(a)$ is derived. Let p_{NC} be the probability that a randomly chosen aircraft is down because of a backorder for a noncannibalizable LRU. The model then approximates the distribution of NFMC aircraft by using a binomial distribution. (The binomial distribution gives the probability of j successes, given J independent trials with p being the probability of success.) Although this approximation correctly captures the expected value, the variance may be underestimated in some cases. The probability that a or fewer aircraft are NFMC is thus approximated by

$$P_{NC}(a) = \sum_{j=0}^{a} \frac{A!}{(A - j)!j!} p_{NC}{}^{j}(1 - p_{NC})^{A-j} \ .$$

Let r_{NC} be the probability that a randomly chosen aircraft will not be grounded due to backorders of noncannibalizable LRUs. Clearly, $p_{NC} = 1 - r_{NC}$. Let r_i be the probability that a randomly chosen plane has no more than $(Q_i - q_i)$ backorders of noncannibalizable LRU i. Thus

$$r_{NC} = \prod_{i=I_C+1}^{I} r_i \ .$$

To determine r_i, the model conditions on the number of backorders of LRU i. Let $r_i(j)$ be the probability that a randomly chosen aircraft has at most $(Q_i - q_i)$ backorders of LRU i, given that there are j backorders of LRU i. Then

$$r_i = \sum_{j=0}^{AQ_i} b_{ij} r_i(j) \ .$$

Let $h_i(k,j)$ be the probability of k backorders of LRU i on the random plane, given j backorders of LRU i distributed across the A aircraft at the base. So

$$r_i(j) = \sum_{k=0}^{Q_i - q_i} h_i(k,j)$$

Finally, $h_i(k,j)$ is given directly by the hypergeometric distribution. (The hypergeometric gives the probability of k successes, given Q trials, when sampling without replacement from an urn containing AQ balls, of which j are "successes.") That is,

$$h_i(k,j) = \binom{j}{k} \binom{AQ_i - j}{Q_i - 1} \Big/ \binom{AQ_i}{Q_i}$$

At this point, all the intermediate quantities have been defined, so that one can work backward from $h_i(k,j)$ to $r_i(j)$ to r_i to r_{NC} to p_{NC} to $P_{NC}(a)$ and finally to $P_{PC}(a)$, the desired probability distribution of FMC aircraft under a policy of partial cannibalization.

COMPUTATION OF PERFORMANCE MEASURES

The performance measures in the Dyna-METRIC performance report are derived directly from the probability distributions of NFMC aircraft under full cannibalization and partial cannibalization policies.

Notation

A – Number of aircraft at the base, after attrition.

I – Number of different types of LRU of which the aircraft are constructed.

$E_i(BO)$ – Expected backorder quantity of LRU i (computed in the previous discussion of "Pipeline Distributions").

$P_{FC}(a)$ – Probability a or fewer aircraft are NFMC at this base, assuming a policy of full cannibalization (computed in the previous discussion of "Pipeline Distributions").

$P_{PC}(a)$ – Probability a or fewer aircraft are NFMC at this base, assuming some types of LRU can be cannibalized and some types cannot (computed in the previous discussion of "Pipeline Distributions").

Target NFMC

When selecting the performance report with either option 11 or 12, the user specifies a percentage p of the aircraft that is an acceptable NFMC percentage. The target NFMC level is simply $pA/100$, truncated to an integer. (If there are 10 aircraft with an acceptable NFMC percentage of 15 percent, the target NFMC level is one because $pA/100 = 15 \times 10/100 = 1.5$).

Total Aircraft

The total aircraft is just A, the total aircraft at the base after attrition.

Probability Less than *p* Percent NFMC

This number is actually the probability of at most p percent NFMC aircraft. Let a be the target NFMC. Then this probability is $P_{FC}(a)$ under full cannibalization and $P_{PC}(a)$ under partial cannibalization.

Probability Achieve Sorties

Suppose the requested sortie rate per aircraft is S_R, and the maximum achievable sortie rate per aircraft is S_T. The minimum number of FMC aircraft that can achieve the requested sortie program is a_M, the smallest integer such that $a_M S_T \geq AS_R$. To achieve the sorties, no more than $(A - a_M)$ aircraft can be NFMC. So the probability of achieving the requested sorties is $P_{FC}(A - a_M)$ for full cannibalization, and $P_{PC}(A - a_M)$ for partial cannibalization.

FMC with *C* Percent Confidence

When selecting the performance report with either option 11 or option 12, the user specifies a confidence level C at which FMC aircraft are to be reported. That is, a_M is to be found, where a_M is the largest integer such that the probability of a_M or more aircraft FMC is at least C. In the terms of our distributions, a_M is the largest integer such that $P_{FC}(A - a_M) > C$ for full cannibalization or $P_{PC}(A - a_M) > C$ for partial cannibalization.

Expected NFMC Aircraft

For a not equal to zero, the probability a aircraft are NFMC is $P_{FC}(a) - P_{FC}(a - 1)$ or $P_{PC}(a) - P_{PC}(a - 1)$. If $a = 0$, the probability of no NFMC aircraft is $P_{FC}(0)$ or $P_{PC}(0)$. Computing the expected value directly, for full cannibalization,

$$E_{FC}(NFMC) = \sum_{a=1}^{A} a\ [P_{FC}(a) - P_{FC}(a - 1)] \ .$$

Similarly, for partial cannibalization,

$$E_{PC}(NFMC) = \sum_{a=1}^{A} a\ [P_{PC}(a) - P_{PC}(a - 1)] \ .$$

Expected Percentage NFMC

This is expected NFMC aircraft expressed as a fraction of available aircraft. For full cannibalization, $E_{FC}(NFMC)/A$; for partial cannibalization, $E_{PC}(NFMC)/A$.

Expected Sorties and Expected Sorties per Aircraft

Let S_R be the requested sortie rate per aircraft and S_T be the maximum achievable sortie rate per aircraft. The minimum number of FMC aircraft that can achieve the requested sorties is a_M, where a_M is the smallest integer such that $a_M S_T \geq AS_R$. If there are fewer than a_M aircraft FMC, each aircraft flies S_T sorties. If there are $a \geq a_M$ aircraft FMC, each aircraft flies AS_R/a sorties.

Let $R_{FC}(a)$ be the probability that a aircraft are up. Then $R_{FC}(A) = P_{FC}(0)$ and $R_{FC}(a) = P_{FC}(A - a) - P_{FC}(A - a - 1)$. Computing expected sorties and expected sorties per aircraft by conditioning on available aircraft, for full cannibalization

$$E(\text{sorties}) = \sum_{a=1}^{a_M - 1} R_{FC}(a) a S_T + \sum_{a=a_M}^{A} R_{FC}(a) A S_R, \text{ and}$$

$$E(\text{sorties per aircraft}) = \sum_{a=1}^{a_M - 1} R_{FC}(a) S_T + \sum_{a=a_M}^{A} R_{FC}(a)\, AS_R/a \ .$$

The expected sorties never exceed and are frequently less than the expected FMC aircraft multiplied by the maximum sortie rate per aircraft. Expected sorties and expected sorties per aircraft for partial cannibalization are derived in a similar manner from $P_{PC}(a)$.

Total Backorders

This is the sum of the LRU expected backorder quantities,

$$\text{Total backorders} = \sum_{i=1}^{I} E_i(BO) \ .$$

The previous discussion on pipeline distributions shows how $E_i(BO)$ is computed.

Attrition Effects

Before proceeding to the problem LRUs report, we should mention a few assumptions about attrition. Details of how the model computes the effects of attrition were discussed earlier under "Expected Daily LRU Removals."

First, all aircraft measures reflect the number of aircraft available *after* attrition effects are taken into account. That is, if the goal is 10 percent NFMC, that means 10 percent of the aircraft remaining after attrition. Similarly, the expected number of NFMC aircraft is the expected number of aircraft down because of shortages of LRUs, and for no other reason. (In particular, the number of aircraft down from attrition is not included in this field.) The model at this point is effectively ignorant of the number of aircraft available before attrition.

Second, the sorties to be flown are the requested sorties per aircraft times the number of aircraft available after attrition. Thus, if the flying program calls for one sortie per aircraft on day 30, and the number of aircraft on day 30 with no attrition would have been 100, but after attrition is only 50, the base is evaluated against a goal of 50 sorties.

Problem LRUs Selection and Ordering

In addition to the overall assessment of capability produced by the option 11 (and option 12) performance report, the user may select option 8, which generates a list of problem LRUs along with a detailed report on their disposition throughout the theaters (on the shelf, in repair, in transit, AWP, etc.).

An LRU is selected for inclusion on this list if for some base there is too low a probability of at most a critical number of backorders. The allowable number of backorders is derived from the option 11 parameter that specifies the target NFMC percent. The allowable backorders level for an LRU at a base is just the acceptable number of NFMC aircraft at the base times the QPA of that LRU at that base. The probability target is the confidence level associated with option 11.

Problem LRUs are sorted according to the magnitude of their component impacts. The component impact for LRU i is an approximation to the number of aircraft that would be grounded if the aircraft were composed of only LRU i. In particular, the problem LRUs are sorted by

$$W_i = \sum_b \frac{\text{Expected backorders of LRU } i \text{ at base } b}{\text{QPA of LRU } i \text{ at base } b} \times [1/\text{Application fraction of LRU } i \text{ at base } b].$$

The application fraction is included in the above equation to scale up the apparent effect of LRUs with application fractions less than one. (However, the component impact is assumed to be zero at bases with application fractions of zero.) The reason the application fraction was included is best illustrated by an example. Suppose LRU A has an application fraction of one and expected backorders of four. Suppose LRU B has an application fraction of a half, and also has expected backorders of four. The QPA of each LRU is one. There are eight aircraft in all. Although both LRUs are grounding four aircraft, LRU B is considered worse because it is grounding *all* the aircraft it possibly can, while LRU A is only grounding half. The introduction of the application fraction into the above equation allows us to reflect such situations when ordering the problem LRUs list.

W_i is used only to order the list and is not written out in the problem LRUs report. The component impact quantities that are reported include:

- base level component impacts—the expected backorders divided by the QPA.
- worldwide component impact—the sum of the base level component impacts.
- base level target impact—the NFMC aircraft goal multiplied by the application fraction.
- worldwide target impact—the sum of the base level target impacts.
- worldwide minimum component impact—the expected backorders if all base pipelines and stocks were to be consolidated, divided by the average QPA. This quantity is used as an approximation to the worldwide component impact that would be attained if base stocks were optimally redistributed.

Problem SRUs and SubSRUs Selection and Ordering

If option 23 is selected in combination with option 8, a detailed disposition report, similar to that written for LRUs, is written for problem SRUs and subSRUs.

A problem LRU has problem SRUs if, for some location, the percent of the LRU's pipeline attributable to AWP exceeds the percentage specified in the second option 23 parameter.

For each location where the AWP value is too large, the worst SRUs are flagged and subsequently reported. By "worst" is meant those SRUs with the largest expected backorders over quantity per LRU. The number of SRUs selected per location (per LRU) is limited by the first option 23 parameter. These are the problem SRUs associated with that problem LRU, and their detailed disposition reports follow the detailed disposition report for the LRU.

Similarly, a problem SRU has problem subSRUs if, for some location, the percent of the SRU's pipeline attributable to AWP exceeds the specified percentage. For each location where the AWP value is too large, the subSRUs with the largest expected backorders over quantity per SRU are flagged as problems. Again, the number of subSRUs selected per location per SRU is limited by the first option 23 parameter. The detailed disposition reports for problem subSRUs follow the detailed disposition report of the associated SRU.

REQUIREMENTS MODE COMPUTATIONS

Dyna-METRIC's requirements mode adds sufficient stock to attain a user-specified performance goal. There are three different stockage algorithms:

- Backorder goal: Add sufficient stock to a location so that the probability of k or fewer backorders of the component is at least as great as a user-supplied confidence level (k is also derived from user-supplied data).
- Full cannibalization NFMC goal: Add sufficient LRU stock to a base so that the probability of N_D or fewer NFMC aircraft is at least as great as a user-supplied confidence level. Full cannibalization of LRUs is assumed.
- No cannibalization NFMC goal: Add sufficient LRU stock to a base so that the expected FMC level exceeds a user-supplied target level. No cannibalization of LRUs is assumed.

Backorder Goal

The problem is to recommend the stock level of a component such that the probability of k or fewer backorders is greater than C. Let p_j be the probability that there are j of the component in the pipeline at the location being analyzed.[1] Let P_{sj} be the probability that there are j or fewer backorders, given a stock level of s. Then

$$P_{sk} = \sum_{j=0}^{s+k} p_j \ .$$

Note that $P_{0k} = \sum_{j=0}^{k} p_j$,

and $P_{s+1,j} = P_{s,j} + p(s + 1 + j)$ for $s > 0$. The algorithm iteratively computes P_{sj} until $P_{sj} \geq C$. Let S_{new} equal the final value of s. Then the recommended stock level is the greater of S_{new} and the input stock level.

[1]These probabilities are computed under the discussion of "Pipeline Distributions."

Full Cannibalization NFMC Goal

The goal is to add sufficient stock so that the probability that A_D or fewer aircraft are NFMC, assuming full cannibalization of LRUs, is greater than C.

Notation

A – Number of aircraft at the base.

I – Number of types of LRU.

Q_i – QPA of LRU i.

q_i – QPA of LRU i that must be operational if the aircraft is not to be FMC.

p_{ij} – Probability that there are j of LRU i in the pipeline.

S_i – Stock level of LRU i. This may change during the operation of the algorithm.

c_i – Cost of another unit of LRU i.

Step 1: Determine the probability that A_D or fewer aircraft are NFMC, assuming full cannibalization. From the section on computing performance measures, this is

$$P_{FC}(A_D) = \prod_{i=1}^{I} \sum_{j=0}^{q_i A_D + (Q_i - q_i)A + S_i} p_{ij} \ .$$

If $P_{FC}(A_D) \geq C$, the current stock levels are adequate and the computation is done. Otherwise, proceed with steps 2 and 3.

Step 2: Let Δ_i be the amount $P_{FC}(A_D)$ would increase if S_i increased by one. Define

$$P_i(k) = \sum_{j=0}^{q_i A_D + (Q_i - q_i)A + k} p_{ij} \ .$$

Then $\Delta_i = P_{FC}(A_D)P_i(S_i + 1)/P_i(S_i) - P_{FC}(A_D)$.

Step 3: Determine which LRU i maximizes Δ_i/c_i. For this LRU, increase S_i by one and add Δ_i to $P_{FC}(A_D)$. Recompute the Δ_i for each LRU, and repeat this step until $P_{FC}(A_D) \geq C$.

No Cannibalization Availability Goal

The goal is to add sufficient stock so that the probability a random plane is FMC is greater than C. A policy of no cannibalization of LRUs is assumed.

Notation

A – Number of aircraft at the base.

I – Number of types of LRU.

Q_i – QPA of LRU i. (This algorithm assumes every LRU must be operational for the aircraft to be FMC.)

S_i – Stock level of LRU i. This may change during the operation of the algorithm.

p_{ij} – Probability that there are j of LRU i in the pipeline.

c_i – Cost of another unit of LRU i.

λ_i – Expected pipeline of LRU i.

$BO(i,s)$ – Expected backorders of LRU i, given a stock level of s.

From the subsection on pipeline distributions, we know that

$BO(i,s) = \lambda_i - s - \sum_{k=0}^{s-1} p_{ik}(k - s)$. And simple algebra gives

$$BO(i,s + 1) = BO(i,s) - 1 + \sum_{k=0}^{s} p_{ik} \ .$$

To approximate aircraft availability, assume that the number of backorders has a binomial distribution with parameter $BO(i,S_i)/AQ_i$. The probability that a random aircraft has no holes of LRU i is $(1 - BO(i,S_i)/AQ_i)$ if Q_i is one, and $(1 - BO(i,S_i)/AQ_i)^Qi$ if Q_i is greater than one. The probability that a random aircraft has no holes of any LRU is

$$B = \prod_{i=1}^{I} (1 - BO(i,S_i)/AQ_i)^{Q_i} \ .$$

If $B \geq C$, the current stock levels are adequate and the computation is done. If not, let Δ_i be the amount B would increase if S_i increased by one. Then

$$\Delta_i = B \times \left[\left(\frac{1 - BO(i,S_i + 1)/AQ_i}{1 - BO(i,S_i)/AQ_i} \right)^{Q_i} - 1 \right] .$$

Determine which LRU i maximizes Δ_i/c_i. For this LRU, increase S_i by one and add Δ_i to B. Recompute Δ_i for each LRU and repeat this step until $B \geq C$.

Stockage Options and Their Algorithms

How do the options selected by the user get translated into the above stockage algorithms? Associated with each option is an NFMC percentage goal, r, and a confidence level, C. Let Q be the quantity per aircraft of the LRU, and q be the number that must be operational if an aircraft is to be FMC.

Option 2: Depot LRU Stockage. For each LRU in turn, the algorithm to satisfy a backorder goal is invoked. Let A be the number of aircraft summed over all the bases. The goal is to add enough depot stock so that the probability of $(1 - r)A(Q - q) + rAQ$ or fewer depot backorders is greater than or equal to C.

Option 3: Base LRU Stockage. For each LRU in turn, the algorithm to satisfy a backorder goal is invoked. Let A_B be the number of aircraft at the base. Then the goal is to add enough base stock so that the probability of $(1 - r)A_B(Q - q) + rA_BQ$ or fewer base backorders is greater than or equal to C.

Option 4: Base LRU Stockage to Achieve Performance Goal. The algorithm to satisfy a full cannibalization NFMC goal is invoked. Let A_B be the number of aircraft at the base. The goal is to add sufficient stock so that the probability that rA_B or fewer aircraft are NFMC is greater than or equal to C.

Option 5: Cirf LRU Stockage. For each LRU in turn, the algorithm to satisfy a backorder goal is invoked. Let A_C be the number of aircraft at the bases served by the cirf.

Then the goal is to add enough cirf stock so that the probability of $(1 - r)A_c(Q - q) + rA_cQ$ or fewer cirf backorders is greater than or equal to C.

Option 6: SRU and SubSRU Stockage. For each SRU and subSRU in turn, the algorithm to satisfy a backorder goal is invoked. Let A be the number of aircraft at or served by the current location. Let Q_S be the quantity per aircraft of the SRU or subSRU. The goal is to add enough stock so that the probability of rAQ_S or fewer backorders is greater than or equal to C.

Option 7: No Cannibalization Base LRU Stockage. The algorithm to satisfy a no cannibalization availability goal is invoked. The goal is to add sufficient stock so that the probability a random plane is FMC is greater than C.

Order of the Stockage Computations. Each of steps 1 through 8 is performed for each time of analysis. That is, step 1 is done for each time, *then* step 2, etc.

1. Depot subSRU stock (if option 6 selected)
2. Cirf subSRU stock (if option 6 selected)
3. Base subSRU stock (if option 6 selected)
4. Depot SRU stock (if option 6 selected)
5. Cirf SRU stock (if option 6 selected)
6. Base SRU stock (if option 6 selected)
7. Depot LRU stock (if option 2 selected)
8. Cirf LRU stock (if option 5 selected)

Next, for each time of analysis, steps 9 through 11 are performed. That is, steps 9 through 11 are done for the first time of analysis, *then* for the second, etc.

9. Base LRU stock to satisfy a backorder goal (if option 3 selected).
10. Base LRU stock to satisfy a full cannibalization NFMC goal (if option 4 selected).
11. Base LRU stock to satisfy a no cannibalization availability goal (if option 17 selected).

DEPOT WORKLOAD

The depot workload report is optional. Its purpose is to recommend depot stock levels and induction schedules for LRUs so that the depot can satisfy all requisitions as they occur (meaning simply that the depot has LRUs available to ship to the base). The computations in general involve expected values and ignore the probabilistic aspects considered throughout most of Dyna-METRIC.

Computation of Maximum Requisitions

Requiring the depot to satisfy all requisitions immediately is a fairly stringent goal. In particular, it does not take into account those cases where th base may have ample stock and not really need the additional stock it automatically requisitions from the depot. This goal also ignores cases where forward transportation to the bases has been cut off: the depot is still required to satisfy the requisitions, even if there is no way to transport the part to the base. Thus, maximum required issues refer to the daily demands on depot supply caused by NRTS or condemnation decisions made by bases and cirfs.

Let N_C be the NRTS rate at cirfs, C_B the condemnation rate at bases, and C_C the condemnation rate at cirfs. Let $d_j(n)$ be the expected number of LRUs completing repair (if NRTS and condemnation decisions occur after attempting repair) or completing administrative delay (if NRTS and condemnation decisions are made before attempting repair) at location j (a base or cirf). Let J_B be the set of bases and J_C be the set of cirfs. The number of requisitions received by the depot on day n is

$$d_S(n) = \sum_{j \in J_B} C_B d_j(n) + \sum_{j \in J_C} (N_C + C_C) d_j(n) .$$

Maximum Inductions and Production

To meet the required requisitions, the depot must either buy stock or repair LRUs. The preference is to meet a requisition by repairing an LRU. The general algorithm, then, assumes that the depot repairs everything it can. Given the resulting levels of repaired LRUs, the model determines the shortfall that must be covered with new stock. Given the new stock, it reexamines the vector of daily inductions and deletes any unnecessary repairs from the schedule.

The first step is to determine how many inductions can be made on each day of the scenario. This is simply the arrivals for depot repair, discussed earlier in this section. Let $i_{max}(n)$ be the maximum number of inductions on day n. Let $r_{max}(n)$ be the maximum number of LRUs successfully completing depot repair on day n, as determined by the general pipeline equations.

Computation of Needed Depot Shelf Stock

Suppose at the start of the scenario the depot pipeline is λ_0 and the depot stock level is S. The total number of requisitions to be satisfied by day n is

$$D_S(n) = \lambda_0 + \sum_{j=1}^{n} d_S(j)$$

The number of requisitions that can be met through day n is the sum of the stock S plus the total successful repairs,

$$R(n) = \sum_{j=1}^{n} r_{max}(j) .$$

Dyna-METRIC determines ΔS, the incremental stock level needed to satisfy requisitions. If no confidence level was specified, ΔS is determined by

$$\Delta S = \max\,[0, \lambda_0 - S, \max_n\,(D_S(n) - R(n) - S)] .$$

If the user specifies a confidence level, the model first determines the maximum $[\lambda_0, \max D_S(n) - R(n)]$, and assumes that the shortfall in demands satisfied by repaired LRUs has that mean. (This is not too bad an assumption if the LRU has a VTMR of one and the expected depot repair times exceed a day.) Next the model determines S_D such that the probability that the shortfall is less than S_D exceeds the given confidence level. Finally, ΔS is taken to be the larger of the ΔS from the first approach and $S_D - S$.

Minimum Inductions and Workload

Given the stock purchased in the last step, we now have $D_S(n) - S - \Delta S \leq R(n)$. At any given point in time, at least $D_S(n) - S - \Delta S$ LRUs must have completed repair if the depot is to meet all requisitions. There may be several repair schedules F such that $F(n) \geq D_S(n) - S - \Delta S$ and $F(n) \leq R(n)$. (Also required is $F(n) \geq F(n-1)$ so that LRUs cannot be "unrepaired.") $F(n) = D_S(n) - S - \Delta S$ and $F(n) = R(n)$ are two possibilities. But what is the best F?

We define as "best" the F that induces the shortest peak repair load (call the winner F_{best}). That is, F_{best} is the F that minimizes

$$\max_n \left(F(n) - F(n-1)\right) .$$

To understand this solution, it is easiest to draw a picture (see Fig. 46). The algorithm to compute F_{best} produces the "taut string solution." Visualize R and $D_S - S - \Delta S$ as two tall walls on a flat surface. Let F be a string affixed at the origin. Then any string F that lies between the two walls is feasible. A "taut string" F_{best} can be determined by standing at $D_S - S - \Delta S$ at the maximum time of analysis and pulling the string taut.

To determine minimum inductions, the model assumes constant repair times. For a part to leave at time n, it must have been inducted a repair time earlier.

The depot workload report includes cumulative and daily values for requisitions, maximum possible inductions, maximum possible production, and the inductions associated with the best minimum repair schedule.

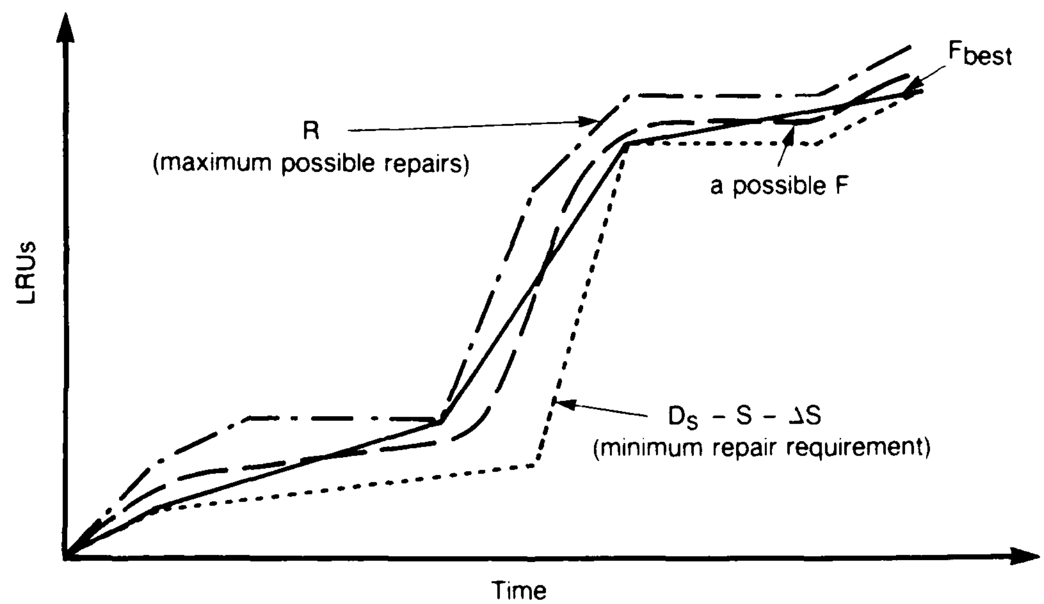

Fig. 46—Minimum Peak Load (Taut String) Solution

Appendix A

FILE REQUIREMENTS AND JOB CONTROL LANGUAGE

The following provides information about the files and job control language (JCL) used to run Dyna-METRIC's five programs. It is tailored to the IBM system we use and must be adapted for other computer systems. Some systems (such as our Vax) require that the files be defined in the source code with OPEN statements (see App. B for details).

File requirements depend on the options selected in the OPT record group. Formats, contents, and logical record lengths (LRECL) for standard option runs are described first, followed by a sample job in JCL. Runs with nonstandard options (7, 10, 16, and 20-22) require modifications to the JCL that are explained next.

STANDARD OPTIONS

1. PART Partition the input data set according to the compiled dimensions. Skip this step if there are fewer than DMLRUGRP LRUs, fewer than DMSRUGRP SRUs, and fewer than DMSUBGRP subSRUs.
   - File 4 - Formatted, input. The Version 4 data set to be partitioned.
   - File 5 - Not used but some systems may require that it be defined.
   - File 6 - Formatted, output. Error messages.
   - File 8 - Formatted, output. The partitioned data set to be used as File 5 in subsequent job steps. LRECL is the greater of 88 and (DMCHANGE×9)+4.

2. ECHO Check the partitioned data set for errors and echo the inputs. Options 13 and 14 supress various parts of the echo.
   - File 5 - Formatted, input. The partitioned data set. (PART's File 8).
   - File 6 - Formatted, output. Error messages and echo of administrative and scenario data. LRECL=133.
   - File 1 - Formatted, output. Table 1 of component data. LRECL=133.
   - File 2 - Formatted, output. Table 2 of component data. LRECL=133.
   - File 3 - Formatted, output. Table 3 of LRU data. LRECL=133.
   - File 4 - Formatted, output. Table 4, application fractions. LRECL=133.

File 7 - Formatted, output. Table 5, quantities per application. LRECL=80.

File 8 - Formatted, output. Table 6, stock levels. LRECL=133.

File 9 - Formatted, output. Constrained repair data. LRECL=133.

3. PIPE Perform the pipeline computations.

File 5 - Formatted, input. The partitioned data set. (PART's File 8).

File 6 - Formatted, output. Selected reports.

File 1 - Formatted, internal. LRECL=80.

File 2 - Unformatted, output. Used as File 2 in REPORT.

File 3 - Formatted, output. Associated with option 9. Stock levels for LRUS at depots and cirfs, and SRUs and subSRUs at all locations. (Base level LRU stock is output in REPORT). Stock levels can be saved on disk rather than routed to the printer. LRECL=80.

File 4 - Unformatted, internal.

File 8 - Unformatted, output. Used as File 2 in MOD.

4. MOD Compress the pipeline file.

File 5 - Formatted, input. The partitioned data set. (PART's File 8).

File 6 - Formatted, output. Error messages.

File 2 - Unformatted, input. PIPE's File 8.

File 8 - Unformatted, output. Used as File 9 in REPORT.

5. REPORT Write reports.

File 5 - Formatted, input. The partitioned data set. (PART's File 8).

File 6 - Formatted, output. Selected reports.

File 2 - Unformatted, input. PIPE's File 2.

File 3 - Formatted, output. Associated with option 9. Base LRU stock levels. (Other stock reports are output in PIPE). Stock levels can be saved on disk rather than routed to the printer. LRECL=80.

File 4 - Unformatted, internal.

File 9 - Unformatted, input. MOD's File 8.

STANDARD JCL EXAMPLE

The following JCL is used to run Dyna-METRIC on the Rand IBM 3032 with standard options. In this example, stock levels (requested with option 9) are printed rather than saved on disk.

```
//T0114RUN JOB (4031,500,1,7),'DYNA-METRIC',CLASS=N,COND=((8,LT))
/**
/**              STEP 1 -- PART
/**
/**     V4PART is the run module in the library PARTLIB.  It reads
/**     UNPART.DATA, the unpartitioned input data set, and creates
/**     PART.DATA, the partitioned data set used in the rest of the program.
/**
//PART  EXEC PGM=V4PART
//STEPLIB  DD DSN=T.T0114.A4031.PARTLIB,DISP=SHR
//FT04F001 DD DSN=T.T0114.A4031.UNPART.DATA,DISP=SHR
//FT05F001 DD DUMMY
//FT06F001 DD SYSOUT=A
//FT08F001 DD DSN=T.T0114.A4031.PART.DATA,DISP=(NEW,CATLG),
//            UNIT=USER,VOL=SER=USER50,SPACE=(TRK,(10,1),RLSE),
//            DCB=(RECFM=FB,LRECL=100,BLKSIZE=3600)
/**
/**              STEP 2 -- ECHO
/**
/**     V4ECHO is the run module in the library ECHOLIB.  It checks
/**     the partitioned data set for errors and echoes the inputs.
/**
//ECHO  EXEC PGM=V4ECHO
//STEPLIB  DD DSN=T.T0114.A4031.ECHOLIB,DISP=SHR
//FT05F001 DD DSN=T.T0114.A4031.PART.DATA,DISP=(OLD,PASS)
//FT06F001 DD SYSOUT=A
//FT01F001 DD SYSOUT=A
//FT02F001 DD SYSOUT=A
//FT03F001 DD SYSOUT=A
//FT04F001 DD SYSOUT=A
//FT07F001 DD SYSOUT=A
//FT08F001 DD SYSOUT=A
//FT09F001 DD SYSOUT=A
/**
/**              STEP 3 -- PIPE
/**
/**     File 2 is passed to REPORT, and File 8 is passed to MOD.
/**     File 3, which contains all stock levels except those of base LRUs,
/**     is routed to the printer.  Use this JCL if stock options are
/**     not selected.
/**
//PIPE  EXEC PGM=V4PIPE,REGION=990K
//STEPLIB  DD DSN=T.T0114.A4031.PIPELIB,DISP=SHR
//FT05F001 DD DSN=T.T0114.A4031.PART.DATA,DISP=(OLD,PASS)
```

```
//FT06F001 DD SYSOUT=A
//FT01F001 DD DSN=T.T0114.A4031.TMP1,DISP=(NEW,DELETE),
//             DCB=(RECFM=FB,LRECL=80,BLKSIZE=3600,BUFNO=1),
//             UNIT=TEMP,VOL=SER=TEMP10,SPACE=(TRK,(10,1),RLSE)
//FT02F001 DD DSN=T.T0114.A4031.TMP2,DISP=(NEW,CATLG),
//             DCB=(RECFM=VBS,LRECL=32756,BLKSIZE=6500,BUFNO=1),
//             UNIT=TEMP,VOL=SER=TEMP10,SPACE=(TRK,(200,10),RLSE)
//FT03F001 DD SYSOUT=A
//FT04F001 DD DSN=T.T0114.A4031.TMPP4,DISP=(NEW,DELETE),
//             DCB=(RECFM=VBS,LRECL=32756,BLKSIZE=6500,BUFNO=1),
//             UNIT=TEMP,VOL=SER=TEMP10,SPACE=(TRK,(200,10),RLSE)
//FT08F001 DD DSN=T.T0114.A4031.TMP8,DISP=(NEW,PASS),
//             DCB=(RECFM=VBS,LRECL=32756,BLKSIZE=6500,BUFNO=1),
//             UNIT=TEMP,VOL=SER=TEMP10,SPACE=(TRK,(200,10),RLSE)
/**
/**              STEP 4 -- MOD
/**
/**    File 2 was passed from PIPE.  File 8 is passed to REPORT.
/**
//MOD    EXEC PGM=V4MOD,REGION=990K
//STEPLIB  DD DSN=T.T0114.A4031.MODLIB,DISP=SHR
//FT05F001 DD DSN=T.T0114.A4031.PART.DATA,DISP=SHR
//FT06F001 DD SYSOUT=A
//FT02F001 DD DSN=T.T0114.A4031.TMP8,DISP=SHR
//FT08F001 DD DSN=T.T0114.A4031.TMP9,DISP=(NEW,PASS),
//             DCB=(RECFM=VBS,LRECL=32756,BLKSIZE=6500,BUFNO=1),
//             UNIT=TEMP,VOL=SER=TEMP10,SPACE=(TRK,(200,10),RLSE)
/**
/**              STEP 5 -- REPORT
/**
/**    File 2 was passed from PIPE, and File 9 was passed from MOD.
/**    File 3, which contains base LRU stock levels, is routed to the
/**    printer.  Use this JCL if stock options are not selected.
/**
//REPORT   EXEC PGM=V4REP,REGION=1990K
//STEPLIB  DD DSN=T.T0114.A4031.REPLIB,DISP=SHR
//FT05F001 DD DSN=T.T0114.A4031.PART.DATA,DISP=SHR
//FT06F001 DD SYSOUT=A
//FT02F001 DD DSN=T.T0114.A4031.TMP2,DISP=(OLD,DELETE)
//FT03F001 DD SYSOUT=A
//FT04F001 DD DSN=T.T0114.A4031.TMPR4,DISP=(NEW,DELETE),
//             DCB=(RECFM=VBS,LRECL=32756,BLKSIZE=6500,BUFNO=1),
//             UNIT=TEMP,VOL=SER=TEMP10,SPACE=(TRK,(200,10),RLSE)
//FT09F001 DD DSN=T.T0114.A4031.TMP9,DISP=(OLD,DELETE)
//
```

Option 7

Option 7 instructs Dyna-METRIC to prepare a report about minimum and maximum depot workloading. Two new files are added to PIPE:

File 10 - Output, formatted. Depot workload report (with table headings) for each time of analysis. LRECL=133. This report is usually sent to the printer.

File 11 - Output, formatted. Similar to File 10, this report is for every internal (scaled) day and includes no headings. It was intended to be input to a statistical program such as SAS. LRECL=133.

Option 10

Option 10 reads (rather than computes) the expected contents of the peacetime pipeline. A data set of expected peacetime pipeline values must already have been prepared in a previous run that invoked either option 16 (saves values from a Dyna-METRIC run) or option 22 (builds the file with user-supplied data). Option 10 requires an adjustment to File 1 in PIPE.

File 1 - Formatted, input. Saved pipeline values (MOD's File 1) from a previous model run with either option 16 or 22 invoked. LRECL=80.

Option 16

Option 16 constructs a file of expected pipeline values to be used as expected peacetime pipelines in a subsequent model run in which option 10 will be used. MOD requires a new file.

File 1 - Formatted, output. The data set used to restart a later model run. Used as File 1 in PIPE when option 10 is selected. LRECL=80.

Option 20

Option 20 reduces the daily LRU removals to reflect the fact that not all requested sorties will be FMC sorties--in particular, sorties may be flown with individual LRUs not operational, and hence not subject to failure (they've already failed, and we don't want to double count them).

PIPE has several changes. File 8 becomes an internal, unformatted file, and two new files are added.

File 7 - Output, formatted. Listing of LRUs that prevent achieving the requested sortie rate as well as the achievable sortie rates associated with them. LRECL is the greater of 80 and (DMCHANGE×8)+4.

File 9 - Unformatted, output. Used as File 2 in MOD.

Option 21

Option 21 allows the user to specify LRU-dependent flying programs by reading a new file in PIPE.

File 12 - Input, formatted. A record is needed for each LRU-base combination. Up to DMTIME different wartime periods are permitted.

Columns 1-16: LRU name
Columns 17-26: peacetime flying hours
Columns 27-36: flying hours for 1st wartime period
Columns 37-46: flying hours for 2nd wartime period, etc.

Records should be ordered as follows:

Record for 1st LRU at 1st base
Record for 1st LRU at 2nd base
.
.
Record for 1st LRU at last base
Record for 2nd LRU at 1st base
.
.
etc.

Option 22

Option 22 interprets a file of user-supplied peacetime pipeline values into a format usable in a subsequent Dyna-METRIC run (with option 10). For components not explicitly included in the input file, the steady state computation is used.

Option 22 requires only two steps, PART and this version of MOD.

File 5 - Formatted, input. The partitioned data set. (PART's File 8).

File 6 - Formatted, output. Error messages.

File 1 - Formatted, output. The data set used to restart a later model run. Used as File 1 in PIPE when option 10 is selected. LRECL=80.

File 9 - Formatted, input. Peacetime pipeline values. These records may appear in any order.

Columns 1-16: LRU, SRU or subSRU name.

Columns 18-21: location name.

Columns 22-32: retrograde pipeline contents (from the named location to a higher echelon).

Columns 33-43: number of this component experiencing administrative delays at named location.

Columns 44-54: number of this component in the repair pipeline at named location.

Columns 55-65: number ordered but not yet received.

Appendix B

INSTALLATION OF DYNA-METRIC VERSION 4

Before compiling the model, modify the source code to size the model and conform to the system on which it is to be run.

MODEL SIZING PARAMETERS

Sizing parameters limit the size of the problem that can be analyzed. They should be set to reasonable values for your system--large enough to be useful but small enough to run and save resources. We typically set the parameters with a global change command within the text editor. For example, a sed command is "s/DMAIRCFT/500/g" and its WYLBUR counterpart is "CHANGE 'DMAIRCFT' TO '500' IN ALL".

The following parameters must be set in each of the model's five programs before they are compiled. Standard values we use are given in the left column.

500	DMAIRCFT	- The maximum number of aircraft (plus one) that may be stationed at a base on any day of the scenario.
10	DMANALYS	- The maximum number of times of analysis.
10	DMBASES	- The maximum number of bases to be analyzed. (A base with an identical base count greater than 1 on its BASE record is considered one base.)
10	DMCHANGE	- The maximum number of occasions that data such as sortie rates and aircraft levels may change during a scenario.
2	DMCIRFS	- The maximum number of cirfs to be analyzed.
2	DMDEPOTS	- The maximum number of depots to be analyzed.
30.	DMEXPMAX	- A number such that exp(-DMEXPMAX) > 0. This standard value works well on our Vax (UNIX) and IBM systems, as well as on some Honeywells.
14	DMLOCS	- The maximum number of locations to be analyzed (DMBASES + DMCIRFS + DMDEPOTS).
100	DMLRUGRP	- The maximum number of LRUs per data partition. (With constrained repair, DMLRUGRP must be at least as large as DMLRUTEQ.)

5	DMLRUSRU	- The maximum number of LRUs to which a single SRU may be indentured.
1000	DMLRUS	- The maximum number of LRUs to be analyzed.
50	DMLRUTEQ	- The maximum number of LRUs that may be assigned to a single constrained repair resource.
5	DMMISSNS	- The maximum number of mission types to be analyzed.
23	DMOPTION	- The maximum number of available program options.
2200	DMPARTS	- The maximum number of parts to be analyzed (DMLRUS + DMSRUS + DMSSRUS).
7500	DMPMFMAX	- The maximum number of occurrences of an LRU at a single base. The number of occurrences is the LRU's QPA times the number of aircraft, plus stock. Note that this number may increase if stock is being purchased. Only one vector is dimensioned by this parameter, so use a large value.
200	DMPRTGRP	- The maximum number of parts per data partition (DMLRUGRP + DMSRUGRP + DMSUBGRP).
50	DMSRUGRP	- The maximum number of SRUs per data partition. This should at least be as large as DMSRULRU.
50	DMSRULRU	- The maximum number of SRUs that may be assigned to a single LRU.
600	DMSRUS	- The maximum number of SRUs to be analyzed.
600	DMSSRUS	- The maximum number of subSRUs to be analyzed.
10	DMSTANDS	- The maximum number of servers per constrained repair resource that may be assigned to a single location.
50	DMSUBGRP	- The maximum number of subSRUs per data partition. This should be at least as large as the maximum number of subSRUs that are (through SRUs) indentured to a single LRU.
50	DMSUBS2S	- The maximum number of subSRUs that may be assigned to a single SRU.
5	DMS2SUBS	- The maximum number of SRUs to which a single subSRU may be indentured.
5	DMTEQTYP	- The maximum number of constrained repair resources per data partition.

30 DMTIME - The maximum number of days in a scenario that can be analyzed without time-scaling the data. (Note: the model automatically time-scales if you select a time of analysis greater than DMTIME. Results are essentially the same as those attained without time scaling, except that some of the dynamics may be "smoothed.")

50 DMTQLRUS - The maximum number of LRUs in a single data partition that may be assigned to constrained repair of any sort.

CHARACTER VS. REAL TYPE DECLARATIONS

Dyna-METRIC is coded in Fortran '66. Beware that Fortran '66 is not completely upwardly compatible with Fortran '77. In particular, Fortran '66 does not support the CHARACTER type declaration. Some Fortran '77 compilers insist that all character variables be declared as CHARACTER rather than REAL. In such cases, you must redeclare the character variables as CHARACTER*4. Identification of the character variable declarations is straightforward because they all follow one of these comments:

```
C COMMON CHARACTER SCALAR:
C COMMON CHARACTER SCALARS:
C COMMON CHARACTER ARRAY:
C COMMON CHARACTER ARRAYS:
C LOCAL CHARACTER SCALAR:
C LOCAL CHARACTER SCALARS:
C LOCAL CHARACTER ARRAY:
C LOCAL CHARACTER ARRAYS:
C CHARACTER SCALAR:
C CHARACTER SCALARS:
C CHARACTER ARRAY:
C CHARACTER ARRAYS:
```

For example, you would need to replace lines that say

```
C COMMON CHARACTER SCALAR:
C
      REAL BLANK
```

with lines that say

```
C COMMON CHARACTER SCALAR:
C
      CHARACTER*4 BLANK
```

EXPONENTIAL UNDERFLOWS

Dyna-METRIC computations often lead to underflows, numbers too small to be represented by the computer. Some systems generate an error message when an underflow occurs. Typically the message is only a warning and the underflow is set to zero. This is fine as far as Dyna-METRIC is concerned; a number so close to zero as to be indistinguishable from zero by the computer may as well be zero. Other systems consider an underflow to be an error and halt program execution without completing the analysis.

Most systems have a way of turning off these messages. On the IBM, we insert the following statement into the main routine of PIPE and REPORT just after the variable definitions:

```
CALL ERRSET(208,256,-1,1,0,0)
```

OPEN STATEMENTS

Some systems (such as our Vax with the UNIX operating system) require that the files used to run Dyna-METRIC (described in Appendix A) be defined in the source code with OPEN statements. OPEN statements should appear after the variable declarations in the main routine of each program. For example, File 2 in PIPE might be opened with the following statement:

```
OPEN (2,FILE='FILE2',FORM='UNFORMATTED')
```

Appendix C

ERROR AND WARNING MESSAGES

Error	Message
1	VERSION 4.4 not entered in columns 20-30 of the second record. "VERSION 4.4" must be entered in columns 20-30 of the second record to indicate a standard Dyna-METRIC data set in Version 4 format. If the data set is in a different format, run it through the data set conversion program before proceeding.
2	ACFT/ATTR/BASE/CIRF/DEPT/FLHR/ILM/MESL/OPT/SRTS/TRNS/TURN records entered more than once. Consolidate all records of the indicated type into a single set, removing duplicate records. Make sure these header words do not start a component name or name a location.
3	ACFT/ATTR/FLHR/ILM/MESL/SRTS/TRNS/TURN records encountered before BASE records. The BASE record group must appear before these other record groups. Move the offending record group after the BASE records and before the first partition of LRU records.
4	APPL/INDT/QPA/SRU/SSRU/STK/VTM records encountered more than once in the same group of parts. The named record group may appear only once per data partition. Consolidate all similar types of records in the current partition into a single set, removing duplicate records. Make sure these header words do not start a component name or name a location.
5	APPL/INDT/QPA/SRU/SSRU/STK/TBED/TEST/TPRT/VTM records encountered before LRU records. Within a data partition, the LRU description records must appear before these other record types. Move the offending set of records after the LRU records to which they are related, but before any subsequent sets of LRU records.
6	CIRF records encountered after BASE records. The CIRF record group must appear before the BASE record group. (This is so the cirf named on a BASE record can be matched against the list of valid cirf names.) Put the groups in proper order.

7 No BASE/OPT records encountered in input.

At least one BASE record and one OPT record are required for Dyna-METRIC to run. Make sure the named record group is in the data set and appears before the first LRU partition.

8 No LRU records encountered in input.

At least one LRU record is required for Dyna-METRIC analysis. Make sure the LRU record group exists and has the header record, LRU. Check that there are no END records appearing before the LRU record group.

9 No LRUs have been assigned to ATE type TTTT.

The TPRT record group that assigns LRUs to repair resource TTTT is missing, misplaced, or blank. Make sure LRUs have been assigned to TTTT and that the TPRT records immediately follow the TEST and TBED record groups associated with TTTT (but precede subsequent TEST record groups and LRU partitions).

10 INDT records encountered before SRU records.

Within a data partition, the SRU records must appear before the INDT records. Move the INDT records after the SRU (and SSRU) records for that partition. Do not specify an indenture structure in any partition that contains no SRUs.

11 TBED records entered more than once for ATE type TTTT.

One set of TBED (server level) records is allowed per repair resource, but multiple sets were found for repair resource TTTT. If all TBED records apply to TTTT, consolidate the records into one set and remove duplicates. If some apply to other repair resources, move them so they immediately follow the TEST record group for that repair resource (preceding subsequent TEST record groups and LRU partitions).

12 TBED/TPRT records encountered before TEST records.

TBED or TPRT records were encountered in a data partition before a TEST record was found. The model does not know to which constrained repair resource the records apply. Make sure that the appropriate set of TEST records precedes these other sets of records.

13 TPRT records entered more than once for ATE type TTTT.

One set of TPRT records is allowed per repair resource, but multiple sets were found for repair resource TTTT. If all TPRT records apply to TTTT, consolidate the records into one set and remove duplicates. If some apply to other repair resources, move them so they immediately follow the TEST record group for that repair resource (preceding subsequent TEST record groups and LRU partitions).

14 TRNS records encountered before DEPT records.

The DEPT record group must appear before the TRNS record group. (This is so the depot described in the transportation records can be matched against the list of valid depots.) Move the TRNS record group so that it follows the DEPT, CIRF, and BASE record groups.

15 Unidentified base on ACFT/ATTR/FLHR/MESL/SRTS/TURN record. Name was BBBB.

A scenario record of the named type has a base named BBBB that does not match any bases named in the BASE record group. Make sure the name is correct and a record exists for BBBB. If BBBB is a cirf or depot, do not include scenario records for it (cirfs and depots cannot be assigned aircraft).

16 Multiple ACFT/ATTR/FLHR/MESL/SRTS/TURN records entered for base BBBB.

One scenario record of each type is permitted per base. Determine which record is correct and remove all other records of the indicated type for base BBBB.

17 Time of new rate on ACFT/ATTR/FLHR/MESL/SRTS/TBED/TURN record for base BBBB is negative.

The days on which new rates (such as aircraft levels and sortie rates) go into effect cannot be negative. Day 0 is peacetime and assumed to be steady-state. To handle a dynamic peacetime, consider the start of the war on a day beyond 1 (anything before that day is "peacetime").

18 Wartime offshore demand rate multiplier for LRU ABCDEFGHIJKLMNOP is negative.

The wartime offshore demand rate multiplier (in the VTM record group) for the named LRU multiplied by its offshore demand rate (in the LRU record group) gives the wartime demand rate at offshore bases. A negative multiplier implies a negative demand rate and is not permitted.

19 Times of new rates on ACFT/ATTR/FLHR/MESL/SRTS/TBED/TURN record for base BBBB are out of order.

The days on which new rates (such as aircraft levels and sortie rates) go into effect must be in ascending order. Leave blank extra (unused) fields for the remaining days on which new rates go into effect.

20 Aircraft level on ACFT record for base BBBB is either negative or greater than NNN.

Aircraft levels cannot be negative nor exceed NNN, the value compiled for the parameter DMAIRCFT (maximum number of aircraft per base). Recompile the model with DMAIRCFT set large enough.

21 Unidentified part on APPL/INDT/QPA/STK/TPRT/VTM record. Name was ABCDEFGHIJKLMNOP.

Only LRUs may be assigned application fractions (APPL), be assigned to a repair resource (TPRT), or have VTM records. An SRU or subSRU will not be recognized. Delete any SRUs and subSRUs in these record groups.

In the INDT record group, each record has a code identifying the type of part. "L" denotes an LRU, "S" an SRU, and "B" a subSRU. If this field has been miscoded, the part is not recognized.

If none of the above apply, then the named component does not match any names in the current partition of LRU, SRU, and subSRU records. Make sure the part name is correct, it has an appropriate description record, and it belongs in this partition.

22 Multiple VTM records entered for part ABCDEFGHIJKLMNOP.

One VTM record is permitted per LRU. Two or more have been found for the named LRU. Delete the duplicates.

23 Application fraction for LRU ABCDEFGHIJKLMNOP on APPL record not between zero and one.

For some base, the named LRU's application fraction is not valid (the application fraction is the probability that a randomly selected aircraft at the given base contains one or more instances of LRU ABCDEFGHIJKLMNOP). To indicate the quantity per aircraft, use the QPA field in the LRU or QPA record group.

24 Onshore/offshore switch on BASE record for base BBBB not equal to zero or one.

Blank or zero indicates an offshore base; one indicates an onshore base. No other values are meaningful.

25 Negative attrition rate on ATTR record for base BBBB.

Attrition rates cannot be negative. To increase the number of aircraft assigned to a base, use the ACFT record group.

26 More than one location named LLLL has been entered.

Each location (base, cirf or depot) must have a unique name. If a base is collocated with a cirf, assign them different names and set the transportation times between them to zero.

27 More than NN bases entered on BASE records.

The number of bases to be analyzed exceeds NN, the value compiled for the parameter DMBASES. Recompile the model with DMBASES set large enough. If some bases are identical (same scenario, stock levels, application fractions, etc.), use the identical base count in the BASE record group and enter only one BASE record.

28 Unidentified cirf serving base BBBB on BASE records.

Make sure a CIRF record exists for the cirf, the cirf name is correct on the BASE record, and the CIRF record group precedes the BASE record group. Do not specify a cirf name on the BASE record if the base has no cirf.

29 Base-to-cirf transportation for base BBBB on BASE record is negative.

The base-to-cirf transportation time in the BASE record group must be greater than or equal to zero. (A part cannot arrive at base BBBB's cirf before it departs from base BBBB.)

30 Cirf-to-base transportation for base BBBB on BASE record is negative.

The cirf-to-base transportation time in the BASE record group must be greater than or equal to zero. (A part cannot arrive at base BBBB before it departs from base BBBB's cirf.)

31 Switch specifying whether the peacetime cirf-to-base pipeline continues to empty for base BBBB on BASE record is neither 0 nor 1.

The cirf availability switch on the BASE record for base BBBB must be blank, zero, or one. No other values are valid.

32 Time on BASE record when cirf-to-base pipeline starts to empty for base BBBB is negative.

The cirf start time on the BASE record for base BBBB is negative. Resupply from the cirf is available in peacetime and again in wartime at the cirf start time. Set the start time to zero if resupply is always available.

33 Time on BASE record when cirf-to-base pipeline is initially cut off for base BBBB is negative.

The cirf cutoff time on the BASE record for base BBBB is negative. Resupply from the cirf is always available in peacetime and again in wartime at the cirf start time. Set the cutoff time to a large number (beyond the scenario) or set the cutoff duration to zero if resupply is always available.

34 Duration on BASE record of cutoff of cirf-to-base pipeline for base BBBB is negative.

The cirf cutoff duration on the BASE record for base BBBB is negative. Resupply from the cirf is always available in peacetime and again in wartime at the cirf start time. Set the cutoff duration to zero if resupply is always available.

35 Time for deployment of RRR repair at location LLLL is negative.

The RRR repair start time on the BASE/CIRF/DEPT record for location LLLL is negative. RRR repair is available in peacetime and again in wartime at the RRR repair start time. Set the start time to zero if repair is always available.

37 Time when resupply for location LLLL first becomes available is negative.

The resupply start time on the BASE/CIRF/DEPT record for location LLLL is negative. Resupply from an outside supplier is available in peacetime and again in wartime at the resupply start time. Set the start time to zero if resupply is always available.

38 Time for deployment of RR repair at location LLLL is negative.

The RR repair start time on the BASE/CIRF/DEPT record for location LLLL is negative. Repair is available in peacetime and again in wartime at the RR repair start time. Set the start time to zero if repair is always available.

40 Switch specifying whether the peacetime resupply pipeline continues to empty for location LLLL is neither 0 nor 1.

The resupply availability switch on the BASE/CIRF/DEPT record for location LLLL must be blank, zero or one. No other values are valid.

41 Switch specifying whether location LLLL cannibalizes SRUs is neither 0 nor 1.

The SRU/subSRU cannibalization switch on the BASE/CIRF/DEPT record for location LLLL must be blank, zero or one. No other values are valid.

42 Time for deployment of SRU repair at location LLLL is negative.

The SRU repair start time on the BASE/CIRF/DEPT record for location LLLL is negative. Repair of subcomponents is available in peacetime and again in wartime at the SRU repair start time. Set the start time to zero if repair is always available.

44 Time when resupply pipeline is initially cut off for location LLLL is negative.

The resupply cutoff time on the BASE/CIRF/DEPT record for location LLLL is negative. Resupply is always available in peacetime and again in wartime at the resupply start time. Set the cutoff time to a large number (beyond the scenario) or set the cutoff duration to zero if resupply is always available.

45 Duration of cutoff of resupply pipeline for location LLLL is negative.

The resupply cutoff duration on the BASE/CIRF/DEPT record for location LLLL is negative. Resupply is always available in peacetime and again in wartime at the resupply start time. Set the cutoff duration to zero if resupply is always available.

46 The number of bases identical to base BBBB specified on the BASE record is negative.

The identical base count on the BASE record for base BBBB is negative. If base BBBB is a unique base, code blank, 0 or 1. If there are N bases like base BBBB (including base BBBB), code N in the field.

47 No bases entered between BASE record and XXXX record.

The record immediately following the BASE header record was the XXXX header record; no base description records were found. There must be at least one base description record. Make sure locations are not named with a keyword such as BASE or CIRF.

48 More than NN cirfs entered on CIRF records.

The number of cirfs to be analyzed exceeds NN, the value compiled for the parameter DMCIRFS. Recompile the model with DMCIRFS set large enough.

49 More than NN depots entered on DEPT records.

The number of depots to be analyzed exceeds NN, the value compiled for the parameter DMDEPOTS. Recompile the model with DMDEPOTS set large enough.

50 Negative flying hours on FLHR record for base BBBB.

The flying hours per sortie record for base BBBB has a negative rate and is thus invalid.

51 More than one LRU named ABCDEFGHIJKLMNOP have been entered.

Each LRU must have a unique name. Delete duplicate records and rename incorrectly named LRUs.

52 Cost of part ABCDEFGHIJKLMNOP is negative.

The cost of the named part on its LRU/SRU/SSRU record is less than zero and thus invalid.

53 More than NNN LRUs entered on LRU records.

The number of LRUs within a partition exceeds NNN, the value compiled for the parameter DMLRUGRP. Make sure DMLRUGRP is large enough and that the data set is partitioned. If a repair resource serves more than NNN LRUs, recompile the model with DMLRUGRP greater than the maximum number of LRUs sharing that resource. If the presence of constrained repair forces a partition that violates DMLRUGRP, check the output from PART for the number of LRUs in the largest partition and reset DMLRUGRP.

54 Unidentified depot serving part ABCDEFGHIJKLMNOP.

The name of the depot serving this part does not match a name in the DEPT record group and it is not blank. Make sure a DEPT record exists for the depot and the depot name is correct on the part's LRU/SRU/SSRU record. Do not specify a depot name if the part is not assigned to a depot.

55 Level of repair for part ABCDEFGHIJKLMNOP is not between one and three.

The level of repair on the named part's LRU/SRU/SSRU record must be 1 (all echelons), 2 (cirf and depot), or 3 (depot only).

56 Quantity per aircraft for part ABCDEFGHIJKLMNOP is negative.

The quantity per aircraft on the named part's QPA record or LRU/SRU/SSRU record must be greater than or equal to zero.

57 Onshore demand rate for part ABCDEFGHIJKLMNOP is negative.

The onshore demand rate on the named part's LRU/SRU/SSRU record is negative and thus invalid.

58 Lone base repair cycle time for part ABCDEFGHIJKLMNOP is negative.

The lone base repair time on the named part's LRU/SRU/SSRU record is negative and thus invalid.

59 Depot repair cycle time for part ABCDEFGHIJKLMNOP is negative.

The depot repair time on the named part's LRU/SRU/SSRU record is negative and thus invalid.

60 NRTS rate for part ABCDEFGHIJKLMNOP at base not served by cirf is not between zero and one.

The lone base NRTS rate (or fraction of parts leaving repair that are sent to a higher echelon to complete repair) on the named part's LRU/SRU/SSRU record must, by definition, be between 0 and 1.

61 Condemnation rate for part ABCDEFGHIJKLMNOP at base not served by cirf is not between zero and one.

The lone base condemnation rate (or fraction of parts leaving repair that cannot be repaired and are condemned) on the named part's LRU/SRU/SSRU record must, by definition, be between 0 and 1.

62 NRTS rate plus condemnation rate for part ABCDEFGHIJKLMNOP at base not served by cirf is not between zero and one.

Condemnation rates are given as a percentage of all parts, not as a percentage of non-NRTS parts. By definition, the sum of the NRTS and condemnation rates at lone bases on the named part's LRU/SRU/SSRU record cannot exceed one.

63 Condemnation rate for part ABCDEFGHIJKLMNOP at a depot is not between zero and one.

The depot condemnation rate (or fraction of parts leaving repair that cannot be repaired and are condemned) on the named part's LRU/SRU/SSRU record must, by definition, be between 0 and 1.

64 Peacetime order and ship time for part ABCDEFGHIJKLMNOP is negative.

The peacetime resupply time on the named part's LRU/SRU/SSRU record is negative and thus invalid.

65 Wartime order and ship time for part ABCDEFGHIJKLMNOP is negative.

The wartime resupply time on the named part's LRU/SRU/SSRU record is negative and thus invalid.

66 Repair ceiling at the depot for part ABCDEFGHIJKLMNOP is negative.

The depot repair limit can equal -1 (no wartime depot repair), 0 (unlimited depot repair), or greater (depot can repair up to that amount daily in wartime). Some other value was entered on the named part's LRU/SRU/SSRU record.

67 Missing or incorrect second description record for part ABCDEFGHIJKLMNOP.

Each LRU, SRU and SSRU requires a pair of records, one immediately after the other, with the same name in the first 16 columns. Make sure the second LRU/SRU/SSRU record for the named part exists and has the same name as the first record.

68 Maintenance type specifier for LRU ABCDEFGHIJKLMNOP is not between zero and N.

The parameter DMCHANGE was compiled with a value of N+1. Valid maintenance type specifiers are 0,1,2,...N. On the VTM record for the named LRU, the maintenance type specifier is not within the range of valid values. Either correct it or recompile the model with DMCHANGE set large enough.

69 Wartime onshore demand rate multiplier for LRU ABCDEFGHIJKLMNOP is negative.

The wartime onshore demand rate multiplier (in the VTM record group) for the named LRU multiplied by its onshore demand rate (in the LRU record group) gives the wartime demand rate at onshore bases. A negative multiplier implies a negative demand rate and is not permitted.

70 Probability a PMC ATE can test LRU ABCDEFGHIJKLMNOP is not between zero and one.

The partial repairability probability on the named LRU's VTM record must be between 0 and 1.

71 Pipeline variance to mean ratio for LRU ABCDEFGHIJKLMNOP is negative.

Variance to mean ratios (VTM record group) cannot be negative because mean pipeline size and variances are, by definition, not negative.

72 Cirf repairable switch for part ABCDEFGHIJKLMNOP does not equal zero or one.

The cirf repairability switch on the named part's LRU/SRU/SSRU record is not 0 or blank (cirf cannot repair) and not 1 (cirf can repair).

73 NRTS rate for part ABCDEFGHIJKLMNOP at bases served by cirfs is not between zero and one.

The cirf-served base NRTS rate (or fraction of parts leaving repair that are sent to a higher echelon to complete repair) on the named part's LRU/SRU/SSRU record must, by definition, be between 0 and 1.

74 Condemnation rate for part ABCDEFGHIJKLMNOP at bases served by cirfs is not between zero and one.

The cirf-served base condemnation rate (or fraction of parts leaving repair that cannot be repaired and are condemned) on the named part's LRU/SRU/SSRU record must be between 0 and 1.

75 NRTS rate plus condemnation rate for part ABCDEFGHIJKLMNOP at bases served by cirfs is not between zero and one.

Condemnation rates are given as a percentage of all parts, not as a percentage of non-NRTS parts. By definition, the sum of the NRTS and condemnation rates at cirf-served bases on the named part's LRU/SRU/SSRU record cannot exceed one.

76 Cirf repair time for part ABCDEFGHIJKLMNOP is negative.

The cirf repair time on the named part's LRU/SRU/SSRU record is negative and thus invalid.

77 NRTS rate for part ABCDEFGHIJKLMNOP at a cirf is not between zero and one.

The cirf NRTS rate (or fraction of parts leaving repair that are sent to a higher echelon to complete repair) on the named part's LRU/SRU/SSRU record must, by definition, be between 0 and 1.

78 Condemnation rate for part ABCDEFGHIJKLMNOP at cirfs is not between zero and one.

The cirf condemnation rate (or fraction of parts leaving repair that cannot be repaired and are condemned) on the named part's LRU/SRU/SSRU record must, by definition, be between 0 and 1.

79 NRTS rate plus condemnation rate for part ABCDEFGHIJKLMNOP at cirfs is not between zero and one.

Condemnation rates are given as a percentage of all parts, not as a percentage of non-NRTS parts. By definition, the sum of the NRTS and condemnation rates at cirfs on the named part's LRU/SRU/SSRU record cannot exceed one.

80 Encountered end of file or a keyword record when looking for second description record for part ABCDEFGHIJKLMNOP.

Each LRU, SRU and SSRU requires a pair of records, one immediately after the other, with the same name in the first 16 columns. Make sure a header word does not start a component name and that there are no END records appearing before the second description record.

81 No LRU data found between LRU record and next keyword record.

At least one pair of LRU description records is required between the LRU header record and the next header record. Make sure LRU names do not start with keywords.

82 Option NN has not been defined.

Select only options described in the input specifications (App. E). Make sure the parameter DMOPTION was compiled large enough (set to 23, the highest defined option).

83 First parameter specified for option NN is not between 0 and 100.

The first parameter for option NN is a percentage and must be between 0 and 100.

84 First parameter for option 8 is negative.

Option 8's first parameter specifies the maximum number of problem LRUs on which to report and cannot be negative. Setting this option to zero is the same as not requesting this report.

86 Second parameter for option NN is negative.

Make sure the OPT record associated with option NN has a valid value for the second parameter as described in the input specifications (App. E).

87 Option NN has been selected more than once.

Each option can be selected once. Delete the duplicate OPT records for option NN.

88 No options have been selected.

At least one OPT record is required for Dyna-METRIC to run. Make sure the OPT record group is in the data set and appears before the first LRU partition.

89 Negative sortie rate on SRTS record foı base BBBB.

Negative sortie rates are invalid. If no sorties are to be generated, enter a rate of zero.

90 Duplicate parts named ABCDEFGHIJKLMNOP.

Each part must have a unique name. Change the name of the duplicate component so it is unique.

91 More than NNN SRUs entered on SRU records.

The number of SRUs within a partition exceeds NNN, the value compiled for the parameter DMSRUGRP. Make sure DMSRUGRP is large enough and that the data set is partitioned. For structural reasons (LRUs sharing the same constrained repair resource, SRUs indentured to the same LRU), the model may be forced to put more than NNN SRUs in the same partition. The output from PART gives the number of SRUs in the largest partition so that DMSRUGRP may be appropriately reset.

92 Negative ATE level specified on TBED record for ATE type N at location LLLL.

A server level on the TBED record for the Nth repair resource at location LLLL is negative and thus invalid.

93 ATE break rate on TBED record for ATE type N at location LLLL is negative.

The backorder rate on the TBED record for the Nth repair resource at location LLLL is negative and thus invalid.

94 Resupply time on TBED record for breaks of ATE type N at location LLLL is negative.

The resupply start time on the TBED record for the Nth repair resource at location LLLL is negative. Resupply is always available in peacetime and again in wartime at the resupply start time. Set the start time to zero if resupply is always available.

95 Time ATE resupply is cut off on TBED record for ATE type N at location LLLL is negative.

The resupply cutoff time on the TBED record for the Nth repair resource at location LLLL is negative. Resupply is always available in peacetime and again in wartime at the resupply start time. Set the cutoff time to a large number (beyond the scenario) or set the cutoff duration to zero if resupply is always available.

96 Duration of ATE resupply cutoff on TBED record for ATE type N at location LLLL is negative.

The resupply cutoff duration on the TBED record for the Nth repair resource at location LLLL is negative. Resupply is always available in peacetime and again in wartime at the resupply start time. Set the cutoff duration to zero if resupply is always available.

97 Unrecognized location name on TBED record for ATE type N. Location name was LLLL.

A TBED (server level) record associated with the Nth repair resource has a location named LLLL that does not match any locations named in the BASE, CIRF and DEPT record groups. Make sure the name is correct and a record exists for LLLL.

98 Multiple TBED records for ATE type N were entered for location LLLL.

More than one TBED record associated with the Nth repair resource was found for location LLLL. Determine which is correct and delete the duplicates.

99 Multiple TEST records entered for ATE TTTT.

One TEST record is permitted per repair resource. Make sure each repair resource has a unique name (other than TTTT) and its own TEST record group. Delete duplicate TEST records associated with TTTT.

100 More than NNN ATEs entered on TEST records.

The number of constrained repair resources within a partition exceeds NNN, the value compiled for the parameter DMTEQTYP. Make sure DMTEQTYP is large enough and that the data set is partitioned. PART puts as few repair resources in each partition as possible. Check the partitioned data set for the largest number of repair resources per partition and recompile the model with DMTEQTYP set at least that large.

101 Availability fraction entered on TEST record for ATE TTTT is not between zero and one.

Server availabilities (the fraction of the time a repair resource is available for repairing or testing LRUs per day) must, by definition, be between zero and one.

102 Only one ATE allowed per TEST record. TTTT on record following not recognizable as a keyword.

Each repair resource requires its own TEST, TBED and TPRT record groups. If TTTT is a repair resource, insert its three associated record groups before the next TEST record group or LRU partition.

103 End of file encountered when looking for ATE description following TEST record.

Each TEST header record requires a TEST data record. Make sure there is something following each TEST header record. Also make sure a repair resource is not illegally named END, causing the model to incorrectly believe it reached the end of file.

104 Keyword record encountered when looking for ATE description following TEST record.

Each TEST header record requires a TEST data record. Make sure there is something following each TEST header record. Also make sure a repair resource is not illegally named with a keyword (header), causing the model to incorrectly believe it reached another record group.

105 More than NN LRUs assigned to ATE TTTT.

The number of LRUs assigned to repair resource TTTT exceeds NN, the value compiled for the parameter DMLRUTEQ. Recompile the model with DMLRUTEQ set large enough.

106 Unrecognized depot on TRNS record. Name was DDDD.

A TRNS record has a depot named DDDD that does not match any depots named in the DEPT record group. Make sure the name is correct and a record exists for DDDD. Also make sure that the base/cirf name on the TRNS record was not switched with the depot name. Do not use TRNS records to specify transportation between a base and cirf: fields for that purpose are located in the BASE record group.

107 Base N to depot DDDD TRNS record encountered more than once.

There are multiple TRNS records describing the transportation between depot DDDD and the Nth base. Determine which record is correct and delete the duplicates.

108 Base N to depot DDDD transportation on TRNS record is negative.

The transportation to depot time on the TRNS record for depot DDDD and the Nth base cannot be negative. (A part cannot arrive at depot DDDD before it departs from base N.)

109 Depot DDDD to base N transportation on TRNS record is negative.

The transportation from depot time on the TRNS record for depot DDDD and the Nth base cannot be negative. (A part cannot arrive at base N before it departs from depot DDDD.)

110 Switch on TRNS record specifying whether peacetime forward pipeline from depot DDDD to base N continues to empty is not 0 or 1.

The transportation availability switch on the TRNS record for depot DDDD and the Nth base must be blank, zero or one. No other values are valid.

111 Time on TRNS record specifying when forward pipeline from depot DDDD to base N starts to empty is negative.

The transportation start time on the TRNS record for depot DDDD and the Nth base is negative. Transportation from the depot is available in peacetime and again in wartime at the transportation start time. Set the start time to zero if transportation is always available.

112 Time on TRNS record specifying when forward pipeline from depot DDDD to base N is cut off is negative.

The transportation cutoff time on the TRNS record for depot DDDD and the Nth base is negative. Transportation from the depot is always available in peacetime and again in wartime at the transportation start time. Set the cutoff time to a large number (beyond the scenario) or set the cutoff duration to zero if transportation is always available.

113 Time on TRNS record specifying how long forward pipeline from depot DDDD to base N is cut off is negative.

The transportation cutoff duration on the TRNS record for depot DDDD and the Nth base is negative. Transportation from the depot is always available in peacetime and again in wartime at the transportation start time. Set the cutoff duration to zero if transportation is always available.

114 Cirf N to depot DDDD TRNS record encountered more than once.

There are multiple TRNS records describing the transportation between depot DDDD and the Nth cirf. Determine which record is correct and delete the duplicates.

115 Cirf N to depot DDDD transportation on TRNS record is negative.

The transportation to depot time on the TRNS record for depot DDDD and the Nth cirf cannot be negative. (A part cannot arrive at depot DDDD before it departs from cirf N.)

116 Depot DDDD to cirf N transportation on TRNS record is negative.

The transportation from depot time on the TRNS record for depot DDDD and the Nth cirf cannot be negative. (A part cannot arrive at cirf N before it departs from depot DDDD).

117 Switch on TRNS record specifying whether peacetime forward pipeline from depot DDDD to cirf N continues to empty is not 0 or 1.

The transportation availability switch on the TRNS record for depot DDDD and the Nth cirf must be blank, zero or one. No other values are valid.

118 Time on TRNS record specifying when forward pipeline from depot DDDD to cirf N starts to empty is negative.

The transportation start time on the TRNS record for depot DDDD and the Nth cirf is negative. Transportation from the depot is available in peacetime and again in wartime at the transportation start time. Set the start time to zero if transportation is always available.

119 Time on TRNS record specifying when forward pipeline from depot DDDD to cirf N is cut off is negative.

The transportation cutoff time on the TRNS record for depot DDDD and the Nth cirf is negative. Transportation from the depot is always available in peacetime and again in wartime at the transportation start time. Set the cutoff time to a large number (beyond the scenario) or set the cutoff duration to zero if transportation is always available.

120 Time on TRNS record specifying how long forward pipeline from depot DDDD to cirf n is cut off is negative.

The transportation cutoff duration on the TRNS record for depot DDDD and the Nth cirf is negative. Transportation from the depot is always available in peacetime and again in wartime at the transportation start time. Set the cutoff duration to zero if transportation is always available.

121 Unrecognized base or cirf name on TRNS record. Name was BBCC.

A TRNS record has a base or cirf named BBCC that does not match any locations named in the BASE and CIRF record groups. Make sure the name is correct and a record exists for BBCC. Also make sure the base/cirf name on the TRNS record was not switched with the depot name. Do not use TRNS and DEPT records for BBCC if it is a supplier to a depot.

122 Negative maximum turn rate on TURN record for base BBBB.

The TURN record for base BBBB has at a negative maximum sorties per aircraft rate and is thus invalid.

123 NRTS/condemnation switch not 0 or 1 for part ABCDEFGHIJKLMNOP.

The NRTS/condemnation policy on the named part's LRU/SRU/SSRU record can only be blank or 0 (NRTS/condemn/discover failed subcomponents after attempting repair) or 1 (NRTS/condemn/discover failed subcomponents before attempting repair).

124 Offshore demand rate for part ABCDEFGHIJKLMNOP is negative.

The offshore demand rate on the named part's LRU/SRU/SSRU record is negative and thus invalid.

125 Cirf-served base repair cycle time for part ABCDEFGHIJKLMNOP is negative.

The repair time at cirf-served bases on the named part's LRU/SRU/SSRU record is negative and thus invalid.

126 Encountered ACFT/ATTR/BASE/CIRF/DEPT/FLHR/ILM/MESL/OPT/SRTS/TRNS/TURN record in parts data.

The named record group must precede the first LRU header record. Move it so that it does.

127 More than NNN subSRUs entered on SSRU records.

The number of subSRUs within a partition exceeds NNN, the value compiled for the parameter DMSUBGRP. Make sure DMSUBGRP is large enough and that the data set is partitioned. For structural reasons (LRUs sharing the same constrained repair resource, SRUs indentured to the same LRU, subSRUs indentured to the same SRU), the model may be forced to put more than NNN subSRUs in the same partition. The output from PART gives the number of subSRUs in the largest partition so that DMSUBGRP may be appropriately reset.

128 INDT does not contain recognizable LRU/SRU/subSRU identification.

The indenture mode specifier on an INDT record does not contain an "L" (LRU), "S" (SRU) or "B" (subSRU).

129 Part ABCDEFGHIJKLMNOP is being indentured to a part of the wrong level.

LRUs may not be indentured to any other part. SRUs may only be indentured to LRUs, subSRUs only to SRUs. A component indentured to both an LRU and an SRU should be treated as two separate components, one an SRU and one a subSRU.

130 More than NN SRUs indentured to LRU ABCDEFGHIJKLMNOP.

The number of SRUs indentured to the named LRU exceeds NN, the value compiled for the parameter DMSRULRU. Recompile the model with DMSRULRU set large enough.

131 SRU ABCDEFGHIJKLMNOP indentured to more than NN LRUs.

The number of LRUs to which the named SRU is indentured exceeds NN, the value compiled for the parameter DMLRUSRU. Recompile the model with DMLRUSRU set large enough.

132 Quantity per application in INDT records for part ABCDEFGHIJKLMNOP is not positive.

Quantities per application, or higher assembly, must be greater than zero. Delete records with zero quantities.

133 Depot replacement percentage for part ABCDEFGHIJKLMNOP is not between zero and one.

The depot replacement percentage (the probability that, if the LRU/SRU to which the named part is indentured enters repair at the depot, the depot will discover that the named part has failed) in the INDT record group must be between zero and one.

134 More than NN subSRUs indentured to SRU ABCDEFGHIJKLMNOP.

The number of subSRUs indentured to the named SRU exceeds NN, the value compiled for the parameter DMSUBS2S. Recompile the model with DMSUBS2S set large enough.

135 SubSRU ABCDEFGHIJKLMNOP indentured to more than NN SRUs.

The number of SRUs to which the named subSRU is indentured exceeds NN, the value compiled for the parameter DMS2SUBS. Recompile the model with DMS2SUBS set large enough.

136 Cirf replacement percentage for part ABCDEFGHIJKLMNOP is not between zero and one.

The cirf replacement percentage (the probability that, if the SRU to which the named subSRU is indentured enters repair at the cirf, the cirf will discover that the subSRU has failed) in the INDT record group must be between zero and one.

137 Base replacement percentage for part ABCDEFGHIJKLMNOP is not between zero and one.

The base replacement percentage (the probability that, if the SRU to which the named subSRU is indentured enters repair at the base, the base will discover that the subSRU has failed) in the INDT record group must be between zero and one.

138 SRU or subSRU ABCDEFGHIJKLMNOP not indentured to any part.

Make sure the named subcomponent has an INDT record if it is indentured to another part. Otherwise, this part is effectively an LRU and should be treated as such; move it to the LRU record group.

139 CIRF/DEPT records encountered after ILM records.

The ILM (maintenance deployment) record group must follow the DEPT, CIRF and BASE record groups. Move the ILM records after the location description records.

140 Multiple ILM records were entered for location LLLL.

One ILM record is permitted per location. Determine which record is correct and delete the duplicates.

141 Unrecognized location name on ILM record. Location name was LLLL.

An ILM record has a location named LLLL that does not match any locations named in the BASE, CIRF and DEPT record groups. Make sure the name is correct and a record exists for LLLL.

142 Record following Times of Analysis Record is not a keyword record. Record read XXXX.

The fourth record in the data set must be a header record. Instead, a record was found with the first four characters XXXX. Add the appropriate header record (usually OPT).

143 Options 10 and 20 cannot both be enabled in the same model run.

Option 10 will not work correctly if option 20 has been selected. Do not select both of these options in the same run.

144 When option 7 is invoked, only the initial stock levels plus any purchased stock are considered. Additional deployed stock is disregarded.

Option 7 produces a depot workload report for each LRU and computes appropriate depot stock levels. It ignores STK records other than those giving the initial levels or those it computes. This warning does not affect results of the run. To be safe, omit STK records that do not reflect initial stock levels.

Appendix D

STOP CODES

Program & Stop Code		Meaning
ECHO	20	An error was encountered in the input data set. Specific error messages were written at the top of the output.
PIPE	21	The file of LRU-specific flying hours that option 21 reads is not in the expected order. See Appendix A for the description of this file and its order.
PIPE	80	Peacetime demands for constrained repair at some location exceed peacetime repair capacity. The message identifies the location and the repair resource. After confirming that appropriate LRU failure rates and repair times have been entered, add sufficient servers to handle peacetime demands.
PIPE	211	Option 21 has an invalid value in the first parameter (the time period associated with the flying hours). See Appendix E for more details on the use of this parameter.
PIPE	222	"VERSION 4.4" not written in the appropriate columns of the second record of the data set. Refer to Appendix E for the proper location. Make sure that the data set is in Version 4 format.
PIPE	1001-1011	The records in the peacetime pipeline file that option 10 reads are not in proper order. Make sure the file was built using option 16 or 22, and that the input used in this run has the same locations entered in the same order as those in the data set with option 16 or 22. The LRUs, SRUs and subSRUs must remain the same, as well as the indenture structure, level of repair, and cirf repairability indicator. Stock levels, transportation parameters, and repair parameters may change, however.
MOD	10	Too many locations in the data set. Make sure MOD was compiled with the same parameters as the rest of the model.
MOD	20	Too many depots in the data set. Make sure MOD was compiled with the same parameters as the rest of the model.
MOD	30	Too many LRUs in a partition. Make sure MOD was compiled with the same parameters as the rest of the model. The input data set must be partitioned.

MOD	40	Too many parts in a partition. Make sure MOD was compiled with the same parameters as the rest of the model. The input data set must be partitioned.
MOD	50	Too many SRUs indentured to one LRU. Make sure MOD was compiled with the same parameters as the rest of the model.
MOD	60	Too many subSRUs indentured to one SRU. Make sure MOD was compiled with the same parameters as the rest of the model.
MOD	70	End of the input data set encountered before expected. Make sure the entire input data set is available to MOD.
REPORT	30	More than DMAIRCFT found at a base. Recompile with a larger value of DMAIRCFT or assign fewer aircraft to the base.
REPORT	31	Perhaps because of stock purchased, the number of a given LRU at a base exceeds DMPMFMAX. (The number of a given LRU at a base is the LRU's quantity per aircraft times the aircraft, plus stock.) Either increase DMPMFMAX or make sure the performance goals associated with the stockage options have reasonable arguments (no NFMC aircraft with 100 percent confidence would require an infinite level of stock, for example).
REPORT	32	The subroutine that computes pipeline probability distributions was unable to do so because some LRU had a negative VTMR.
REPORT	33	The subroutine that computes backorder probability distributions was unable to do so because it encountered a negative stock level.
REPORT	111	None of the times of analysis in the input data set match the times of analysis that were used by PIPE to compute expected pipeline values.
REPORT	241	When option 4 is selected, the model adds units of stock based on a marginal analysis technique, where the next unit added is that with the greatest benefit/cost ratio. STOP 241 occurs when the benefit/cost ratio for every LRU is computationally indistinguishable from zero. Select option 3 with more stringent parameters than option 4 so that when option 4 starts, there is sufficient stock that the benefit/cost ratio can be distinguished from zero. Or translate all costs from dollars to, say, millions of dollars and see if that increases the size of the benefit/cost ratio enough that the marginal analysis can proceed.

Appendix E

INPUT SPECIFICATIONS

General notes:

- All fields must be nonnegative unless otherwise specified.
- A blank is the same as a zero.
- Times are specified as days, but any consistent unit is acceptable.
- Peacetime values, represented by day 0, are assumed to have been in effect forever. Wartime begins on day 1 and lasts throughout the scenario.

Order and Structure of Input Data Sets

Data Set	Record Group
Administrative Data:	
Three records that specify the run's heading, administrative delays, and times of analysis.	TOP
Option selections.	OPT
Location Descriptions:	
Depot descriptions (optional)	DEPT
Cirf descriptions (optional)	CIRF
Base descriptions	BASE
Depot transportation (optional)	TRNS
Scenario Data:	
Aircraft levels	ACFT
Sortie rates	SRTS
Maximum sortie rates	TURN
Attrition rates (optional)	ATTR
Flying hours per sortie (optional)	FLHR
Maintenance types (optional)	ILM
Mission requirements (optional)	MESL
Component Data:	
LRU descriptions	LRU
SRU descriptions (optional)	SRU
SubSRU descriptions (optional)	SSRU
Indenture relationships (if have SRUs/subSRUs)	INDT
Quantity per aircraft (optional)	QPA
Optional LRU data:	
Application fractions	APPL
Variance to mean data	VTM
Constrained repair data:	
Constrained repair availability	TEST
Server Levels	TBED
LRU repair assignments	TPRT
Stock Levels.	STK

USING IMPLIED DECIMAL POINTS

Three classes of data are used in Dyna-METRIC: alphanumeric (a), integer (i), and real (x). The fields associated with real data use an implied decimal point, which may be overridden with an explicit decimal point.

To use the implied decimal point, we need to know where the point is located. Throughout this appendix, "X" denotes a digit to the left of the decimal point and "x" denotes a digit to the right. For example, a format representation of "XXXxx" for columns 5-9 means there is an implied point between columns 7 and 8.

An easier method is to include a decimal point explicitly in the data. The data need only fit within the columns to which they are assigned and need not be right-justified. This method has the advantage of allowing a greater range of entries but the disadvantage of using one of the columns in the field to contain the decimal point, thus reducing by one the available significant digits that can be used.

Whichever method the user chooses depends on the values to be coded. Suppose ten and a half is to be coded in the field defined "XXXxx". We could use either method, entering " 1050" to use the implied point or overriding it with "10.5 " or " 10.5". In this example the range of entries is .01 to 999.99 for the implied point with up to five significant digits, and .0001 to 9999. for the explicit decimal point with up to four significant digits.

Below are examples illustrating how various values can be coded.

Format	Value to be Coded	Explicit Decimal Point	Implied Decimal Point
XXxxxx	2.5	"2.5 "	" 25000"
	999.03	"999.03"	--
	99.0359	--	"990359"
XXX	2.8	"2.8"	--
	725.	--	"725"
	23.	"23."	" 23"

Aircraft Levels

Header record: ACFT

Restriction: Must follow the BASE record group.

General description:
These records specify how many aircraft are assigned to each base during peacetime and on each day of war. Bases for which ACFT records are not given will be assigned no aircraft.

```
Columns
         1         2         3         4
123456789012345678901234567890123456789012345678

aaaaXXXXiiiiXXXXiiiiXXXXiiiiXXXXiiiiXXXXiiiiXXXX
|   |   |   |   |   |   |   |   |   |   |   |
|   |   |   |   |   |   |   |   |   |   |   Sixth aircraft level
|   |   |   |   |   |   |   |   |   |   Day sixth level starts
|   |   |   |   |   |   |   |   |   Fifth aircraft level
|   |   |   |   |   |   |   |   Day fifth level starts
|   |   |   |   |   |   |   Fourth aircraft level
|   |   |   |   |   |   Day fourth level starts
|   |   |   |   |   Third aircraft level
|   |   |   |   Day third level starts
|   |   |   Second aircraft level
|   |   Day second level starts
|   First aircraft level
Base name
```

Detailed description:
The maximum number of aircraft levels is limited by the value of DMCHANGE, selected when the model was compiled. Extra fields for more than six levels follow the format of the previous fields. Each day a level starts must be greater than the previous day.

For example, a record containing

```
BAS1   0   1   18  8   6
```

indicates that base BAS1 has no aircraft in peacetime. Starting on day 1 of war, BAS1 has 18 aircraft. On day 8, the level reduces to six aircraft that lasts throughout the scenario.

Columns	Format	
1-4	aaaa	Base name. The name of the base for which aircraft levels are specified. Must be named in the BASE record group. Enter at most one record per base.
5-8	XXXX	First aircraft level.
9-12	iiii	Day second level starts.
13-16	XXXX	Second aircraft level.
17-20	iiii	Day third level starts.
21-24	XXXX	Third aircraft level.
25-28	iiii	Day fourth level starts.
29-32	XXXX	Fourth aircraft level.
33-36	iiii	Day fifth level starts.
37-40	XXXX	Fifth aircraft level.
41-44	iiii	Day sixth level starts.
45-48	XXXX	Sixth aircraft level.

Application Fractions

Header record: APPL

Restriction: Must follow the LRU record group.

General description:
These records specify the fraction of each base's aircraft that contain a given LRU. Bases for which application fractions are not specified default to application fractions of one (all aircraft at the base contain the LRU).

```
Columns
          1         2         3         4
1234567890123456789012345678901234567890123456

aaaaaaaaaaaaaaaa aaaaXXXxx aaaaXXXxx aaaaXXXxx
|                |   |     |   |     |   |
|                |   |     |   |     |   Third application fraction
|                |   |     |   |     Third base name
|                |   |     |   Second application fraction
|                |   |     Second base name
|                |   First application fraction
|                First base name
LRU name

           4    5         6         7
           890123456789012345678901234567

           aaaaXXXxx aaaaXXXxx aaaaXXXxx
           |   |     |   |     |   |
           |   |     |   |     |   Sixth application fraction
           |   |     |   |     Sixth base name
           |   |     |   Fifth application fraction
           |   |     Fifth base name
           |   Fourth application fraction
           Fourth base name
```

Detailed description:
Each application fraction must be between zero and one. Base names that do not match names entered in the BASE record group will generate a warning message. Application fractions for up to six bases may be entered on a single record. If more are needed, multiple records are allowed.

For example, records containing

```
WIDGET           BAS1  .50 BAS2  .50 BAS3  .50 BAS4  .50
WIDGET           BAS5 1.00 BAS6 1.00
```

indicate that LRU WIDGET appears on half the aircraft at the first four bases and on all the aircraft at the last two bases. (Note that the data could have been entered on one record.)

Columns	Format	
1-16	a16	LRU name. The name of the LRU for which application fraction data are specified. Must be named in the LRU record group. Enter as many records as needed per LRU.
18-21	aaaa	First base name. Name of the base to which first application fraction applies.
22-26	XXXxx	First application fraction. Fraction of the aircraft stationed at first base that contain the LRU.
28-31	aaaa	Second base name.
32-36	XXXxx	Second application fraction.
38-41	aaaa	Third base name.
42-46	XXXxx	Third application fraction.
48-51	aaaa	Fourth base name.
52-56	XXXxx	Fourth application fraction.
58-61	aaaa	Fifth base name.
62-66	XXXxx	Fifth application fraction.
68-71	aaaa	Sixth base name.
72-76	XXXxx	Sixth application fraction.

Attrition Rates

Header record: ATTR

Restriction: Must follow the BASE and ACFT record groups.

General description:
These records specify the fraction of aircraft that are attrited per sortie at each base on each day of the war (there is no peacetime attrition). Aircraft at bases for which ATTR records are not given do not experience attrition.

```
Columns
         1         2         3         4         5
12345678901234567890123456789012345678901234567890123 4

aaaaXXxxxiiiiXXxxxiiiiXXxxxiiiiXXxxxiiiiXXxxxiiiiXXxxx
|   |    |   |    |   |    |   |    |   |    |   |
|   |    |   |    |   |    |   |    |   |    |   Sixth attrition rate
|   |    |   |    |   |    |   |    |   |    Day sixth rate starts
|   |    |   |    |   |    |   |    |   Fifth attrition rate
|   |    |   |    |   |    |   |    Day fifth rate starts
|   |    |   |    |   |    |   Fourth attrition rate
|   |    |   |    |   |    Day fourth rate starts
|   |    |   |    |   Third attrition rate
|   |    |   |    Day third rate starts
|   |    |   Second attrition rate
|   |    Day second rate starts
|   First attrition rate
Base name
```

Detailed description:
The maximum number of attrition rates is limited by the value of DMCHANGE, selected when the model was compiled. Extra fields for more than six rates follow the format of the previous fields. Each day a rate starts must be greater than the previous day.

For example, a record containing

```
BAS1  .01   5  .02
```

indicates that base BAS1 loses .01 aircraft per sortie (one aircraft per hundred sorties) flown on the first four days of war. On day 5 and for the rest of the scenario, BAS1 loses two aircraft per hundred sorties flown.

Columns	Format	
1-4	aaaa	Base name. The name of the base for which attrition rates are specified. Must be named in the BASE record group. Enter at most one record per base.
5-9	XXxxx	First attrition rate.
10-13	iiii	Day second rate starts.
14-18	XXxxx	Second attrition rate.
19-22	iiii	Day third rate starts.
23-27	XXxxx	Third attrition rate.
28-31	iiii	Day fourth rate starts.
32-36	XXxxx	Fourth attrition rate.
37-40	iiii	Day fifth rate starts.
41-45	XXxxx	Fifth attrition rate.
46-49	iiii	Day sixth rate starts.
50-54	XXxxx	Sixth attrition rate.

Base Descriptions

Header record: BASE

Restriction: Must follow the DEPT and CIRF record groups (if any).

General description:
These records describe the availability of repair and resupply at each base, name the base's cirf (if any), and describe the transportation resources connecting the base and cirf. A record is required for each base.

```
Columns
         1         2         3         4
1234567890123456789012345678901234567890

aaaaaaaaXXXxxXXXxxXXXxxiXXXxxXXXxxXXXxxi
|   |   |    |    |    ||    |    |    |
|   |   |    |    |    ||    |    |    Resupply availability switch
|   |   |    |    |    ||    |    Resupply start
|   |   |    |    |    ||    Cirf cutoff duration
|   |   |    |    |    |Cirf cutoff
|   |   |    |    |    Cirf availability switch
|   |   |    |    Cirf start
|   |   |    Cirf-to-base transportation time
|   |   Base-to-cirf transportation time
|   Cirf name
Base name

           4         5         6         7         8
           1234567890123456789012345678901234567890

           XXXxxXXXxxXXXxxXXXxxXXXxxiXXXxxXXXxxiiii
           |    |    |    |    |    ||    |    |  |
           |    |    |    |    |    ||    |    |  Onshore switch
           |    |    |    |    |    ||    |    Identical base count
           |    |    |    |    |    ||    Reparable arrival time
           |    |    |    |    |    |Sustained demand start time
           |    |    |    |    |    SRU cannibalization switch
           |    |    |    |    SRU repair start
           |    |    |    RRR repair start
           |    |    RR repair start
           |    Resupply cutoff duration
           Resupply cutoff
```

Detailed description:
Cirf resupply and cutoff parameters also apply to the retrograde (base-to-cirf) pipeline if the cutoff direction switch in the TOP record group is set.

Detailed description:

Columns	Format	
1-4	aaaa	Base name. The name of the base. This may be any string of four characters as long as it is neither a keyword (such as "BASE") nor the name of another base, cirf or depot.
5-8	aaaa	Cirf name. The name of the cirf that serves this base. Must be named in the CIRF record group. If the base is not served by a cirf, this field must be blank.
9-13	XXXxx	Base-to-cirf transportation time (in days).
14-18	XXXxx	Cirf-to-base transportation time (in days).
19-23	XXXxx	Cirf start. Day resupply from the cirf first becomes available.
24	i	Cirf availability switch. Set to 1 if parts ordered from the cirf in peacetime continue to arrive at the base prior to the cirf start time. Set to 0 or blank if they do not.
25-29	XXXxx	Cirf cutoff. Day resupply from the cirf is first cut off. The peacetime resupply pipeline from the cirf is also cut off, even if cirf availability is set.
30-34	XXXxx	Cirf cutoff duration. Number of days resupply from the cirf is cut off.
35-39	XXXxx	Resupply start. Day resupply of components ordered from a supplier other than the cirf or depot first becomes available.
40	i	Resupply availability switch. Set to 1 if parts ordered in peacetime from a supplier other than the cirf or depot continue to arrive at the base prior to the resupply start time. Set to 0 or blank if they do not.
41-45	XXXxx	Resupply cutoff. Day resupply of components ordered from a supplier other than the cirf or depot is cut off. The peacetime resupply pipeline is also cut off, even if resupply availability is set.

Columns	Format	
46-50	XXXxx	Resupply cutoff duration. Number of days that resupply of components ordered from a supplier other than the cirf or depot is cut off.
51-55	XXXxx	RR repair start. Day the base can start repairing LRUs coded RR. (This field is disregarded if an ILM record is included for the base.)
56-60	XXXxx	RRR repair start. Day the base can start repairing LRUs coded RRR. (This field is disregarded if an ILM record is included for the base.)
61-65	XXXxx	SRU repair start. Day the base can start repairing SRUs and subSRUs.
66	i	SRU cannibalization switch. Set to 1 if the base cannibalizes SRUs and subSRUs; set to 0 or blank if it does not. SRUs may only be cannibalized between identical LRUs that are AWaiting Parts. Similarly, subSRUs may only be cannibalized between identical SRUs that are AWP.
67-71	XXXxx	Sustained demand start time. Day that components begin to break according to their sustained demand rates (entered in the VTM record group) as opposed to their wartime demand rates. If set to 0 or blank, the wartime rates remain in effect for the entire wartime scenario.
72-76	XXXxx	Reparable arrival time. Day that peacetime reparables are deployed to the base. Prior to this day, peacetime reparables are not available to be repaired, even if repair facilities are available.
77-79	iii	Identical base count. Number of bases identical to this one to be analyzed. Do not explicitly enter records for the other bases. Blank, 0 or 1 implies a unique base.
80	i	Onshore switch. Set to 1 to indicate an onshore base; set to 0 to indicate an offshore base. (LRUs and SRUs can have two different demand rates, one for each base type.)

Cirf Descriptions

Header record: CIRF

Restriction: Must precede the BASE record group.

General description:
These records describe the availability of repair and resupply at each cirf.

```
Columns
      3         4         5         6         7
1234...0123456789012345678901234567890123456789012345

aaaa       XXXxxiXXXxxXXXxxXXXxxXXXxxXXXxxi     XXXxx
|          |    ||    |    |    |    |    |     |
|          |    ||    |    |    |    |    |     Reparable arrival time
|          |    ||    |    |    |    |    SRU cannibalization switch
|          |    ||    |    |    |    SRU repair start
|          |    ||    |    |    RRR repair start
|          |    ||    |    RR repair start
|          |    ||    Resupply cutoff duration
|          |    |Resupply cutoff
|          |    Resupply availability switch
|          Resupply start
Cirf name
```

Detailed description:

Columns	Format	
1-4	aaaa	Cirf name. The name of the cirf. This may be any string of four characters as long as it is neither a keyword (such as "CIRF") nor the name of another base, cirf or depot.
35-39	XXXxx	Resupply start. Day resupply of components ordered from a supplier other than the depot first becomes available.
40	i	Resupply availability switch. Set to 1 if parts ordered in peacetime from a supplier other than the depot continue to arrive at the cirf prior to the resupply start time Set to 0 or blank if they do not.

Columns	Format	
41-45	XXXxx	Resupply cutoff. Day resupply of components ordered from a supplier other than the depot is cut off. The peacetime resupply pipeline from the depot is also cut off, even if resupply availability is set.
46-50	XXXxx	Resupply cutoff duration. Number of days that resupply of components ordered from a supplier other than the depot is cut off.
51-55	XXXxx	RR repair start. Day the cirf can start repairing LRUs coded RR. (This field is disregarded if an ILM record is included for the cirf.)
56-60	XXXxx	RRR repair start. Day the cirf can start repairing LRUs coded RRR. (This field is disregarded if an ILM record is included for the cirf.)
61-65	XXXxx	SRU repair start. Day the cirf can start repairing SRUs and subSRUs.
66	i	*SRU cannibalization switch.* Set to 1 if the cirf cannibalizes SRUs and subSRUs; set to 0 or blank if it does not. SRUs may only be cannibalized between identical LRUs that are AWaiting Parts. Similarly, subSRUs may only be cannibalized between identical SRUs that are AWP.
72-76	XXXxx	Reparable arrival time. Day that peacetime reparables are deployed to the cirf. Prior to this day, peacetime reparables are not available to be repaired, even if repair facilities are available.

Depot Descriptions

Header record: DEPT

Restriction: Must precede the BASE record group.

General description:
These records describe the availability of repair and resupply at each depot.

```
Columns
      3         4         5         6         7
1234...0123456789012345678901234567890123456789012345 6

aaaa        XXXxxiXXXxxXXXxxXXXxxXXXxxXXXxxi     XXXxx
|           |    ||    |    |    |    |    |     |
|           |    ||    |    |    |    |    |     Reparable arrival time
|           |    ||    |    |    |    |    SRU cannibalization switch
|           |    ||    |    |    |    SRU repair start
|           |    ||    |    |    RRR repair start
|           |    ||    |    RR repair start
|           |    ||    Resupply cutoff duration
|           |    |Resupply cutoff
|           |    Resupply availability switch
|           Resupply start
Depot name
```

Detailed description:

Columns	Format	
1-4	aaaa	Depot name. The name of the depot. This may be any string of four characters as long as it is neither a keyword (such as "DEPT") nor the name of another base, cirf or depot.
35-39	XXXxx	Resupply start. Day resupply of components ordered from an outside supplier first becomes available.
40	i	Resupply availability switch. Set to 1 if parts ordered in peacetime continue to arrive at the depot prior to the resupply start time. Set to 0 or blank if they do not.
41-45	XXXxx	Resupply cutoff. Day resupply of components ordered from an outside supplier is cut off. The peacetime resupply pipeline is also cut off, even if resupply availability is set.

Columns	Format	
46-50	XXXxx	Resupply cutoff duration. Number of days that resupply of parts ordered from an outside supplier is cut off.
51-55	XXXxx	RR repair start. Day the depot can start repairing LRUs coded RR. (This field is disregarded if an ILM record is included for the depot.)
56-60	XXXxx	RRR repair start. Day the depot can start repairing LRUs coded RRR. (This field is disregarded if an ILM record is included for the depot.)
61-65	XXXxx	SRU repair start. Day the depot can start repairing SRUs and subSRUs.
66	i	SRU cannibalization switch. Set to 1 if the depot cannibalizes SRUs and subSRUs; set to 0 or blank if it does not. SRUs may only be cannibalized between identical LRUs that are AWaiting Parts. Similarly, subSRUs may only be cannibalized between identical SRUs that are AWP.
72-76	XXXxx	Reparable arrival time. Day that peacetime reparables are deployed to the depot. Prior to this day, peacetime reparables are not available to be repaired, even if repair facilities are available.

Flying Hours per Sortie

Header record: FLHR

Restrictions: Must follow the BASE record group.

General description:
These records specify how many flying hours are required per sortie at each base during peacetime and each day of the war. Aircraft at bases for which FLHR records are not given are assumed to fly sorties of one hour each.

```
Columns
         1         2         3         4
12345678901234567890123456789012345678901234567 8

aaaaXXXXiiiiXXXXiiiiXXXXiiiiXXXXiiiiXXXXiiiiXXXX
|   |   |   |   |   |   |   |   |   |   |   |
|   |   |   |   |   |   |   |   |   |   |   Sixth flying hours level
|   |   |   |   |   |   |   |   |   |   Day sixth level starts
|   |   |   |   |   |   |   |   |   Fifth flying hours level
|   |   |   |   |   |   |   |   Day fifth level starts
|   |   |   |   |   |   |   Fourth flying hours level
|   |   |   |   |   |   Day fourth level starts
|   |   |   |   |   Third flying hours level
|   |   |   |   Day third level starts
|   |   |   Second flying hours level
|   |   Day second level starts
|   First flying hours level
Base name
```

Detailed description:
The maximum number of flying hour levels is limited by the value of DMCHANGE, selected when the model was compiled. Extra fields for more than six levels follow the format of the previous fields. Each day a level starts must be greater than the previous day.

For example, a record containing

```
BAS1  2.   4 1.5
```

indicates that aircraft at base BAS1 fly sorties of two hours each in peacetime and through the first three days of war. From day 4 on, they fly sorties of one and a half hours each.

Columns	Format	
1-4	aaaa	Base name. The name of the base for which flying hours per sortie are specified. Must be named in the BASE record group. Enter at most one record per base.
5-8	XXXX	First flying hours level.
9-12	iiii	Day second level starts.
13-16	XXXX	Second flying hours level.
17-20	iiii	Day third level starts.
21-24	XXXX	Third flying hours level.
25-28	iiii	Day fourth level starts.
29-32	XXXX	Fourth flying hours level.
33-36	iiii	Day fifth level starts.
37-40	XXXX	Fifth flying hours level.
41-44	iiii	Day sixth level starts.
45-48	XXXX	Sixth flying hours level.

Maintenance Deployment

Header record: ILM

Restriction: Must follow the BASE record group.

General description:
These records name the different maintenance types that apply to the LRUs and specify when each becomes available at each location.

```
Columns (first record)
         1         2         3
12345678901234567890123456789012345

     aaaa aaaa aaaa aaaa aaaa aaaa
     |    |    |    |    |    |
     |    |    |    |    |    Name of maintenance type 5
     |    |    |    |    Name of maintenance type 4
     |    |    |    Name of maintenance type 3
     |    |    Name of maintenance type 2
     |    Name of maintenance type 1
     Name of maintenance type 0


Columns (remaining records)
         1         2         3
12345678901234567890123456789012345

aaaa XXXx XXXx XXXx XXXx XXXx XXXx
|    |    |    |    |    |    |
|    |    |    |    |    |    Day maintenance type 5 starts
|    |    |    |    |    Day maintenance type 4 starts
|    |    |    |    Day maintenance type 3 starts
|    |    |    Day maintenance type 2 starts
|    |    Day maintenance type 1 starts
|    Day maintenance type 0 starts
Location name
```

Detailed description:
The maximum number of maintenance types is limited by the value of DMCHANGE, selected when the model was compiled. Extra fields for more than six levels follow the format of the previous fields.

Maintenance types are designated by any string of four letters, not necessarily unique. If the ILM record group is omitted, maintenance type 0 defaults to RR and maintenance type 1 defaults to RRR, with all other maintenance types blank. Start times must be zero or greater. A start time between zero and one implies day 1.

For example, records containing

```
     ILM1 ILM2 JEIM
BAS1   3.   5.   0.
```

indicate that at base BAS1, maintenance type ILM1 becomes available on day 3 of war, ILM2 becomes available on day 5, and JEIM has always been available, even in peacetime.

First record:

Columns	Format	
6-9	aaaa	Name of maintenance type 0.
11-14	aaaa	Name of maintenance type 1.
16-19	aaaa	Name of maintenance type 2.
21-24	aaaa	Name of maintenance type 3.
26-29	aaaa	Name of maintenance type 4.
31-34	aaaa	Name of maintenance type 5.

Remaining record(s):

Columns	Format	
1-4	aaaa	Location name. The name of the location for which maintenance type start times are specified. Must be named in the BASE, CIRF or DEPT record group. Enter at most one record per location.
6-9	XXXx	Day maintenance type 0 starts.
11-14	XXXx	Day maintenance type 1 starts.
16-19	XXXx	Day maintenance type 2 starts.
21-24	XXXx	Day maintenance type 3 starts.
26-29	XXXx	Day maintenance type 4 starts.
31-34	XXXx	Day maintenance type 5 starts.

Indenture Relationships

Header record: INDT

Restriction: Must follow the LRU, SRU and (if any) SSRU record groups.

General description:

These records specify the relationships between LRUs and their indentured SRUs, and between SRUs and their subSRUs. Do not include records for LRUs that have no indentured SRUs. All SRUs must be indentured to at least one LRU and all subSRUs to at least one SRU.

```
Columns
         1         2         3
123456789012345678901234567890123456789

aaaaaaaaaaaaaaaaiaiii XXxxx XXxxx XXxxx
|               |||   |     |     |
|               |||   |     |     Base replacement fraction
|               |||   |     Cirf replacement fraction
|               |||   Depot replacement fraction
|               ||Quantity per higher assembly
|               |LRU/SRU/subSRU indicator
|               Indenture mode specifier
Part name
```

Detailed description:

Records appear in families consisting of parents and children. The first record in a family names the parent: either an LRU to which SRUs are indentured or an SRU to which subSRUs are indentured. The records following the parent name the children or indentured parts (SRUs or subSRUs). SRUs may appear more than once: as a child indentured to at least one parent LRU, and as a parent to which subSRUs are indentured.

For example, records containing

```
ENGINE          1L
AUGMENTOR        S  1
GEARBOX          S  1
AUGMENTOR       1S
VALVE            3  2
```

indicate that LRU ENGINE has two SRUs indentured to it, each with a quantity of one on that LRU. In addition, SRU AUGMENTOR has a subSRU indentured to it that has a quantity per next higher assembly of two.

Columns	Format	
1-16	a16	Part name. Name of LRU, SRU or subSRU.
17	i	Indenture mode specifier. Set to 1 for parent LRUs and SRUs (those having parts indentured to them). Set to blank or 0 for child SRUs and subSRUs (those that are indentured to the most recently named parent). SRUs may be indentured only to LRUs and subSRUs only to SRUs.
18	a	LRU/SRU/subSRU indicator. Enter "L" for LRUs, "S" for SRUs, and "B" for subSRUs.
19-21	iii	Quantity per higher assembly. Quantity of this part on the next higher assembly. For SRUs, its number on the parent LRU; for subSRUs, its number on the parent SRU. Quantities must be greater than zero. This field is disregarded for parent records (as designated in the indenture mode specifier).
23-27	XXxxx	Depot replacement fraction. Fraction of the part arriving at a depot on its next higher assembly that will be removed and replaced. This field is disregarded for parent records.
29-33	XXxxx	Cirf replacement fraction (subSRUs only). Fraction of the subSRU arriving at a cirf on its higher SRU assembly that will be removed and replaced.
35-39	XXxxx	Base replacement fraction (subSRUs only). Fraction of the subSRU on its higher SRU assembly in base repair that will be removed and replaced.

LRU Descriptions

Header record: LRU

Restriction: Must follow the scenario record groups.

General description:
These records describe the failure, repair and resupply characteristics of each LRU. A pair of these records is required for each LRU.

```
Columns (first record of pair)
         1         2         3         4
1234567890123456789012345678901234567890

aaaaaaaaaaaaaaa aaaa i iiiiiiiiXXxxxxx
|                |    | ||  |  |||
|                |    | ||  |  ||Onshore demand rate
|                |    | ||  |  |NRTS/condemnation/failed SRU policy
|                |    | ||  |  Demands per sortie indicator
|                |    | ||  Minimum quantity
|                |    | |Quantity per aircraft
|                |    | Cirf reparability switch
|                |    Level of repair
|                Depot name
LRU name

        4         5         6         7
        1234567890123456789012345678901234567

        XXxxxxxXXXxx XXxx XXxxXXXxx XXxx XXxx
        |      |      |    |   |     |    |
        |      |      |    |   |     |    Cirf-served base
        |      |      |    |   |     |        condemnation rate
        |      |      |    |   |     Cirf-served base NRTS rate
        |      |      |    |   Cirf-served base repair time
        |      |      |    Lone base condemnation rate
        |      |      Lone base NRTS rate
        |      Lone base repair time
        Offshore demand rate
```

Detailed description:
For LRUs not assigned to constrained repair, repair (cycle) time includes time awaiting maintenance and in work (and time awaiting parts if there are no SRUs). Otherwise, it is the time the repair resource is exclusively dedicated to the LRU (excluding time spent in the queue).

Condemnation rates are given as a percentage of all LRUs, not as a percentage on non-NRTS LRUs. Thus the sum of the condemnation and NRTS rates for a particular echelon cannot exceed one (100%).

```
Columns (second record of pair)
                 1         2         3         4
1234567890123456789012345678901234567890123456

aaaaaaaaaaaaaaaa XXxx XXxx XXxxXXXxx XXXx XXxx
|                |    |    |   |     |    |
|                |    |    |   |     |    Depot condemnation rate
|                |    |    |   |     Depot repair limit
|                |    |    |   Depot repair time
|                |    |    Cirf condemnation rate
|                |    Cirf NRTS rate
|                Cirf repair time
LRU name

          4  5         6         7
          78901234567890123456789012345

          XXXXxXXXXx XXXXXXXX aaaaaaa i
          |    |     |        |       |
          |    |     |        |       No cannibalization indicator
          |    |     |        Work unit code
          |    |     Cost
          |    Wartime resupply time
          Peacetime resupply time
```

First record of pair:

Columns	Format	
1-16	a16	LRU name. The name of the LRU. This may be any string of 16 characters as long as it has neither the name of another LRU, SRU or subSRU nor the first four characters of a keyword (such as "BASE1 ENGINE").
18-21	aaaa	Depot name. The name of the depot that repairs this LRU. Must be named in the DEPT record group. If the LRU is not repaired by a depot, this field must be blank.
23	i	Level of repair. Set to 1 if the LRU can be repaired at a base, cirf or depot. Set to 2 if the LRU can be repaired at a cirf or depot but not a base. Set to 3 if the LRU can only be repaired at a depot. Regardless of setting, LRU always experiences base administrative delay.

Columns	Format	
25	i	Cirf reparability switch. Set to 1 if a cirf can repair the LRU; set to 0 or blank if it cannot. This field is disregarded if the level of repair is 3 (depot only).
26-28	iii	Quantity per aircraft. Number of this LRU per aircraft at all bases. Different QPAs across bases are specified in the QPA record group and override this field.
29-31	iii	Minimum quantity. Minimum number of this LRU required for the aircraft to be mission capable (i.e., the QPA from the previous field less the number that may be broken without impairing the aircraft's capability). Different minimum quantities across bases are specified in the QPA record group and override this field.
32	i	Demands per sortie indicator. Set to 1 if the demand rates (described below) are per sortie; set to 0 or blank if the demand rates are per flying hour. (Mode of failure for indentured SRUs is the same as that of the LRU.)
33	i	NRTS/condemnation/failed SRU policy. Indicates when the decision to NRTS or condemn the LRU is made, and when failed SRUs indentured to the LRU are discovered. Set to 1 if these occur before attempting repair; set to 0 or blank if the LRU first enters repair.
34-40	XXxxxxx	Onshore demand rate (peacetime). At onshore bases, the fraction of the LRU that breaks per flying hour/sortie (determined in column 32) in peacetime. (The wartime demand rate is determined by multiplying this field by the onshore demand rate multiplier in the VTM record group.)
41-47	XXxxxxx	Offshore demand rate (peacetime). At offshore bases, the fraction of the LRU that breaks per flying hour/sortie (determined in column 32) in peacetime. (The wartime demand rate is determined by multiplying this field by the offshore demand rate multiplier in the VTM record group.)
48-52	XXXxx	Lone base repair time (in days). Repair time at bases not served by a cirf.

Columns	Format	
54-57	XXxx	Lone base NRTS rate. Fraction of removals at bases not served by a cirf that are declared NRTS.
59-62	XXxx	Lone base condemnation rate. Fraction of removals at bases not served by a cirf that are declared condemned.
63-67	XXXxx	Cirf-served base repair time (in days). Repair time at bases served by a cirf.
69-72	XXxx	Cirf-served base NRTS rate. Fraction of removals at bases served by a cirf that are declared NRTS.
74-77	XXxx	Cirf-served base condemnation rate. Fraction of removals at bases served by a cirf that are declared condemned.

Second record of pair:

Columns	Format	
1-16	a16	LRU name. The name of the LRU. Must match LRU name given on first record of pair.
18-21	XXxx	Cirf repair time (in days). Repair time at cirfs.
23-26	XXxx	Cirf NRTS rate. Fraction of removals at cirfs that are declared NRTS.
28-31	XXxx	Cirf condemnation rate. Fraction of removals at cirfs that are declared condemned.
32-36	XXXxx	Depot repair time (in days). Repair time at the depot.
38-41	XXXx	Depot repair limit. Maximum number of the LRU that can be repaired at the depot each day during wartime. (Applies only to LRUs not assigned to constrained repair.) A value of 0 means no limit, while a value of -1 means no depot repair.

Columns	Format	
43-46	XXxx	Depot condemnation rate. Fraction of removals at the depot that are declared condemned.
47-51	XXXXx	Peacetime resupply time (in days). The expected time for the highest echelon repairing the LRU to procure a replacement during peacetime.
52-56	XXXXx	Wartime resupply time (in days). The expected time for the highest echelon repairing the LRU to procure a replacement during wartime.
58-65	XXXXXXXX	Cost. Unit cost of the LRU.
67-73	aaaaaaa	Work unit code. This field may contain any string of 7 characters.
75	i	No cannibalization indicator. Set to 1 if the LRU cannot be cannibalized; set blank or 0 if it can. This switch affects only the results in the performance report (selected with option 11 in the OPT record group) for the partial cannibalization assumption.

Mission Requirements

Header record: MESL

Restriction: Must follow the BASE record group.

General description:
These records specify which missions the aircraft at each base fly during peacetime and on each day of war. Aircraft at bases for which MESL records are not given are assumed to fly all missions.

```
Columns
         1         2         3         4         5
12345678901234567890123456789012345678901234567890

aaaa aaaaaiiiaaaaaiiiaaaaaiiiaaaaaiiiaaaaaiiiaaaaa
|    |    |  |    |  |    |  |    |  |    |  |
|    |    |  |    |  |    |  |    |  |    |  Sixth set of missions
|    |    |  |    |  |    |  |    |  |    Day sixth set starts
|    |    |  |    |  |    |  |    |  Fifth set of missions
|    |    |  |    |  |    |  |    Day fifth set starts
|    |    |  |    |  |    |  Fourth set of missions
|    |    |  |    |  |    Day fourth set starts
|    |    |  |    |  Third set of missions
|    |    |  |    Day third set starts
|    |    |  Second set of missions
|    |    Day second set starts
|    First set of missions
Base name
```

Detailed description:
The maximum number of mission sets is limited by the value of DMCHANGE, selected when the model was compiled. Extra fields for more than six sets follow the format of the previous fields.

The number of missions within a set is limited by the value of DMMISSNS, also selected when the model was compiled (we assume a value of 5). Should this parameter not equal 5, the "set of missions" fields will alter in length.

Each column in the set of missions fields represents a mission. The value of these columns must be either 1 (the base flies that particular mission) or 0 (the base does not fly that mission). Each day a set starts must be greater than the previous day. Whether or not an LRU is essential for a mission is determined in the VTM record group.

For example, a record containing

BAS1 10011 311111

indicates that at base BAS1, aircraft fly the first, fourth and fifth missions in peacetime and through the end of day 2. On day 3 and for the rest of the scenario, BAS1's aircraft fly all five missions.

Columns	Format	
1-4	aaaa	Base name. The name of the base for which missions are specified. Must be named in the BASE record group. Enter at most one record per base.
6-10	aaaaa	First set of missions.
11-13	iii	Day second set starts.
14-18	aaaaa	Second set of missions.
19-21	iii	Day third set starts.
22-26	aaaaa	Third set of missions.
27-29	iii	Day fourth set starts.
30-34	aaaaa	Fourth set of missions.
35-37	iii	Day fifth set starts.
38-42	aaaaa	Fifth set of missions.
43-45	iii	Day sixth set starts.
46-50	aaaaa	Sixth set of missions.

Option Selection

Header record: OPT

General description:
These records define the options that generate Dyna-METRIC's reports and specify the parameters that further define the options.

```
Columns
         1         2
12345678901234567890

    iiiiiXXXxx
    |  |  |
    |  |  Second parameter
    |  First parameter
    Option number
```

Detailed description:
Options and their associated parameter values are described below.

Columns	Format	
5-7	iii	Option number.
8-10	iii	First parameter.
11-15	XXXxx	Second parameter.

Option Description

1 Wartime depot daily repair limit.
The maximum number of each component that can finish depot repair on each day of war is a percentage of the peacetime daily completions. For example, 200 means that twice as much can be repaired per day in wartime as in peacetime; 50 means that half as much can be repaired; 0 means no wartime depot repair. The first parameter contains the percentage (0 to 999). The second parameter is not used.

CAUTION: Option 1 overrides the depot repair limit entered in the LRU, SRU and SSRU record groups and echoed in Table 2. Instead, use option 18 to report the values that were used. Using option 1 without a peacetime flying program implies wartime repair limits of zero, effectively cutting off depot repair. Option 1 must be run with Option 11 (performance report and/or Option 18 (daily demands report).

Option Description

2 Depot LRU stockage.
For each LRU, add enough depot stock so as not to exceed a target number of degraded theater aircraft with a specified confidence level. The LRU may degrade the aircraft because of depot repair or retrograde transportation delays. The first parameter is the percent of theater aircraft that may be degraded (0 to 100). The second parameter is the confidence level (between 0 and 1).

CAUTION: This option reports only the cost of the stock; also select option 9 to report the stock levels. Do not request 0% degraded aircraft with 1.00 confidence.

3 Base LRU stockage for individual LRU goal.
For each LRU, add enough stock at each base so that an individual LRU does not exceed a target number of degraded aircraft with a specified confidence level. The LRU may degrade the aircraft because of base repair or serviceable transportation delays. The first parameter is the percent of aircraft at each base that may be degraded (0 to 100). The second parameter is the confidence level (between 0 and 1).

CAUTION: This option reports only the cost of the stock; also select option 9 to report the stock levels. Stock is not bought beyond the (non-zero) cost limits selected with option 4. Do not request 0% degraded aircraft with 1.00 confidence.

4 Base LRU stockage for overall aircraft goal.
Across all LRUs, add enough stock at each base so that taken together, the LRUs do not exceed a target number of degraded aircraft at a specified confidence level. Any LRU may degrade the aircraft because of base repair or serviceable transportation delays. The first parameter is the percent of aircraft at each base that may be degraded (0 to 100). The second parameter is the confidence level (between 0 and 1).

CAUTION: This option reports only the cost of the stock; also select option 9 to report the stock levels. Do not request 0% degraded aircraft with 1.00 confidence.

The difference between options 3 and 4 is best described with an example. Let's specify that no aircraft be degraded with .90 confidence for 10 LRUs. Option 3 selects stock for each LRU so that the probability of not having a stock shortfall for that part is 90%. With option 4, that probability is the product of each LRU's probability of not having a stock shortfall ($.90^{10}$ or 35%). Option 4 selects additional stock until the target confidence level is attained.

Option Description

5 Cirf LRU stockage.
For each LRU, add enough cirf stock so as not to exceed a target number of degraded aircraft at the cirf's bases, with a specified confidence level. The LRU may degrade the aircraft because of cirf repair or retrograde transportation delays. The first parameter is the percent of aircraft at the cirf's bases that may be degraded (0 to 100). The second parameter is the confidence level (between 0 and 1).

CAUTION: This option reports only the cost of the stock; also select option 9 to report the stock levels. Do not request 0% degraded aircraft with 1.00 confidence.

6 SRU and subSRU stockage.
For each SRU and subSRU, add enough base, cirf and depot stock so as not to exceed a target number of degraded aircraft with a specified confidence. The SRU or subSRU may degrade the aircraft because of component repair delays. The first parameter is the percent of aircraft that may be degraded (0 to 100). The second parameter is the confidence level (between 0 and 1).

CAUTION: This option reports only the cost of the stock; also select option 9 to report the stock levels. Do not request 0% degraded aircraft with 1.00 confidence.

7 Depot workload report.
Report the minimum depot repairs required to satisfy each arriving requisition. If requisitions cannot be met by depot repairs alone, additional depot stock is purchased with a confidence of satisfying the requisitions (specified in the second parameter as a value between 0 and 1). The first parameter is not used. This option uses nonstandard files, described in App. A.

CAUTION: Do not specify time-varying stock levels.

8 Problem LRUs list.
Report those LRUs that individually have a high confidence of grounding more than a target number of aircraft (the target and confidence level are specified in option 11). Only the first parameter, the maximum number of LRUs to be reported, is used.

9 Stock purchased report.
Report the stock levels as purchased by options 2, 3, 4, 6, 7 and 17. PIPE reports LRU stock levels at depots and cirfs as well as SRU and subSRU stock levels. REPORT gives LRU stock levels at bases.

Option Description

10 Read in peacetime pipelines.
Initialize the peacetime pipelines using previously saved or measured data. The first parameter is the option used to construct the input file of pipeline values: either 16 (base this run on values saved by a previous run) or 22 (base this run on user-supplied data). The second parameter is not used. Option 10 uses nonstandard files, described in App. A.

CAUTION: Do not specify time-varying stock levels.

11 Performance Report (based on input and previously purchased stock).
Assess the capability of aircraft at each time of analysis based on input stock and stock bought before that time of analysis. Reported for both full and partial cannibalization assumptions are: the predicted number of degraded aircraft and achieved sorties, the probability of achieving a specified level of aircraft, the number of FMC aircraft at a minimum requested confidence level, and the probability of having less than the target level of aircraft degraded due to component support. The first parameter is the percent of aircraft that may be degraded (0 to 100). The second parameter is the confidence level (between 0 and 1).

12 Performance Report (based on input and currently purchased stock).
The same as option 11, only the assessment is based on input stock and stock bought through that time of analysis. The first parameter is the percent of aircraft that may be degraded (0 to 100). The second parameter is the confidence level (between 0 and 1).

13 Suppress echo.
Do not echo input data except for the summary of options selected.

14 Suppress parts echo.
Do not echo LRU, SRU or subSRU descriptive data.

15 Detailed parts disposition.
Write a detailed parts disposition report at each time of analysis.

16 Save data for restart.
Save the pipeline status at the last time of analysis for use with option 10 in later runs. If a subsequent run is to be started with user-supplied data, select option 22 instead. This option uses nonstandard files, described in App. A.

17 Buy base LRU stock to a no cannibalization goal.
Economically buy LRU stock at each base so as to meet a target percentage of fully mission capable aircraft, assuming LRUs cannot be cannibalized. The first parameter is not used. The second parameter is the percent of aircraft that must be operational (a value between 0 and 1).

Option Description

18 Daily demands report.
Produce a report detailing demands for repair and supply at each base, cirf and depot at each time of analysis.

19 Not available.

20 Achievable sorties on PMC and FMC aircraft.
Recompute the requested sorties affecting an LRU if, for any base at any time of analysis, the backorders for that LRU preclude successfully achieving the requested sortie rate with a specified confidence. The sortie rate is successively multiplied by the percent in the first parameter (0 to 100) until the confidence level in the second parameter (between 0 and 1) is achieved. This option uses nonstandard files, described in App. A.

CAUTION: Do not request 100% for the first parameter or 1.00 for the second.

Suppose we request 100 sorties per day with .90 confidence on day 1, but the actual confidence of achieving them is only .50. Option 20 scales down the requested sortie rate according to the first parameter (say, 95). Thus, the sortie rate becomes 95% of 100, or 95. Suppose the confidence of achieving 95 sorties is still lower than .90. The sortie rate is again adjusted to 95% of its current level, or 90.25, and continues in that fashion until the requested confidence level is achieved.

21 LRU-dependent flying hour programs.
Read LRU-dependent flying hour programs from a separate file. The first parameter is the time period associated with the flying hours (e.g., 30 means monthly, 1 means daily). The second parameter is not used. See App. A for format of the input data and nonstandard files used by this option.

22 Build restart file based on user-supplied pipeline.
With user-supplied data, build a file to initialize pipelines in later runs with option 10. If a subsequent run is to be started with data saved from a Dyna-METRIC run, select option 16 instead. See Appendix A for format of the input data and nonstandard files used by this option.

23 Add problem SRUs and subSRUs to the problem LRUs report.
Include the status of the worst SRUs per problem LRU and the worst subSRUs per problem SRU in option 8's output. The first parameter is the maximum number of subcomponents per component per location to be reported. The second parameter is the fraction of the parent component's total pipeline due to AWP only that must be exceeded (for some location) for the subcomponents to be reported.

Quantity per Aircraft

Header record: QPA

Restriction: Must follow the LRU, SRU and SSRU record groups.

General description:
These records specify the quantity per aircraft of each component for each base. Bases for which QPA records are not entered use the quantity per aircraft data specified in the LRU, SRU and SSRU record groups.

```
Columns
         1         2         3         4
1234567890123456789012345678901234567

aaaaaaaaaaaaaaa aaaaiiiiiiaaaaiiiiiiaaaaiiiiii
|                |   |  |  |   |  |  |   |  |
|                |   |  |  |   |  |  |   |  Third minimum quantity
|                |   |  |  |   |  |  |   Third quantity per aircraft
|                |   |  |  |   |  |  Third base name
|                |   |  |  |   |  Second minimum quantity
|                |   |  |  |   Second quantity per aircraft
|                |   |  |  Second base name
|                |   |  First minimum quantity
|                |   First quantity per aircraft
|                First base name
Part name

         4       5         6         7
         890123456789012345678901234567

         aaaaiiiiiiaaaaiiiiiiaaaaiiiiii
         |   |  |  |   |  |  |   |  |
         |   |  |  |   |  |  |   |  Sixth minimum quantity
         |   |  |  |   |  |  |   Sixth quantity per aircraft
         |   |  |  |   |  |  Sixth base name
         |   |  |  |   |  Fifth minimum quantity
         |   |  |  |   Fifth quantity per aircraft
         |   |  |  Fifth base name
         |   |  Fourth minimum quantity
         |   Fourth quantity per aircraft
         Fourth base name
```

Detailed description:
Minimum quantities (specified for LRUs only) cannot be larger than the quantity per aircraft. Minimum quantities of 0 or blank imply that the LRU is nonessential. Base names that do not match names entered in the BASE record group will generate a warning message. QPAs for up to six bases may be entered on a single record. If more are needed, multiple records are allowed.

For example, a record containing

```
COMPUTER         BAS1  2  1BAS2  2  1
```

indicates that there are two COMPUTERs on the aircraft at both bases, but only one is essential for the aircraft to be mission capable.

Columns	Format	
1-16	a16	Part name. The name of the part for which QPA data are specified. Must be named in the LRU, SRU or subSRU record group. Enter as many records as needed per part.
18-21	aaaa	First base name. Name of the base to which the first QPA applies.
22-24	iii	First quantity per aircraft.
25-27	iii	First minimum quantity (LRUs only).
28-31	aaaa	Second base name.
32-34	iii	Second quantity per aircraft.
35-37	iii	Second minimum quantity (LRUs only).
38-41	aaaa	Third base name.
42-44	iii	Third quantity per aircraft.
45-47	iii	Third minimum quantity (LRUs only).
48-51	aaaa	Fourth base name.
52-54	iii	Fourth quantity per aircraft.
55-57	iii	Fourth minimum quantity (LRUs only).
58-61	aaaa	Fifth base name.
62-64	iii	Fifth quantity per aircraft.
65-67	iii	Fifth minimum quantity (LRUs only).
68-71	aaaa	Sixth base name.
72-74	iii	Sixth quantity per aircraft.
75-77	iii	Sixth minimum quantity (LRUs only).

Sortie Rates

Header record: SRTS

Restriction: Must follow the BASE record group.

General description:
These records specify the average number of daily sorties required per aircraft at each base during peacetime and on each day of war. Aircraft at bases for which SRTS records are not given will fly no sorties.

```
Columns
          1         2         3         4
123456789012345678901234567890123456789012345678

aaaaXXXxiiiiXXXxiiiiXXXxiiiiXXXxiiiiXXXxiiiiXXXx
|   |   |   |   |   |   |   |   |   |   |   |
|   |   |   |   |   |   |   |   |   |   |   Sixth sortie rate
|   |   |   |   |   |   |   |   |   |   Day sixth rate starts
|   |   |   |   |   |   |   |   |   Fifth sortie rate
|   |   |   |   |   |   |   |   Day fifth rate starts
|   |   |   |   |   |   |   Fourth sortie rate
|   |   |   |   |   |   Day fourth rate starts
|   |   |   |   |   Third sortie rate
|   |   |   |   Day third rate starts
|   |   |   Second sortie rate
|   |   Day second rate starts
|   First sortie rate
Base name
```

Detailed description:
Sortie rates may not exceed the base's maximum sortie rate entered in the TURN record group. The number of sortie rates is limited by the value of DMCHANGE, selected when the model was compiled. Extra fields for more than six levels follow the format of previous fields. Each day a sortie rate starts must be greater than the previous day.

For example, a record containing

```
BAS1  1.   8   3.
```

indicates that the aircraft at base BAS1 are required to average one daily sortie in peacetime and through day 7 of the war. Day 8 begins a new daily average of three sorties per aircraft that lasts throughout the scenario.

Columns	Format	
1-4	aaaa	Base name. The name of the base for which sortie requirements are specified. Must be named in the BASE record group. Enter at most one record per base.
5-8	XXXx	First sortie rate.
9-12	iiii	Day second rate starts.
13-16	XXXx	Second sortie rate.
17-20	iiii	Day third rate starts.
21-24	XXXx	Third sortie rate.
25-28	iiii	Day fourth rate starts.
29-32	XXXx	Fourth sortie rate.
33-36	iiii	Day fifth rate starts.
37-40	XXXx	Fifth sortie rate.
41-44	iiii	Day sixth rate starts.
45-48	XXXx	Sixth sortie rate.

SRU Descriptions

Header record: SRU

Restriction: Must follow the LRU record group.

General description:
These records describe the failure, repair and resupply characteristics of each SRU. A pair of these records is required for each SRU.

```
Columns (first record of pair)
         1         2         3         4
1234567890123456789012345678901234567890

aaaaaaaaaaaaaaa aaaa i iiii    iXXxxxxx
|               |    | ||      ||
|               |    | ||      |Onshore demand rate
|               |    | ||      NRTS/condemnation/failed subSRU policy
|               |    | |Quantity per aircraft
|               |    | Cirf reparability switch
|               |    Level of repair
|               Depot name
SRU name
```

```
4         5         6         7
12345678901234567890123456789012345678

XXxxxxxXXXxx XXxx XXxxXXXxx XXxx XXxx
|      |     |    |   |      |    |
|      |     |    |   |      |    Cirf-served base
|      |     |    |   |      |        condemnation rate
|      |     |    |   |      Cirf-served base NRTS rate
|      |     |    |   Cirf-served base repair time
|      |     |    Lone base condemnation rate
|      |     Lone base NRTS rate
|      Lone base repair time
Offshore demand rate
```

Detailed description:
Repair (cycle) time includes time awaiting maintenance and in work.

Condemnation rates are given as a percentage of all SRUs, not as a percentage on non-NRTS SRUs. Thus the sum of the condemnation and NRTS rates for a particular echelon cannot exceed one (100%).

```
Columns (second record of pair)
         1         2         3         4
1234567890123456789012345678901234567890123456

aaaaaaaaaaaaaaaa XXxx XXxx XXxxXXXxx XXXx XXxx
|                |    |    |   |     |    |
|                |    |    |   |     |    Depot condemnation rate
|                |    |    |   |     Depot repair limit
|                |    |    |   Depot repair time
|                |    |    Cirf condemnation rate
|                |    Cirf NRTS rate
|                Cirf repair time
SRU name

           4  5         6         7
           78901234567890123456789012345

           XXXXxXXXXx XXXXXXXX aaaaaaa
           |    |     |        |
           |    |     |        Work unit code
           |    |     Cost
           |    Wartime resupply time
           Peacetime resupply time
```

First record of pair:

Columns Format

1-16	a16	SRU name. The name of the SRU. This may be any string of 16 characters as long as it has neither the name of another LRU, SRU or subSRU nor the first four characters of a keyword (such as "BASE1 SWITCH").
18-21	aaaa	Depot name. The name of the depot that repairs this SRU. Must be named in the DEPT record group. If the SRU is not repaired by a depot, this field must be blank.
23	i	Level of repair. Set to 1 if the SRU can be repaired at a base, cirf or depot. Set to 2 if the SRU can be repaired at a cirf or depot but not a base. Set to 3 if the SRU can only be repaired at a depot.
25	i	Cirf reparability switch. Set to 1 if a cirf can repair the SRU; set to 0 or blank if it cannot. This field is disregarded if the level of repair is 3 (depot only).

Columns	Format	
26-28	iii	Quantity per aircraft. Number of the SRU per aircraft at all bases. Different QPAs across bases are specified in the QPA record group and override this field.
33	i	NRTS/condemnation/failed subSRU policy. Indicates when the decision to NRTS or condemn the SRU is made, and when failed subSRUs indentured to the SRU are discovered. Set to 1 if these occur before attempting repair; set to 0 or blank if the SRU first enters repair.
34-40	XXxxxxx	Onshore demand rate. At onshore bases, the fraction of the SRU that breaks in peacetime per flying hour/sortie (as per the demands per sortie indicator on the parent LRU's record). (The wartime rate is effectively the product of this rate and the onshore demand rate multiplier on the VTM record for the parent LRU.)
41-47	XXxxxxx	Offshore demand rate. At offshore bases, the fraction of the SRU that breaks in peacetime per flying hour/sortie (according to 230the demands per sortie indicator on the parent LRU's record). (The wartime rate is effectively the product of this rate and the offshore demand rate multiplier on the VTM record for the parent LRU.)
48-52	XXXxx	Lone base repair time (in days). Repair cycle time at bases not served by a cirf.
54-57	XXxx	Lone base NRTS rate. Fraction of removals at bases not served by a cirf that are declared NRTS.
59-62	XXxx	Lone base condemnation rate. Fraction of removals at bases not served by a cirf that are declared condemned.
63-67	XXXxx	Cirf-served base repair time (in days). Repair cycle time at bases served by a cirf.
69-72	XXxx	Cirf-served base NRTS rate. Fraction of removals at bases served by a cirf that are declared NRTS.
74-77	XXxx	Cirf-served base condemnation rate. Fraction of removals at bases served by a cirf that are declared condemned.

Second record of pair:

Columns	Format	
1-16	a16	SRU name. The name of the SRU. Must match SRU name given on first record of pair.
18-21	XXxx	Cirf repair time (in days). Repair cycle time at cirfs.
23-26	XXxx	Cirf NRTS rate. Fraction of removals at cirfs that are declared NRTS.
28-31	XXxx	Cirf condemnation rate. Fraction of removals at cirfs that are declared condemned.
32-36	XXXxx	Depot repair time (in days). Repair cycle time at the depot.
38-41	XXXx	Depot repair limit. Maximum number of the SRU that can be repaired at the depot each day during wartime. A value of 0 means no limit, while a value of -1 means no depot repair.
43-46	XXxx	Depot condemnation rate. Fraction of removals at the depot that are declared condemned.
47-51	XXXXx	Peacetime resupply time (in days). The expected time for the highest echelon repairing the SRU to procure a replacement during peacetime.
52-56	XXXXx	Wartime resupply time (in days). The expected time for the highest echelon repairing the SRU to procure a replacement during wartime.
58-65	XXXXXXXX	Cost. Unit cost of the SRU.
67-73	aaaaaaa	Work unit code. This field may contain any string of 7 characters.

SubSRU Descriptions

Header record: SSRU

Restriction: Must follow the LRU and SRU record groups.

General description:
These records describe the failure, repair and resupply characteristics of each subSRU. A pair of these records is required for each subSRU.

```
Columns (first record of pair)
         1         2         3         4
1234567890123456789012345678901234567890

aaaaaaaaaaaaaaaa aaaa i iiii    i
|                |    | ||      |
|                |    | ||      NRTS/condemnation policy
|                |    | |Quantity per aircraft
|                |    | Cirf reparability switch
|                |    Level of repair
|                Depot name
SubSRU name

          4         5         6         7
          1234567890123456789012345678901234567

                 XXXxx XXxx XXxxXXXxx XXxx XXxx
                 |     |    |   |     |    |
                 |     |    |   |     |    Cirf-served base
                 |     |    |   |     |        condemnation rate
                 |     |    |   |     Cirf-served base NRTS rate
                 |     |    |   Cirf-served base repair time
                 |     |    Lone base condemnation rate
                 |     Lone base NRTS rate
                 Lone base repair time
```

Detailed description:
Repair cycle time includes time awaiting maintenance and in work.

Condemnation rates are given as a percentage of all subSRUs, not as a percentage on non-NRTS subSRUs. Thus the sum of the condemnation and NRTS rates for a particular echelon cannot exceed one (100%).

In place of demand rates, subSRUs have replacement fractions in the INDT record group.

```
Columns (second record of pair)
         1         2         3         4
1234567890123456789012345678901234567890123456

aaaaaaaaaaaaaaaa XXxx XXxx XXxxXXXxx XXXx XXxx
|                |    |    |   |     |    |
|                |    |    |   |     |    Depot condemnation rate
|                |    |    |   |     Depot repair limit
|                |    |    |   Depot repair time
|                |    |    Cirf condemnation rate
|                |    Cirf NRTS rate
|                Cirf repair time
SubSRU name

               4  5         6         7
               78901234567890123456789012345

               XXXXxXXXXx XXXXXXXX aaaaaaa
               |    |     |        |
               |    |     |        Work unit code
               |    |     Cost
               |    Wartime resupply time
               Peacetime resupply time
```

First record of pair:

Columns	Format	
1-16	a16	SubSRU name. The name of the subSRU. This may be any string of 16 characters as long as it has neither the name of another LRU, SRU or subSRU nor is the first four characters of a keyword (such as "BASE1 CARD5").
18-21	aaaa	Depot name. The name of the depot that repairs this subSRU. Must be named in the DEPT record group. If the subSRU is not repaired by a depot, this field must be blank.
23	i	Level of repair. Set to 1 if the subSRU can be repaired at a base, cirf or depot. Set to 2 if the subSRU can be repaired at a cirf or depot but not a base. Set to 3 if it can only be repaired at a depot.
25	i	Cirf reparability switch. Set to 1 if a cirf can repair the subSRU; set to 0 or blank if it cannot. This field is disregarded if the level of repair is 3 (depot only).

Columns	Format	
26-28	iii	Quantity per aircraft. Number of the subSRU per aircraft at all bases. Different QPAs across bases are specified in the QPA record group and override this field.
33	i	NRTS/condemnation/failed subSRU policy. Indicates when the decision to NRTS or condemn the subSRU is made. Set to 1 if these occur before attempting repair; set to 0 or blank if the subSRU first enters repair.
48-52	XXXxx	Lone base repair time (in days). Repair cycle time at bases not served by a cirf.
54-57	XXxx	Lone base NRTS rate. Fraction of removals at bases not served by a cirf that are declared NRTS.
59-62	XXxx	Lone base condemnation rate. Fraction of removals at bases not served by a cirf that are declared condemned.
63-67	XXXxx	Cirf-served base repair time (in days). Repair cycle time at bases served by a cirf.
69-72	XXxx	Cirf-served base NRTS rate. Fraction of removals at bases served by a cirf that are declared NRTS.
74-77	XXxx	Cirf-served base condemnation rate. Fraction of removals at bases served by a cirf that are declared condemned.

Second record of pair:

Columns	Format	
1-16	a16	SubSRU name. The name of the subSRU. Must match subSRU name given on first record of pair.
18-21	XXxx	Cirf repair time (in days). Repair cycle time at cirfs.
23-26	XXxx	Cirf NRTS rate. Fraction of removals at cirfs that are declared NRTS.

Columns	Format	
28-31	XXxx	Cirf condemnation rate. Fraction of removals at cirfs that are declared condemned.
32-36	XXXxx	Depot repair time (in days). Repair cycle time at the depot.
38-41	XXXx	Depot repair limit. Maximum number of the subSRU that can be repaired at the depot each day during wartime. A value of 0 means no limit, while a value of -1 means no depot repair.
43-46	XXxx	Depot condemnation rate. Fraction of removals at the depot that are declared condemned.
47-51	XXXXx	Peacetime resupply time (in days). The expected time for the highest echelon repairing the subSRU to procure a replacement during peacetime.
52-56	XXXXx	Wartime resupply time (in days). The expected time for the highest echelon repairing the subSRU to procure a replacement during wartime.
58-65	XXXXXXXX	Cost. Unit cost of the subSRU.
67-73	aaaaaaa	Work unit code. This field may contain any string of 7 characters.

Stock Levels

Header record: STK

Restriction: Must follow all other record groups.

General description:
These records specify each component's level of stock at the depots, cirfs and bases.

```
Columns
          1         2         3         4
1234567890123456789012345678901234567890123456

aaaaaaaaaaaaaaaaiaaaaiiiii aaaaiiiii aaaaiiiii
|               ||    |     |    |     |   |
|               ||    |     |    |     |  Third stock level
|               ||    |     |    |    Third stock location
|               ||    |     |   Second stock level
|               ||    |    Second stock location
|               ||   First stock level
|               |First stock location
|               Replacement switch
Part name

          4         5         6         7    8
          7890123456789012345678901234567890

           aaaaiiiii aaaaiiiii aaaaiiiiiiiiii
           |   |     |   |     |   |    |
           |   |     |   |     |   |   Day levels start
           |   |     |   |     |  Sixth stock level
           |   |     |   |    Sixth stock location
           |   |     |  Fifth stock level
           |   |    Fifth stock location
           |  Fourth stock level
          Fourth stock location
```

Detailed description:
Location names that do not match names entered in the BASE, CIRF or DEPT record groups will generate a warning message. Stock levels for up to six bases may be entered on a single record. If more are needed, multiple records are allowed.

Stock levels may be changed over time by including multiple records for the part and location. The replacement switch indicates whether the new level is to replace, be added to, or subtracted from the current level. The new levels go into effect on the Day Levels Start. Records must appear in ascending order of columns 77-80.

For example, records containing

PART A	BAS1	1	BAS2	2	0
PART B	BAS1	5	BAS2	5	0
PART A	BAS1	7			4

indicate the stock levels of two parts at two bases. One stock level changes during the wartime scenario (day 4), that of PART A at BAS1.

Columns	Format	
1-16	a16	Part name. Name of LRU, SRU or subSRU to which the stock levels apply. Must be named in the LRU, SRU or SSRU record group. Enter as many records as needed per part.
17	i	Replacement switch. Switch indicating whether the stock level replaces (blank or 0), is added to (1), or is subtracted from (2) the stock level previously given for all locations on this record.
18-21	aaaa	First stock location. Name of the location to which first stock level applies.
22-26	iiiii	First stock level. Stock assigned to the first location.
28-31	aaaa	Second stock location.
32-36	iiiii	Second stock level.
38-41	aaaa	Third stock location.
42-46	iiiii	Third stock level.
48-51	aaaa	Fourth stock location.
52-56	iiiii	Fourth stock level.
58-61	aaaa	Fifth stock location.
62-66	iiiii	Fifth stock level.
68-71	aaaa	Sixth stock location.
72-76	iiiii	Sixth stock level.

Columns	Format	
77-80	iiii	Day levels start. Day on which the stock levels on this record go into effect. Each day must be greater than or equal to the previous day on any prior stock record (i.e., all stock records must appear in ascending order based on the value of this field). Leave this field blank if using option 7 or option 10 (i.e., do not deploy stock with these options).

Server Levels

Header record: TBED

Restriction: Must follow the TEST record group to which it applies.

General description:
These records specify how many servers per constrained repair resource are available at each location during peacetime and on each day of war. Locations for which TBED records are not entered are assigned no servers.

```
Columns
         1         2         3         4         5
123456789012345678901234567890123456789012345678901

aaaa XxxxxXXXXXXXXXXiiiiiiiiiiiiiiiiiiiiiiiiii
|    |    |   |  |  |  |  |  |  |  |  |  |  |  |  |
|    |    |   |  |  |  |  |  |  |  |  |  |  |  |  Sixth server level
|    |    |   |  |  |  |  |  |  |  |  |  |  |  Day sixth level starts
|    |    |   |  |  |  |  |  |  |  |  |  |  Fifth server level
|    |    |   |  |  |  |  |  |  |  |  |  Day fifth level starts
|    |    |   |  |  |  |  |  |  |  |  Fourth server level
|    |    |   |  |  |  |  |  |  |  Day fourth level starts
|    |    |   |  |  |  |  |  |  Third server level
|    |    |   |  |  |  |  |  Day third level starts
|    |    |   |  |  |  |  Second server level
|    |    |   |  |  |  Day second level starts
|    |    |   |  |  First server level
|    |    |   |  Resupply cutoff duration
|    |    |   Resupply cutoff
|    |    Resupply time
|    Backorder rate
Location name
```

Detailed description:
The maximum number of server levels is limited by the value of DMCHANGE, selected when the model was compiled. Extra fields for more than six levels follow the format of the previous fields. Each day a new level starts must be greater than the previous day. No server level should exceed DMSTANDS.

Each constrained repair resource is described by its own TEST, TPRT and TBED record groups. Thus, the TBED record group may appear more than once in an input data set. It must follow the TEST record group to which it applies. Servers may be test stands or some other repair resource whose quantity or availability may limit the number of components repaired in a given time period.

Columns	Format	
1-4	aaaa	Location name. The name of the location for which server levels are specified. Must be named in the BASE, CIRF or DEPT record group. Enter at most one record per location.
6-10	Xxxxx	Backorder rate. Expected number of server failures per operating day that cause a server to become partially mission capable (able to repair only a subset of the LRUs assigned to it).
11-14	XXXX	Resupply time. Average time at this location to obtain a replacement module for the most recently named repair resource because of a failure that made a server partially mission capable.
15-17	XXX	Resupply cutoff. Day that resupply of modules for the most recently named repair resource at this location is cut off.
18-20	XXX	Resupply cutoff duration. Number of days resupply of modules for the most recently named repair resource at this location is cut off.
21-23	iii	First server level.
24-26	iii	Day second level starts.
27-29	iii	Second server level.
30-32	iii	Day third level starts.
33-35	iii	Third server level.
36-38	iii	Day fourth level starts.
39-41	iii	Fourth server level.
42-44	iii	Day fifth level starts.
45-47	iii	Fifth server level.
48-50	iii	Day sixth level starts.
51-53	iii	Sixth server level.

Constrained Repair Availability

Header record: TEST

Restriction: Must follow all record groups except STK.

General description:

These records describe the availability of different types of constrained repair, such as test equipment, skilled personnel, and equipment disassembly fixtures.

```
Columns
         1         2         3         4
1234567890123456789012345678901234567890

aaaa           XXxxxXXxxxXXxxxXXxxxXXxxx
|              |    |    |    |    |
|              |    |    |    |    Five server availability
|              |    |    |    Four server availability
|              |    |    Three server availability
|              |    Two server availability
|              One server availability
Repair resource name
```

Detailed description:

The number of server availabilities is limited by the value of DMCHANGE, selected when the model was compiled. Extra fields for more than five (DMCHANGE - 1) availabilities follow the format of the previous fields.

Each constrained repair resource is described by its own TEST record group (the header record and one data record), TBED record group (the header and a record for each location's server levels), and TPRT record group (the header and a list of LRUs the repair resource can test). Thus, the TEST record group may appear more than once in an input data set.

Columns	Format	
1-4	aaaa	Repair resource name. The name of the repair resource. This may be any string of four characters as long as it is neither a keyword (such as "TEST") nor the name of another repair resource.
16-20	[illegible]	One server availability. Fraction of the time a server is up if only one server is stationed at a location.

Columns	Format	
21-25	XXxxx	Two server availability. Fraction of the time each server is up if two servers are stationed at a location.
26-30	XXxxx	Three server availability. Fraction of the time each server is up if three servers are stationed at a location.
31-35	XXxxx	Four server availability. Fraction of the time each server is up if four servers are stationed at a location.
36-40	XXxxx	Five server availability. Fraction of the time each server is up if five servers are stationed at a location.

Administrative Data

Header record: none

Restriction: Must precede all other record groups.

General description:
These records provide general information about the input data, including a heading, labels for each mission, times of analysis, and a switch for setting exponential or deterministic distributions for repair/transportation times. Also given are base, cirf and depot administrative times.

First record:

```
Columns
         1         2     6         7         8
12345678901234567890.....01234567890123456789 0

aaaaaaaaaaaaaaaaaaaa.....aaaaaaaaaaaaaaaaaaaaa
|
Heading
```

Columns Format

1-80	20a4	Heading. The heading appearing at the top of the output reports. May be any string of up to 80 characters.

Second record:

```
Columns
         1         2         3         4         5
12345678901234567890123456789012345678901234567890

iiXXXxxXXXxxXXXxx  VERSION 4.4  aaaaaaaaaaaaaaa
|||    |    |      |            |  |  |  |  |
|||    |    |      |            |  |  |  |  Fifth mission label
|||    |    |      |            |  |  |  Fourth mission label
|||    |    |      |            |  |  Third mission label
|||    |    |      |            |  Second mission label
|||    |    |      |            First mission label
|||    |    |      Data set version
|||    |    Depot administrative time
|||    Cirf administrative time
||Base administrative time
|Exponential repair switch
Cutoff direction switch
```

Detailed description, second record:

The number of missions is limited by the value of DMMISSNS, selected when the model was compiled. Extra fields for more than five missions follow the format of the previous fields. Mission labels are used in output reports. If not specified, labels default to "M1", "M2", etc.

Columns	Format	
1	i	Cutoff direction switch Set to 1 if the cutoff parameters in the BASE and TRNS record groups apply to both forward and retrograde transportation. Set to 0 or blank if they apply to forward transportation only.
2	i	Exponential repair switch. Set to 1 if transportation and repair delays have an exponential distribution. Set to 0 or blank if they have a deterministic distribution.
3-7	XXXxx	Base administrative time. The deterministic delay (in days) experienced by LRUs removed at the flight line prior to entering base level repair.
8-12	XXXxx	Cirf administrative time. The deterministic delay (in days) experienced by LRUs and SRUs that have been NRTSed by the base, after arrival at the cirf and prior to entering cirf level repair.
13-17	XXXxx	Depot administrative time. The deterministic delay (in days) experienced by LRUs and SRUs that have been NRTSed to the depot from bases and cirfs, after arrival at the depot and prior to entering depot level repair.
20-30	11a	Data set version. Must contain "VERSION 4.4".
33-35	aaa	First mission label.
36-38	aaa	Second mission label.
39-41	aaa	Third mission label.
42-44	aaa	Fourth mission label.
45-47	aaa	Fifth mission label.

Third record:

```
Columns
         1         2
12345678901234567890

iiiiiiiiiiiiiiiiiiii
|   |   |   |   |
|   |   |   |   Fifth time of analysis
|   |   |   Fourth time of analysis
|   |   Third time of analysis
|   Second time of analysis
First time of analysis
```

Detailed description, third record:

Times of analysis are ~he days for which output reports are requested. The number of analysis times is limited by the value of DMANALYS, selected when the model was compiled. Extra fields for more than five times of analysis follow the format of the previous fields.

No two times of analysis should be the same. All times to the right of and including a zero time of analysis are ignored. However, times do not need to be in numerical order.

Columns	Format	
1-4	iiii	First time of analysis.
5-8	iiii	Second time of analysis.
9-12	iiii	Third time of analysis.
13-16	iiii	Fourth time of analysis.
17-20	iiii	Fifth time of analysis.

LRU Repair Assignments

Header record: TPRT

Restriction: Must follow the TEST and TBED record groups to which it applies.

General description:
These records specify which LRUs are assigned to the constrained repair resource named in the immediately preceding TEST record group.

```
Columns
        1
1234567890123456

aaaaaaaaaaaaaaaa
|
LRU name
```

Detailed description:
Each constrained repair resource is described by its own TEST, TBED and TPRT record groups. Thus, there may be more than one TPRT record group in an input data set. It must follow the TEST and TBED record groups to which it applies.

Columns	Format	
1-16	a16	LRU name. The name of the LRU assigned to the most recently named repair resource. Must be named in the LRU record group. Enter a record for each LRU that applies to the repair resource.

Depot Transportation

Header record: TRNS

Restriction: Must follow the DEPT and BASE record groups.

General description:
These records describe transportation resources connecting bases and cirfs with depots. If a record is not entered for some base-depot or cirf-depot pair, transportation between the two is assumed to be instantaneous and never cut off.

```
Columns
         1         2         3         4
12345678901234567890123456789012345678901

aaaa aaaa XXxxx XXxxx i XXXxx XXXxx XXXxx
|    |    |     |     | |     |     |
|    |    |     |     | |     |     Transportation cutoff duration
|    |    |     |     | |     Transportation cutoff
|    |    |     |     | Transportation start
|    |    |     |     Transportation availability switch
|    |    |     Transportation time from depot
|    |    Transportation time to depot
|    Depot name
Base or cirf name
```

Detailed description:
Transportation and cutoff parameters also apply to the retrograde (base/cirf-to-depot) pipeline if the cutoff direction switch in the TOP record group is set.

Columns	Format	
1-4	aaaa	Base or cirf name. The name of the base or cirf. Must be named in the BASE or CIRF record group.
6-9	aaaa	Depot name. The name of the depot. Must be named in the DEPT record group. Enter at most one record per base-depot or cirf-depot combination.
11-15	XXXxx	Transportation time to depot. Base/cirf-to-depot transportation time (in days).
17-21	XXXxx	Transportation time from depot. Depot-to-base/cirf transportation time (in days).

Columns	Format	
23	i	Transportation availability switch. Set to 1 if parts ordered from the depot in peacetime continue to arrive at the base/cirf prior to the transportation start time. Set to 0 or blank if they do not.
25-29	XXXxx	Transportation start. Day transportation from the depot first becomes available.
31-35	XXXxx	Transportation cutoff. Day transportation from the depot is cut off. The peacetime resupply pipeline from the depot is also cut off, even if transportation availability is set.
37-41	XXXxx	Transportation cutoff duration. Number of days transportation is cut off from the depot.

Maximum Sortie Rates

Header record: TURN

Restriction: Must follow the BASE record group.

General description:
These records specify, at each base, the maximum number of sorties a mission capable aircraft can fly per day during peacetime and on each day of war. Aircraft at bases for which TURN records are not given will fly no sorties.

```
Columns
         1         2         3         4
123456789012345678901234567890123456789012345678

aaaaXXXXiiiiXXXXiiiiXXXXiiiiXXXXiiiiXXXXiiiiXXXX
|   |   |   |   |   |   |   |   |   |   |   |
|   |   |   |   |   |   |   |   |   |   |   Sixth maximum sortie rate
|   |   |   |   |   |   |   |   |   |   Day sixth rate starts
|   |   |   |   |   |   |   |   |   Fifth maximum sortie rate
|   |   |   |   |   |   |   |   Day fifth rate starts
|   |   |   |   |   |   |   Fourth maximum sortie rate
|   |   |   |   |   |   Day fourth rate starts
|   |   |   |   |   Third maximum sortie rate
|   |   |   |   Day third rate starts
|   |   |   Second maximum sortie rate
|   |   Day second rate starts
|   First maximum sortie rate
Base name
```

Detailed description:
The number of sortie rates is limited by the value of DMCHANGE, selected when the model was compiled. Extra fields for more than six rates follow the format of the previous fields. Each day a new rate starts must be greater than the previous day.

For example, a record containing

```
BAS1  1.   8   3.
```

indicates that at base BAS1, a single aircraft can fly at most one sortie per day in peacetime and through day 7 of the war. Beginning on day 8 (and for the rest of the scenario), a single aircraft can fly up to 3 sorties per day.

Columns	Format	
1-4	aaaa	Base name. The name of the base for which the maximum sortie rates are specified. Must be named in the BASE record group. Enter at most one record per base.
5-8	XXXx	First maximum sortie rate.
9-12	iiii	Day second rate starts.
13-16	XXXx	Second maximum sortie rate.
17-20	iiii	Day third rate starts.
21-24	XXXx	Third maximum sortie rate.
25-28	iiii	Day fourth rate starts.
29-32	XXXx	Fourth maximum sortie rate.
33-36	iiii	Day fifth rate starts.
37-40	XXXx	Fifth maximum sortie rate.
41-44	iiii	Day sixth rate starts.
45-48	XXXx	Sixth maximum sortie rate.

Variance to Mean Data

Header record: VTM

Restriction: Must follow the LRU record group.

General description:
These records specify an LRU's maintenance type and give its wartime demand rate multipliers, pipeline variance to mean ratio, and the probability that a partially mission capable repair resource is able to repair it. LRUs for which VTM records are not given are assumed:

- maintenance type 0 ("RR" if ILM record group not used)
- essential for all missions
- reparable on a partially mission capable repair resource
- to have wartime demand rates equal to peacetime demand rates
- to have pipelines with Poisson distributions.

```
Columns
          1         2         3         4         5
12345678901234567890123456789012345678901234567890123456789012345678

aaaaaaaaaaaaaaaa i XXxx XXxx XXxx XXxx aaaaa XXxxxx XXxxxx
|                | |    |    |    |    ||||| |      |
|                | |    |    |    |    ||||| |      Offshore sustained
|                | |    |    |    |    ||||| |         demand rate
|                | |    |    |    |    ||||| Onshore sustained
|                | |    |    |    |    |||||        demand rate
|                | |    |    |    |    ||||5th mission essentiality
|                | |    |    |    |    |||4th mission essentiality
|                | |    |    |    |    ||3rd mission essentiality
|                | |    |    |    |    |2nd mission essentiality
|                | |    |    |    |    1st mission essentiality
|                | |    |    |    Partial reparability
|                | |    |    Variance to mean ratio
|                | |    Offshore demand rate multiplier
|                | Onshore demand rate multiplier
|                Maintenance type specifier
LRU name
```

Detailed description:
The number of missions is limited by the value of DMMISSNS, selected when the model was compiled (we assume a value of 5). Should this parameter not equal 5, this format may be invalid.

Columns	Format	
1-16	a16	LRU name. The name of the LRU for which variance to mean data are specified. Must be named in the LRU record group. Enter at most one record per LRU.
18	i	Maintenance type specifier. An integer from 0 to 8 corresponding to the LRU's maintenance type defined on an ILM record. If no ILM records, a value of 0 assigns the LRU to default type RR and a value of 1 assigns it to type RRR.
20-23	XXxx	Onshore demand rate multiplier. A number multiplied by the peacetime onshore demand rate (in the LRU record group) to obtain the wartime onshore demand rate.
25-28	XXxx	Offshore demand rate multiplier. A number multiplied by the peacetime offshore demand rate (in the LRU record group) to obtain the wartime offshore demand rate.
30-33	XXxx	Variance to mean ratio. The variance to mean ratio of the LRU's pipelines. A value between 0 and 1 implies pipelines with a binomial distribution; a value of 1 implies pipelines with a Poisson distribution; a value greater than 1 implies pipelines with a negative binomial distribution.
35-38	XXxx	Partial reparability. Fraction of the time that a partially mission capable repair resource is able to repair the LRU. It must be a number between 0 and 1. This field is disregarded if the LRU is not assigned to a repair resource.
40	a	First mission essentiality. Set to 1 if the LRU is essential for the first mission. Set to blank or 0 if it is not essential.
41	a	Second mission essentiality. Set to 1 if the LRU is essential for the second mission. Set to blank or 0 if it is not essential.
42	a	Third mission essentiality. Set to 1 if the LRU is essential for the third mission. Set to blank or 0 if it is not essential.

Columns	Format	
43	a	Fourth mission essentiality. Set to 1 if the LRU is essential for the fourth mission. Set to blank or 0 if it is not essential.
44	a	Fifth mission essentiality. Set to 1 if the LRU is essential for the fifth mission. Set to blank or 0 if it is not essential.
46-51	XXxxxx	Onshore sustained demand rate. At onshore bases, the demand rate beginning on the day set in the sustained demand start time in the BASE record group. If that start time is blank or 0, the wartime demand rates remain in effect for the entire wartime scenario.
53-58	XXxxxx	Offshore sustained demand rate. At offshore bases, the demand rate beginning on the day set in the sustained demand start time in the BASE record group. If that start time is blank or 0, the wartime demand rates remain in effect for the entire wartime scenario.

Appendix F

MODELING TIPS TO AVOID PITFALLS

To model a cirf.

- add a CIRF record.
- second TOP record: set cirf administrative time.
- BASE record: set various cirf-associated fields.
- TRNS record (if also modeling a depot): specify the transportation link between the cirf and depot.
- OPT record (for requirements computation): select option 5 for cirf LRU stockage.
- LRU/SRU/SSRU records: set repair times, NRTS rates, and condemnation rates for cirfs and cirf-served bases. Select the proper setting for the cirf repairability switch--default is for components NOT to go to the cirf. Check level of repair indicator.
- INDT records: set replacement factor for subSRUs.

To model a depot.

- add a DEPT record.
- second TOP record: set depot administrative time.
- TRNS record: specify the transportation links between the depot and its bases and cirfs.
- OPT record (for requirements computation): select option 2 or 7 for depot LRU stockage.
- LRU/SRU/SSRU records: set the depot name, depot repair time and condemnation rate, and (optionally) depot repair limit. Check the level of repair indicator.
- INDT records: set replacement factor for SRUs and subSRUs.

Constrained repair.

- LRU records: use on-stand test times instead of repair cycle times.
- include TEST, TBED and TPRT records.
- VTM records: specify the probability that PMC stands can repair a given LRU.

VTM records.

Default values differ between omitted VTM records and blank VTM records. Make sure fields are set appropriately.

	Defaults if record is omitted for LRU	Defaults if record is included but left blank
mission essential to:	all missions	no missions
maintenance type:	RR	RR
reparable on PMC stands:	yes	no
pipeline distribution:	Poisson	deterministic
wartime demand rate =:	peacetime	zero

Mission essentiality.

- second TOP record: mission names.
- MESL records: indicate the missions each base flies and when.
- VTM records: specify the missions for which an LRU is essential.

Maintenance type specification.

- VTM records: if 0 (RR) or 1 (RRR).
- ILM records: for other values.

Appendix G

INDEX TO DATA FIELDS

	Record Group
Administrative	
repair and transportation switch	TOP
times of analysis	TOP
Base	BASE
administrative delay	TOP
cannibalization: LRUs	LRU
SRUs and subSRUs	BASE
cirf assignment, cutoff, transportation	BASE
components	
NRTS & condemnation rates	LRU/SRU/SSRU
repair start: LRUs	BASE or ILM
SRUs and subSRUs	BASE
repair/test time	LRU/SRU/SSRU
reparable arrival time	BASE
replacement fraction: subSRUs	INDT
depot assignment (see LRU)	
depot cutoff, transportation	TRNS
identical count	BASE
mission requirements	MESL
offshore/onshore identifier	BASE
resupply time, cutoff	BASE
stock buy	OPT
stock levels	STK
Cannibalization	
LRUs	LRU
SRUs and subSRUs	BASE, CIRF and DEPT
stock buys for LRUs	OPT
Cirf	CIRF
administrative delay	TOP
base assignment, cutoff, transportation	BASE
cannibalization: LRUs	LRU
SRUs and subSRUs	CIRF
components	
cirf reparable	LRU/SRU/SSRU
NRTS & condemnation rates	LRU/SRU/SSRU
repair start: LRUs	CIRF or ILM
SRUs and subSRUs	CIRF
repair/test time	LRU/SRU/SSRU
reparable arrival time	CIRF
replacement fraction: subSRUs	INDT
depot assignment (see LRU)	
depot cutoff and transportation	TRNS
resupply time, cutoff	CIRF
stock buy	OPT
stock levels	STK

Constrained repair	
assigned LRUs	TPRT
availability	TEST
backorder rate	TBED
partially mission capable	VTM
resupply	TBED
server levels	TBED
Cutoff	
base resupply	BASE
cirf-to-base	BASE
cirf resupply	CIRF
contrained repair resupply	TBED
depot resupply	DEPT
depot-to-base or -cirf	TRNS
direction	TOP
supplier to base/cirf/depot	BASE/CIRF/DEPT
Demand rate	
multiplier for LRUs	VTM
offshore and offshore LRUs/SRUs	LRU/SRU
subSRUs	INDT
sustained, for LRUs	VTM
sustained start time	BASE
Deployment	
aircraft	ACFT
base repair	BASE, ILM, or TBED
cirf repair	CIRF, ILM, or TBED
constrained repair	TBED
depot repair	DEPT, ILM, or TBED
stock	STK
Depot	DEPT
administrative delay	TOP
base cutoff, transportation	TRNS
cannibalization: LRUs	LRU
SRUs and subSRUs	DEPT
cirf cutoff, transportation	TRNS
components	
assigned to depot	LRU/SRU/SSRU
NRTS & condemnation rates	LRU/SRU/SSRU
repair limit	LRU/SRU/SSRU
repair start: LRUs	DEPT or ILM
SRUs and subSRUs	DEPT
repair/test time	LRU/SRU/SSRU
reparable arrival time	DEPT
replacement fraction: SRUs and subSRUs	INDT
repair limit (computed)	OPT
resupply time, cutoff	DEPT
stock buy	OPT
stock levels	STK
workload report	OPT

Flying program	
aircraft levels	ACFT
attrition rates	ATTR
flying hours per sortie	FLHR
maximum sorties per aircraft	TURN
mission requirements	MESL
sorties per aircraft	SRTS
turn rate	TURN
LRU	LRU and VTM
application fractions	APPL
cannibalization	LRU
cirf reparability	LRU
constrained repair assignment	TPRT
cost	LRU
demand rate	LRU
demand rate multiplier	VTM
demands per sortie	LRU
depot assignment	LRU
level of repair	LRU
maintenance type	VTM
mission essentiality	VTM
NRTS and condemnation rates	LRU
NRTS decision (AWP concurrent/after repair)	LRU
partial reparability	VTM
pipeline probability distribution	VTM
problem report	OPT
quantity per aircraft	LRU or QPA
repair limit, depot	LRU or OPT
repair start time	BASE/CIRF/DEPT or ILM
repair/test time	LRU
resupply time	LRU
SRUs indentured	INDT
stock buys	OPT
stock levels	STK
sustained demand rates	VTM
variance to mean ratio	VTM
work unit code	LRU
Maintenance	
base deployment	BASE or ILM
cirf deployment	CIRF or ILM
depot deployment	DEPT or ILM
LRU type	VTM
type definition	ILM
Mission	
labels	TOP
LRU essentiality	VTM
requirements	MESL
Option selection	OPT

Probability distribution	
pipeline	VTM
repair and distribution	TOP
Quantity per aircraft	
LRUs	LRU or QPA
minimum quantity: LRUs	LRU or QPA
SRUs and subSRUs	SRU/SSRU or QPA
Quantity per higher assembly	INDT
Removal	
of SRU before/after test/repair	LRU
of subSRU before/after test/repair	SRU
Reports	
daily demands	OPT
depot workload for LRUs	OPT
performance	OPT
problem LRUs/SRUs/subSRUs	OPT
requirements computation	OPT
stockage	OPT
Scenario (see Flying program)	
SRU	SRU
assigned LRUs	INDT
cannibalization	BASE/CIRF/DEPT
cirf reparability	SRU
cost	SRU
demand rates	SRU
depot: assignment	SRU
repair limit	SRU
replacement fraction	INDT
level of repair	SRU
NRTS and condemnation rates	SRU
NRTS decision (AWP concurrent/after repair)	SRU
problem report	OPT
quantity: per aircraft	SRU or QPA
per higher assembly	INDT
repair start time	BASE/CIRF/DEPT
repair/test time	SRU
resupply time	SRU
stock buy	OPT
stock levels	STK
subSRUs indentured	INDT
work unit code	SRU
Stock	
computation	OPT
cost of components	LRU/SRU/SSRU
deployment	STK
levels	STK
report	OPT

SubSRU	SSRU
assigned SRUs	INDT
base replacement fraction	INDT
cannibalization	BASE/CIRF/DEPT
cirf: reparability	SSRU
replacement fraction	INDT
cost	SSRU
demand rate	INDT
depot: assignment	SSRU
repair limit	SSRU
replacement fraction	INDT
level of repair	SRU
NRTS and condemnation rates	SSRU
NRTS decision (before/after test)	SSRU
quantity: per aircraft	SSRU or QPA
per higher assembly	INDT
repair start time	BASE/CIRF/DEPT
repair/test time	SSRU
resupply time	SSRU
stock buy	OPT
stock levels	STK
work unit code	SSRU
Supplier	
cutoff from base/cirf/depot	BASE/CIRF/DEPT
resupply time of components	LRU/SRU/SSRU
Test equipment (see constrained repair)	
Transportation	
between base and cirf	BASE
between base and depot	TRNS
between cirf and depot	TRNS
cutoff (see Cutoff)	
probability distribution	TOP

REFERENCES

Berman, M. B., T. F. Lippiatt, and R. Sims, *Impact of Dynamic Wartime Environments on Resource Requirements: An Executive Summary*, The RAND Corporation, unpublished research, May 1978.

Bigelow, J. H., and K. Isaacson, *Models to Assess the Peacetime Material Readiness and Wartime Sustainability of U. S. Air Forces: A Progress Report*, The RAND Corporation, N-1896-MRAL, October 1982.

Carrillo, M. J., "Design and Capabilities of Dyna-TAB," *Proceedings of the USAF Logistics Capability Assessment Symposium* (LOGCAS 82), March 1982.

Cinlar, E., *Introduction to Stochastic Processes*, Prentice-Hall, New Jersey, 1975.

Clarke, Lt. Col. Ronald, "Real-Time Unit Level WRSK Capability Assessment System," *Proceedings of the USAF Logistics Capability Assessment Symposium* (LOGCAS 82), March 1982.

Crawford, G., *Palm's Theorem for Nonstationary Processes*, The RAND Corporation, R-2750-RC, October 1981.

Crawford, Gordon B., *Variability in the Demands for Aircraft Spare Parts: Its Magnitude and Implications*, The RAND Corporation, R-3318-AF, November 1987.

DeGroot, M., *Optimal Statistical Decisions*, McGraw-Hill, New York, 1970.

DRC, *Combat Analysis Capability (CAC) Functional Description*, Dynamic Research Corporation, E-8726-U, November 1983.

DRC, *User's Manual, Sustainability Assessment Module* (SAM, D087C), Dynamic Research Corporation, E-13109-U, July 1987.

Feller, W., *An Introduction to Probability Theory and Its Applications*, Vol. 1, 3rd edition, Wiley, New York, 1968.

Gage, T. W., and Capt. A. J. Ogan, "How Much Do We Need?—What Do We Take? A Tool to Help—Now! " *Air Force Journal of Logistics*, VIII(1), Winter 1984.

Gross, D., and C. M. Harris, *Fundamentals of Queueing Theory*, John Wiley & Sons, New York, 1974.

Hillestad, R. J., *Dyna-METRIC: Dynamic Multi-Echelon Technique for Recoverable Item Control*, The RAND Corporation, R-2785-AF, July 1982.

Hillestad, R. J., and M. J. Carrillo, *Models and Techniques for Recoverable Item Stockage When Demand and the Repair Process Are Nonstationary—Part I: Performance Measurement*, The RAND Corporation, N-1482-AF, May 1980.

Isaacson, Karen E., and P. Boren, *Dyna-METRIC Version 5: A Capability Assessment Model Including Constrained Repair and Management Adaptations*, The RAND Corporation, R-3612-AF, 1988.

Miller, L. W., R. E. Stanton, and G. B. Crawford, *Dyna-SIM: A Nonstationary Queuing Simulation with Application to the Automated Test Equipment Problem*, The RAND Corporation, N-2807-AF, July 1984.

Mood, A., F. Graybill, and D. Boes, *Introduction to the Theory of Statistics*, 3rd edition, McGraw-Hill, New York, 1974.

Ogan, Capt. A. J., "An Alternative Approach to Base Level Capability Assessment," *Logistics Spectrum*, Fall 1983.

Operational Requirements, Management, and Computation of War Reserve Materiel (WRM), Hq. Air Force Logistics Command, AFLCR 57-18, 1979.

Palm, C., "Analysis of Erlang Traffic Formula for Busy-Signal Arrangements," *Ericson Techniques*, No. 5, 1938.

Parzen, E., *Stochastic Processes*, Holden-Day, San Francisco, 1962.

Pyles, R. A., *The Dyna-METRIC Readiness Assessment Model: Motivation, Capabilities and Use*, The RAND Corporation, R-2286-AF, July 1984.

Rice, D. B., *Defense Resource Management Study: Case Studies of Logistics Support Alternatives*, Office of Secretary of Defense, February 1979.

Sherbrooke, C. C., "METRIC: A Multi-Echelon Technique for Recoverable Item Control," *Operations Research*, Vol. 16, 1968.

Stevens, R. J., and Jack M. Hill, *Estimating the Variance-to-Mean Ratio for Recoverable Items in the ALS Marginal Analysis Algorithms*, Working Paper No. 49, Hq. Air Force Logistics Command, XRS, March 1973.