

20030225008

(2)

AD-A257 838



A RAND NOTE

DTIC
ELECTE
DEC 7 1992
S C D

TSAR User's Manual—A Program for Assessing the Effects of Conventional and Chemical Attacks on Sortie Generation: Vol. I, Program Features, Logic, and Interactions

Donald E. Emerson

September 1990

92-30887

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

RAND

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.

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

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

A RAND NOTE

N-3011-AF

TSAR User's Manual—A Program for Assessing the Effects of Conventional and Chemical Attacks on Sortie Generation: Vol. I, Program Features, Logic, and Interactions

**Donald E. Emerson
with Louis H. Wegner**

September 1990

**Prepared for the
United States Air Force**

DTIC QUALITY INSPECTED 8

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

RAND

PREFACE

This Note is one of a four-volume set that collectively describes the latest versions of the TSAR (Theater Simulation of Airbase Resources) and TSARINA (TSAR INputs using AIDA) computer models, which were developed at The RAND Corporation to assess the effect of attacks on the sortie generation capabilities of airbases. These new versions replace earlier versions, including the versions documented in 1985. Among the more significant new features are those that permit representation of (1) austere dispersed operating bases, (2) attacks on the minimum operating surface (MOS) defined after prior attacks, (3) multi-step parts and equipment repairs, (4) repair of damaged aircraft shelters, (5) improved fidelity in the runway repair representation, and (6) damage generated by the delayed detonation of unexploded ordnance (UXO). This development was carried out under the Project Air Force Resource Management Program project entitled "TSAR/TSARINA."

The TSAR model provides an analytic context within which a variety of airbase improvements may be tested. New passive defenses, new chemical defenses, new maintenance doctrine, improved base repair and recovery capabilities, increased stock levels for parts and equipment, and concepts for improved theater-wide resource management can be examined for their effect on aircraft sortie generation. The TSAR model has also proven useful for evaluating initiatives that would improve weapons and weapons-delivery systems, enhance multibase support, upgrade the reliability and maintainability of new aircraft designs, and revise training curricula to broaden the capabilities of maintenance specialists. These models have been briefed to several Air Force organizations during the development process and are currently used at several Air Force agencies, aerospace corporations, and at selected overseas sites.

This volume of the *User's Manual* provides a full description of the logic used in the TSAR model, as well as an understanding of interrelations among the many elements of the logic. The companion Notes include:

- N-3010-AF *TSARINA—A Computer Model for Assessing Conventional and Chemical Attacks on Airbases*
- N-3012-AF *TSAR User's Manual—A Program for Assessing the Effects of Conventional and Chemical Attacks on Sortie Generation: Vol. II—Data Input, Program Operation and Redimensioning, and Sample Problem*

N-3013-AF *TSAR User's Manual—A Program for Assessing the Effects of
Conventional and Chemical Attacks on Sortie Generation: Vol.
III—Variable and Array Definitions and Other Program Aids*

ACKNOWLEDGMENTS

I particularly wish to acknowledge Dr. Louis Wegner of RAND, co-author of the 1985 set of TSAR/TSARINA manuals, for his important contributions to these models in the early 1980s. Dr. Wegner's contributions include (1) the data structure for runway and taxiway damage, (2) the efficient algorithms for runway and taxiway repair, (3) the elegant, compact representation of deposition, evaporation, and vapor transport of chemical agents, and (4) the modeling of chemical casualties. Dr. Wegner also wrote Sec. V in the 1985 TSARINA manual and portions of Sec. IX of the 1985 TSAR manual; these sections remain essentially unchanged in the current versions.

These latest TSAR/TSARINA user's manuals, and the latest model software, have benefited from many helpful suggestions offered over the years by TSAR/TSARINA users. Their help has ranged from ideas for clarifying the documentation, to pinpointing obscure coding errors. And several were ideas for additional capabilities, many of which are reflected in the new features made available in these latest versions. I wish especially to note the suggestions of RAND colleagues John Folkesson and Michael Kennedy, the personnel at Orlando Technology (notably the late Dale Robinson), Ted Hayes (while at JAYCOR), Larry George of the Lawrence Livermore Laboratory, and Captains David Deiner and Robert O'Neill of the Air Force's Center for Studies and Analysis; all have helped to make possible a more effective product.

CONTENTS

PREFACE	iii
ACKNOWLEDGMENTS	v
FIGURES AND TABLES	xi
GLOSSARY	xiii
Section	
I. SUMMARY OF TSAR CAPABILITIES	1
II. TSAR DOCUMENTATION	12
III. STRUCTURAL OVERVIEW	14
1. Input and Initialization	14
2. Simulation	16
3. Output	21
4. Data Storage	24
IV. ON-EQUIPMENT AIRCRAFT MAINTENANCE	27
1. Task Representation	28
2. Task Sequencing	33
3. Task Descriptions	34
4. Shops and Unscheduled Task Collections	36
5. Illustrative Shop-Task Sequence	38
6. TSAR Determination of Aircraft Maintenance Requirements	41
7. Limiting Unscheduled Maintenance	43
8. Aircraft Operations at Dispersed Operating Bases	44
9. Rear Maintenance Base	46
10. Aircraft Maintenance Management	48
11. On-Equipment Task Initiation	50
12. Cannibalization	53
13. Check Flights to Validate Maintenance Actions	56
14. Postflight Inspections and Morning Preflight Inspections	57
15. Deferred Maintenance	58
16. Aircraft Status Projection	59
17. Preparing the Projections of Supply and Demand	61
V. PREFLIGHT TASKS AND MUNITIONS BUILDUP	63
1. Management of Preflight Tasks	65
2. Mission Assignment	66
3. Aircraft Reconfiguration	67
4. Munitions Loading	68
5. Refueling	69
6. Munitions Buildup	69

VI.	OFF-EQUIPMENT MAINTENANCE—PARTS AND EQUIPMENT REPAIR	72
	1. Initialization of Parts Inventory and Pipeline Data	72
	2. Initialization of Stocks for Battle Damage Repairs	76
	3. On-Base Parts Repair	77
	4. Off-Base Parts Repair	81
	5. Support Equipment Repair	82
VII.	AIRCRAFT SORTIE DEMAND AND AIRCREW MANAGEMENT	86
	1. Generating Sortie Demand Data	87
	2. Launching the Aircraft	88
	3. Aircrew Management	89
VIII.	AIRBASE ATTACK AND RECOVERY	91
	1. Attack Effectiveness Estimated with TSARINA	91
	2. Representation of an Airbase in TSAR	92
	3. Data Transfer from TSARINA to TSAR	100
	4. Overview of Attack Time Computations	102
	5. Data Processing for Attacks	103
	6. Postattack Recovery and Reconstruction of Buildings and Other Facilities	114
	7. Completion or Interruption of Civil Engineering Repairs	119
IX.	AIRBASE OPERATIONS IN A CHEMICAL ENVIRONMENT	121
	1. Introduction	121
	2. Computation of Allowable Work Time and Required Rest Time	123
	3. Application of Work Time Criteria to TSAR Tasks	125
	4. Surface Contamination and Vapor from Chemical Attacks	127
	5. Individual Chemical Protection Equipment	131
	6. Representation of Collective Protection Shelters	132
	7. Estimation of Casualties from Chemical Effects	136
X.	COMMUNICATIONS	144
	1. Scheduled Shipments from CONUS	144
	2. Responsive Shipments from CONUS	145
	3. Intratheater Shipments	145
	4. Intratheater Resource Status Reports	147
XI.	THEATER RESOURCE MANAGEMENT	149
	1. Management of Fillers, Aircraft Reallocation, and Diversion	152
	2. Periodic Review and Reallocation of Resources	155
	3. Acquisition of Spare Parts	157
	4. Disposition of Newly Repaired or Newly Acquired Parts at a Central Supply Point	159
	5. Repair Priority Determination at a CIRF	159
	6. Use of the Theater Parts Repair Facility as a Backup	161
XII.	DATA INPUT	163
XIII.	SIMULATION CONTROL	167

XIV. SUPPORT SERVICES	173
Subroutine ASSAY	174
Subroutine BREAK	174
Subroutine CALCLO	175
Subroutine CHECK	175
Subroutine CKTEMP	176
Function CWMOPP	177
Subroutine CWSHIFT	177
Subroutine CWTIME	177
Subroutine ENDAC	178
Function FTIME	178
Subroutine GOREST	179
Subroutine HEAP	180
Subroutine INTRUP	181
Subroutine KILLAC	181
Subroutine LOSSES	182
Subroutine NORRPT	182
Subroutine REDPEO	183
Subroutine RESET	184
Subroutine SHIFT	184
Subroutine STRTSK	186
Subroutine TRIAGE	186
Function TTIME	187
Subroutine UPDATE	188
Subroutine WAIT	188
Additional Services	188
XV. OUTPUT	190
1. Output Controlled by the Variable PRINT	192
2. Output Controlled by the Variable STATFQ	192
3. Provisions for Customized User Output	194
4. Provisions for Subsequent Postprocessor Analysis and Graphics	196
REFERENCES	201

FIGURES

1.	Basic structure of the TSAR simulation	15
2.	TSAR input and initialization	17
3.	Subroutines used to conclude tasks.	19
4.	Subroutines used to manage irregularly occurring airbase activities.	20
5.	Subroutines used to manage scheduled and periodic activities.	22
6.	Subroutines used for airbase attack assessments.	23
7.	Structural representation of on-equipment maintenance.	39
8.	Layout of runways and taxiway network.	94
9.	Variations of facility repair time	118
10.	Illustrative fit to the surface deposition contours by rectangular layers	128
11.	Indoor and outdoor vapor concentration for a ventilation rate of one air exchange per hour.	139

TABLES

1.	Events stored in the period heap	25
2.	On-equipment maintenance activities	29
3.	Cannibalization constraints.	55
4.	TSAR default location assumptions	99
5.	Example repair times for a facility with size = 100	119
6.	Periodic theater-wide resource checks	156
7.	Acquisition of spare parts.	158
8.	Changes available for TSAR control variables	170
9.	Output data controlled by the variable PRINT	193

GLOSSARY

ABDR	Aircraft Battle Damage Repair
AGE	Aerospace Ground Equipment and other support equipment used for carrying out various tasks
AIDA	Airbase Damage Assessment model; the forerunner of TSARINA
AIS	Avionics Intermediate Shops; special test equipment used for repairing avionic LRUs and SRUs
AMU	Aircraft Maintenance Unit; the organization providing maintenance for an aircraft squadron
ATC	Air Traffic Control
BAI	Battlefield Air Interdiction
BKEP	Ballistic Kinetic Energy Penetrator
BLSS	Base-Level Self-Sufficiency stock of aircraft spare parts, composed of the stocks for peacetime, plus additional material to meet wartime demands
BW	Bacteriological Warfare
CAP	Combat Air Patrol
CAS	Close Air Support
CBU	Cluster Bomblet Unit
CILC	Centralized Intermediate Logistics Concept
CIRF	Centralized Intermediate Repair Facility
COB	Collocated Operating Base
COMO	Combat-Oriented Maintenance Organization
CONUS	Continental United States
CP	Collective Protection
CPS	Collective Protection Shelter
CRS	Component Repair Squadron; a wing-level organization responsible for parts repair
CW	Chemical Warfare
DOB	Dispersed Operating Base
EMS	Equipment Maintenance Squadron; a wing-level organization responsible for equipment maintenance and repair

EOD	Explosive Ordnance Disposal
FIFO	First In, First Out
FRAG	FRAGmentary order that specifies flight requirements
GP	General-Purpose bomb
ILM	Intermediate Logistics Maintenance; on-base parts repair supporting the AMUs
IPE	Individual Protection Equipment for a chemical environment
LCOM	Logistics Composite Model
LRU	Line Replaceable Unit; an aircraft spare part with distinguishable subordinate components
MAE	Mean Area of Effectiveness
MMD	Mean Mass Diameter
MOB	Main Operating Base
MOPP	Mission-Oriented Protective Posture (the chemical protection ensemble)
MOS	Minimum Operating Surface
MP	Monitoring Point
MVDC	Mobility, Visibility, Dexterity, and Communications
NMCS	Not Mission Capable because of lack of Spare parts
NORS	Not Operationally Ready because of lack of Spare parts; same as NMCS
NRTS	Not Repairable This Station
OST	Order and Ship Time in days; time for a NRTSed or condemned part to be replaced
PAA	Program Authorization, Aircraft
POL	Petroleum, Oils, and Lubricants; often used as an abbreviation for aircraft fuel
POS	Peacetime Operating Stock; an organization's stock of aircraft spare parts for aircraft maintenance in peacetime
QPA	Quality Per Aircraft; number of parts of the same kind used on an aircraft.
RAM	Rapid Area Maintenance; special mobile teams used for repairing aircraft battle damage

RN Random Number

RR Flight line maintenance that removes and replaces malfunctioning aircraft parts with serviceable components; generally implies no local repair

RRR Flight line maintenance that removes, repairs, and replaces aircraft spare parts (actually, usually removes and replaces with a serviceable unit, and then repairs the malfunctioning unit)

RRR Rapid Runway Repair

SAMSOM Support Availability Multi-System Operations Model

SCL Standard Combat Load that designates the aircraft configuration and the mission dependent munitions to be loaded

SE Support Equipment, usually referred to as AGE in TSAR

SRU Shop Replaceable Unit; a component of an LRU

TBM Tactical Ballistic Missile

TRAP Tanks, Racks, Adaptors, and Pylons

TSAR Theater Simulation of Airbase Resources

TSARINA TSAR INputs using AIDA

UXO Unexploded Ordnance

WRM War Reserve Material

WRSK Wartime Readiness Spares Kit

I. SUMMARY OF TSAR CAPABILITIES

The TSAR (Theater Simulation of Airbase Resources) model simulates a system of interdependent theater airbases, supported by shipments from the Continental United States (CONUS) and by intratheater transportation, communication, and resource management systems. By capturing the interdependencies among 11 classes of resources, the simulation permits decisionmakers to examine the implications of many possible improvements in terms of their effects upon the sortie generation capabilities of a system of airbases. The simulation also allows examination of the effects of damage inflicted by enemy airbase attacks using both conventional and chemical weapons and the results of efforts to restore operations.

The classes of resources treated in TSAR are (1) aircraft, (2) aircrews, (3) ground personnel, (4) support equipment (AGE), (5) aircraft parts, (6) aircraft shelters, (7) munitions, (8) tanks, racks, adaptors, and pylons (TRAP), (9) petroleum, oils, and lubricants (POL), (10) building materials, and (11) airbase facilities. Many different types of each class of resource may be distinguished. When parts are included in the simulation, initial stocks may be specified, or TSAR will initialize the parts data according to standard algorithms for peacetime operating stock (POS), base-level self-sufficiency stock (BLSS), and wartime readiness spares kit (WRSK), and will also initialize the stock location in the depot pipeline.

TSAR is a Monte Carlo discrete-event simulation model that analyzes the interrelations among available resources and the capability of the airbases to generate aircraft sorties in a dynamic, rapidly evolving wartime environment. On-equipment maintenance tasks, parts and equipment repair jobs, munitions assembly, and facilities repair tasks are simulated at each of several airbases. If desired, the constraints imposed by wearing individual chemical protection equipment (IPE) during the conduct of these activities may be simulated. A broad range of policy options that would increase initial resources, modify maintenance doctrine, or improve theater resource management may be assessed using TSAR. Provisions are also included that provide the user a capability to assess dynamic variations in key management policies.

TSAR is readily adaptable to initial conditions encompassing a broad range of complexity. When specific features are not needed for the examination of a particular

issue, they simply need not be used. Thus, TSAR permits one to represent either a single base, a set of independent airbases, or a set of interdependent airbases, without any adjustment or modification of the program. Since each base may have unique characteristics, one may analyze situations that involve main operating bases (MOBs), collocated operating bases (COBs), as well as dispersed operating bases (DOBs), each with their particular characteristics. And TSAR readily accommodates users who do not wish to examine the effects of airbase attacks using conventional or chemical weapons or who may wish to ignore the possible restraints imposed by shortages of aircrews, shelters, ground personnel, equipment, aircraft parts, munitions, TRAP, and/or fuel; users may also consider or ignore the special problems associated with the air traffic control constraints on flight operations and with operations in a chemical environment. TSAR adapts automatically to all such problem representations.

TSAR provides users a means by which a rich variety of improvements for theater airbases may be tested in a common context. By comparing how such improvements affect the system's capabilities for generating effective combat sorties, TSAR can assess new passive defenses, new maintenance doctrines, dispersed aircraft operating locations, modified manning levels, enhanced cross-training, improved clothing and facilities for chemical protection, improved procedures and equipment for increasing runway utilization, increased stock levels for parts and equipment, and many others, as well as several concepts for theater-wide resource management. TSAR has also provided an effective context for assessing new weapon concepts and improved reliability and maintainability of prospective aircraft designs.

An important objective in the original design formulation was to achieve a sufficiently high speed of operation that the extensive (often trial and error) sequence of runs so frequently necessary in research and analysis would be economically practical. Adaptation of existing models (e.g., LCOM [1], SAMSOM [2]) was rejected because modifications would have been extensive and execution times prohibitive for problems of the size that were contemplated. The TSAR program is written in the widely available FORTRAN language. It achieves a substantially higher speed by virtue of more efficient processing and by taking advantage of recent core storage increases of modern computers. In its current formulation, TSAR makes no intermediate use of auxiliary high-speed storage units (e.g., disks, tapes) except for the TSARINA assessments of air attacks and the initial conditions for multiple trials.

In TSAR, several types of aircraft can be assigned to each airbase. The aircraft of a given type at any airbase may be supported by a common pool of resources (personnel and equipment), or, as in the combat-oriented maintenance organization (COMO) concept, the aircraft may be organized into two or three subgroups (squadrons) each supported by its own set of resources (AMU, aircraft maintenance unit). Small groups of aircraft may be transferred to DOBs at prearranged times and subsequently directed to return to their host base. The aircraft are launched on sorties in response to a set of user-supplied sortie demands differentiated by base, aircraft type, mission, and priority; if a base is not specified, the sortie demands are allocated to the base best able to generate the necessary sorties. Flights may be scheduled or they may be scrambled on demand using aircraft that have been placed on alert. Aircraft may be launched late, when permitted, or they may ground abort, and flights may be canceled if required by air traffic control constraints. An early morning inspection may be imposed on all ready aircraft at designated airbases at a user-specified time.

When launched, aircraft may air abort or may be lost on a combat mission; when an aircraft returns it may be damaged, require decontamination, still have munitions, be due for phased (periodic) maintenance, and have several unscheduled maintenance task requirements. These maintenance tasks are normally done at the aircraft's operating base, but an aircraft may be ferried to a rear base for certain specified maintenance tasks. A check flight may be required following specified maintenance tasks to validate the maintenance action. If an aircraft is operating from a DOB and a malfunction occurs that is detected and must be corrected at its host base, the aircraft recovers at the host; if not detected, it must be ferried to the host after recovery at the DOB. When aircraft are lost, a replacement may be ordered from CONUS, or if aircraft are set aside in the theater as fillers, they may provide rapid replacements for lost aircraft and, if specified, for aircraft ferried to the rear for maintenance. When filler aircraft are used to replace losses, a replacement for the filler force is ordered from CONUS if such resources are available.

When an airbase runway has been closed because of an airbase attack, aircraft scheduled to land are diverted to other bases, preferably to one that normally operates the same type of aircraft. An aircraft scheduled to land at a DOB whose runway is closed looks for another DOB with the same host; if none are available, the aircraft recovers at the host or, if it is closed, at another host of like-type aircraft. If base sortie generation capabilities are forecast daily (an option), the base best able to support the aircraft is

selected. During the period that a runway remains closed, that airbase's sortie demands can be allocated to functioning airbases with the appropriate type of aircraft in proportion to either the aircraft available or, if base capabilities are forecast daily, the bases' sortie generation capabilities. When a runway has been reopened, that base's aircraft recover at their parent base on completion of their next combat sortie, if base sortie generation capabilities are not forecast or, if they are, when their parent base's sortie generation capability per available aircraft is within a specified percentage of that at the temporary base.

When an aircraft lands, it may be refueled at a hot-pit hydrant. Each aircraft is assigned to an aircraft shelter if one is available; if not, it is parked on one or another of the designated ramps. A postflight inspection, dependent on the type of mission flown, may be imposed, and chemical decontamination of the aircraft is scheduled if required. The next mission assignment for each aircraft is selected tentatively when the aircraft lands; that selection takes into account the known demand on that base for sorties and the projected capability of the aircraft at that base to meet those demands. The selection also takes into account which of that aircraft's unscheduled maintenance tasks would need to be accomplished for the different missions and when that particular aircraft could probably be readied for the different missions. All tasks that are not essential for the tentative mission assignment may be deferred and the available resources concentrated on required tasks. If aircraft are eventually found not to be needed for the mission for which they were readied, they are reassigned and reconfigured for a more appropriate mission. If phase maintenance is to be simulated, it may be deferred during specific times during the scenario and will be done at night when not deferred.

On-equipment maintenance tasks may require several people, specialized equipment, and spare parts; each task is either a single set of such requirements—i.e., a simple task—or a network of tasks, each with its own demand for personnel, equipment, and parts. When resources are limited, those aircraft most likely to be readied first (given sufficient resources) may be given priority. The basic input data that govern the probabilities for unscheduled maintenance tasks (other than battle damage repairs) may be used directly for the simulation or varied statistically to reflect unexpected differences between planned levels and "actual" wartime experience. Furthermore, these task probabilities—i.e., the breakrates—may either have a fixed rate or be varied daily by shop and aircraft type as a function of achieved sortie rate or other user-specified

adjustments. When aircraft are to operate from DOBs, the basic input data also specifies the likelihood that a malfunction that cannot be corrected at a DOB will be detected before the aircraft lands so that it may be flown directly to the host base.

If a required part is not available, (1) the broken one that is removed may be repaired on base, (2) the part may be cannibalized from another aircraft, (3) a part may be obtained by lateral resupply from a specified subset of bases, or (4) the part may be ordered from a central source within the theater. When a part is cannibalized, it may itself be broken. When a part cannot be repaired on base (is NRTS), it may be sent to a neighboring base or to a centralized facility in the theater designated to perform intermediate maintenance—i.e., a centralized intermediate repair facility (CIRF). When parts cannot be repaired within the theater, the user may request a replacement from a depot in CONUS. Parts may either be a simple part or a line replaceable unit (LRU) that has a defective shop replaceable unit (SRU). Simple parts may be repaired on base using either a unique procedure or a procedure selected at random from two or more repair procedures. Each procedure may include a single step or a sequence of repair steps. For LRUs, the resource requirements to diagnose and replace the faulty SRU are specified separately for each SRU. Faulty SRUs withdrawn from an LRU may themselves be repaired on base or NRTSed to another location for repair.

The various types of support equipment used in on-equipment and off-equipment jobs, in munitions assembly and loading tasks, and by base civil engineers are themselves subject to malfunction and repair. Equipment repair uses either a unique procedure or is done using one of several procedures selected at random. Again, each procedure may consist of one or more repair steps. The special complexities of full and partial mission capability of AIS test equipment used to repair LRUs and SRUs for late-model aircraft may also be simulated.

Each maintenance task, parts repair job, and equipment repair job is done by the personnel and equipment associated with a particular work center or shop. The user may group the resources and tasks into up to 25 different "shops" exclusive of those associated with the scheduled preflight maintenance tasks. Because each shop may be assigned several different types of personnel and equipment, those engaged in on-equipment and off-equipment tasks may be the same or different depending upon how the user wishes to define the base's maintenance policies.

There is substantial flexibility in defining the rules by which aircraft maintenance tasks are processed. The user may permit the activities of certain groups of shops to proceed simultaneously or may require that the activities of several such groups of shops proceed in a specified order. The user may also control these prescriptions for simultaneous and sequential operations separately for each aircraft type at each base. Furthermore, for those groups of shops that are permitted to proceed simultaneously, certain exceptions may be specified in the form of lists of activities that are incompatible with each particular task. These features permit alternative maintenance operating doctrines to be simulated and to be examined for their influence on sortie generation capabilities. Work speed-up and other procedures to shorten on-equipment, preflight, and off-equipment activities may also be specified.

Scheduled preflight tasks are also associated with the shop structure. These tasks involve aircraft refueling and the loading of both basic defensive munitions and mission-dependent munitions. The likelihood that the basic munitions and the mission-dependent munitions are retained from the previous sortie can be specified independently for the two classes of munitions. After mission assignment, aircraft configuration is checked and, if necessary, the aircraft is reconfigured; this may involve one or two separate tasks, each of which may require TRAP, personnel, and equipment. The loading of the mission-dependent munitions may also involve one or two separate tasks, each with its distinct requirements.

When munitions assembly tasks are simulated, munitions demands are projected periodically to define which types of munitions need to be assembled. Such jobs may require both personnel and equipment, much like other tasks that are simulated in TSAR, as well as components from which the munitions are to be assembled. When munitions assembly is simulated, initial stocks and components, as well as shipments, are distinguished as to whether the munitions are assembled.

Chemical protective clothing may be required to be worn at all times for any or all tasks, whether or not a chemical attack has occurred, or only when required by the chemical environment. Different types of chemical ensembles may be prescribed at different airbases. The increased task times that result from restrictions on mobility, visibility, dexterity, and communication and the buildup of excessive body temperature because of the poor heat-transfer properties of such clothing may be defined uniquely for each task. If the work crew temperatures rise too high, the crew may suffer heat

exhaustion and will be hospitalized; if they do not collapse they may have to wait until they have cooled down to a specified level.

Several features permit the user to simulate various workaround procedures that can alleviate resource constraints. One such feature permits the user to specify alternative resource requirements for any on-equipment task, parts repair job, equipment repair job, weapons assembly task, or civil engineering job; for example, one might specify that a three-man crew could do a normal four-man job in 50 percent more time. Similarly, when TRAP or munitions shortages do not permit the normal or preferred munitions to be loaded for a mission, alternative loadings may be specified. A third workaround feature permits the user to designate that certain types of personnel have been cross-trained and that they may replace or assist certain other specialists. This personnel substitutability feature is operative only at specified bases and only for those on-equipment tasks, part and equipment repairs, munitions assembly tasks, and civil engineering jobs that have been specified.

The effects of damage and chemical contamination due to airbase attacks may be simulated. Input data generated by TSARINA [3] normally define the time and location of the attacks, the damage to individual aircraft shelters and other facilities, the contamination at different locations, damage to the runways and taxiways, and the percentage of conventional damage suffered by the personnel, equipment, parts, munitions, TRAP, and POL at each facility. (Only simple conventional attacks can be defined for TSAR without using the TSARINA airbase damage assessment model.) When aircraft, aircraft shelters, or other facilities sustain conventional damage, some portion of the personnel, equipment, and parts at these locations may also be lost. Damage to runways and taxiways may interrupt flight operations, and damage to other key facilities can degrade air traffic control performance. Following a chemical attack, the likelihood that personnel sustain an incapacitating or lethal dose is based on the warning time for the attack, the arrival time of the chemical contaminants, and the degree of personnel protection. Aircraft may be assigned a specific shelter when they land, but the aircraft may be partially exposed when certain shop operations are underway at the time of airbase attack. Alert aircraft may be given priority for assignment to a specific set of shelters, and the damage to these aircraft may be distinguished from that for other aircraft. Aircraft in excess of those that may be placed in shelters are assumed to be parked on designated parking ramps and to sustain a loss rate appropriate for that ramp.

TSAR decrements the various resources to the extent implied by the damage and chemical casualty data. If personnel have generated excessive heat because of their chemical protective clothing, they are required to rest until their temperatures have fallen to the specified level. If personnel sustain casualties, other personnel may be required to provide buddy care for a specified time, to simulate helping the casualties obtain medical assistance. After user-stipulated delays, to roughly account for the disruptive effects of the attack, personnel resume their activities unless a specific facility is required and has been damaged; these delays can be varied in relation to the strength and extent of the attack.

Replacement resources (aircraft, pilots, personnel, parts, munitions, TRAP, and building materials) may be ordered from CONUS when losses are sustained. Resources available for replacing losses may be specified, and the time required to replace the loss may be specified, independently.

After an airbase attack, civil engineering personnel, equipment, and building materials may be allocated to repair the runway and taxiway network. The location and number of such repairs are based on the numbers of unexploded ordnance (UXO), mines, and craters from all previous attacks that have not yet been repaired, plus those delivered by the most recent attack. Designated craters from an air attack may be concentrated on the minimum operating surface (MOS) from the prior attack if enemy intelligence is assumed adequate. When the unexploded ordnance has been removed from one subsection of the intended MOS, mine clearance may begin; and when clearance has been completed on that subsection, crater repair is commenced. UXO may be timed to explode at a predetermined time; if explosive ordnance disposal (EOD) or civil engineering personnel or equipment are working on that weapon, or are in the vicinity, they may sustain casualties. The order in which the MOS subsections are cleared is selected for efficient utilization of the available civil engineering resources. The prioritization of taxiway repairs is designed to maximize the rate at which undamaged shelters obtain access to the section of runway that is being repaired.[4] When the MOS has been cleared, the user may specify that the MOS should be extended, that the entire surface should be cleared, or that the main runway should be cleared when the MOS is on a secondary runway; several extended clearance options are available. Resources to repair damaged aircraft shelters and the other facilities are allocated according to a priority specified by the user. Operation of these facilities is resumed when they once again are functional.

In addition to simulating a set of airbases, the user may also specify the existence of a theater reserve of filler aircraft, a centralized theater distribution center, or a centralized theater repair facility at which some or all intermediate maintenance is conducted. At the user's option, the filler aircraft can be used to replace aircraft losses and aircraft that have been withdrawn to a rear base for maintenance. When additional aircraft resources are specified as available in CONUS, they supplement the filler force. The filler aircraft are managed so as to maintain the inventories at the operating bases to the extent possible. Specific bases may be designated as host bases, and elements of their aircraft may be directed to transfer to a DOB at a prearranged time and to operate from that austere base until requiring maintenance at the host base or until recalled.

The centralized theater distribution facility can receive spare parts from CONUS and either retain them until demanded by a base or transship (some or all) to the base with the earliest projected requirement. Such a facility can also be used to direct the lateral shipment of parts and other resources from one base to another. A theater parts repair facility, sometimes referred to as a CIRF, is assigned maintenance personnel, equipment, and spare parts (LRUs and SRUs). Parts are shipped to and from the CIRF from the operating bases and are processed in the manner prescribed by the user's choice of which theater management rules are to govern these operations.

The simplest rules for CIRF operation prescribe that faulty parts are repaired in the order in which they arrive and that they are returned to the sender. The user may also invoke a variety of more complex theater management algorithms, not only for selecting what to repair and how to dispose of parts when they have been repaired but for reallocating personnel, equipment, and parts among the several operating bases. Repair priorities can be based on existing and projected demands and on the relative necessity of parts for the various missions. Shipment priorities are related to the current and projected demands, on-base reparable, and en route serviceables. When central stocks are insufficient to meet a base's demand, another base can be directed to ship the required part, if both the requesting base and the donor base meet certain conditions relative to the importance of the demand and the availability of stock.

Daily forecasts can be prepared (an option) of each base's capabilities for generating different kinds of missions with different types of aircraft. These forecasts provide the basis for various aircraft management decisions. One application is in selecting which base is to be assigned the sortie demands for which no base has been

specified. These data can also be used for assignment decisions when aircraft must be diverted and when they are transferred from base to base to balance maintenance workloads.

The theater-wide management of the various resources is supported by a user-specified scheduled transportation system that may be subjected to delays, cancellations, and losses. TSAR also permits the user to represent a theater-wide reporting system that can be used to provide the central management authority with periodic resource status reports from the several operating bases; these reports may be delayed, incomplete, or lost.

When these transportation and communication systems are coupled with the sets of rules for distributing and redistributing resources among the operating bases, various concepts of theater resource management may be represented and examined in the context of realistic transportation and communication imperfections. In its current formulation, TSAR already includes certain alternatives for the theater management rules and has been designed to permit additions or modifications to be readily accommodated.

TSAR (and the companion TSARINA model) naturally have limitations and omissions that will inconvenience some potential users. The more obvious limitations derive from the manner in which the problem was bounded in designing TSAR. Some users will be bothered that TSAR treats friendly sorties simply as delays during which the aircraft are not present at an airbase; others will wish that active airbase defenses had been included as an integral part of the simulation, rather than being required to consider active defense tradeoffs externally to TSAR/TSARINA analysis; and still others will find that these tools would be more useful if the production-oriented batch processing of spare parts, as they are handled at depots, also were modeled.

Each of these design limitations could be a serious obstacle for some potential users, but none of these bounds was chosen casually or accidentally. All problems must be bounded, and we believe the choice of boundaries need not inhibit many useful and important analyses. Furthermore, it would be conceptually fairly easy to substantially extend or eliminate these boundaries because TSAR's data structure is sufficiently detailed to be compatible with many such additions. But most such additions would entail difficult design and programming efforts and would further increase TSAR's execution time and expand its input data collection problems.

The last of the limitations that should be highlighted is TSAR's data input requirements. As one elects to include more and more of the real world considerations that TSAR permits the user to include, these requirements become substantial. That is not a property of TSAR but of the richness of the user's problem definition; any approach to dealing with a problem at a comparable level of detail would require equivalent information. TSAR's main contribution to this dilemma is that it will function comfortably at many levels of detail, and the user may quite simply select or reject most of its features and the related data requirements. This extreme adaptability is illustrated on the last page of Vol. II, where 20 card images (2 blank, all trivial) show how really simple a problem can be. One important benefit of this flexibility is that analysts can test the potential sensitivity of their results for a particular effect for which the data would be difficult or costly to secure, using invented data that spans a reasonable range of uncertainty. If the results are reasonably insensitive to that variation, there is an argument for neglecting the effect; if they are sensitive, there is an argument for mounting the effort to secure the needed data.

II. TSAR DOCUMENTATION

TSAR documentation has been designed with four classes of readers in mind:

- Those seeking only a broad overview of TSAR's capabilities,
- Those without a background in programming who seek a full understanding of the logic in the TSAR simulation,
- Those responsible for preparing input materials and for operating TSAR, and
- Those interested in modifying and extending the existing program logic or trying to understand apparent errors.

Sections I through III of the introductory volume [5] for Early-TSAR were designed for the first group, and the entire volume is appropriate for the second; at the present time we do not plan to update that volume to describe the various improvements and extensions that have been introduced into the present TSAR. Only the three-volume *TSAR User's Manual* [6,7,8], designed primarily for the last two groups, has been updated; in addition to the earlier reports, we suggest that the first group read Sec. I in this volume, and the second group read the first 11 sections.

The *TSAR User's Manual* is built around four main sets of mutually supporting explanatory materials. It is intended primarily for those responsible for operating TSAR and for programmers who wish to extend the program logic, or are seeking to understand an apparent error. The TSARINA Manual [3] has also been rewritten and should be considered mandatory reading for anyone planning to use TSAR for examining the effects of airbase attacks.

This first volume of the *TSAR User's Manual* provides a succinct but complete discussion of the processing logic involved in the major subsets of tasks; eight sections deal with the simulation proper, and four deal primarily with housekeeping chores. The second major set of materials is the discussion of the input requirements, procedures, and formats that are found in Vol. II; these detailed discussions provide the only complete explanations for some of the numerous control options available with TSAR and must be considered *mandatory* reading for anyone planning to operate TSAR. The third set of

materials is the complete listings and explanations of the very substantial data base maintained within the simulation; these are located in Vol. III. The source code for the program, together with its *extensive comments*, provides the fourth set of materials. The discussions in subroutines INIT, CONTRL, READFT, IPARTS, DOSHIP, BOMB, and REORGN are particularly extensive and will prove helpful in tailoring TSAR for specific applications; any other questions regarding TSAR logic will probably be answered by the explanations included throughout all major TSAR subroutines.

The way these materials can be best used undoubtedly will vary widely. If the immediate concern is to decide on TSAR's adequacy for installation, the user's reading should probably begin with the introductory volume. If that decision has been made and the problem is to apply TSAR, the user might best begin with a cursory reading of the first sections of the introductory volume, or this volume, and then turn to the discussion of the input procedures for the sample problem in Sec. XX, Vol. II. As the user begins to understand how TSAR can be applied and starts to develop the needed input data, the user will want to refer to the detailed discussions of the data input procedures in Secs. XVIII and XIX, Vol. II. When questions arise as to how TSAR will deal with particular aspects of the problem, the user can consult the appropriate section in this volume.

As the first TSAR data base is built, the user will be well advised to create a dictionary of resources (as illustrated in Sec. XX, Vol. II) and to hold down the number of aircraft for the trial problem; that number can easily be increased later. This will minimize the time and trouble to locate, understand, and eliminate the errors that will inevitably creep into a user's first data base. One to three airbases, with six to eight aircraft per base, can provide quite useful and very rapid hands-on experience. And use of most of the output options (see Sec. XV), especially the SCROLL feature (CT2/1), will also provide useful insights into the simulation. As the user's first trial problem begins to generate these outputs, the user may need to refer to Sec. XXI, Vol. II, where the output formats are explained with illustrations from the sample problem.

When all appears to be behaving logically for a simple trial problem, it will be time to explore some of the more complex control variables that the user may elect to apply; only when those variables are mastered will it be appropriate to increase the size of the user's aircraft fleet (and reduce the volume of output data).

III. STRUCTURAL OVERVIEW

TSAR simulates the interaction between many types of events and many classes of resources and provides a wide variety of output information. To fully understand the simulation, one must understand what the events are, how decisions are made about when they begin and end, and what resources are required. Of particular importance are the internally generated events—events generated by other events in the model—that must be defined, initiated, and concluded, and that sometimes must wait or be interrupted. On-equipment aircraft maintenance tasks, off-equipment parts repair jobs, equipment repair jobs, munitions assembly jobs, and civil engineering reconstruction jobs generate such events.

In broadest terms, the TSAR simulation can be divided into three phases: (1) input and initialization, (2) simulation, and (3) output. The MAIN executive routine initiates these computational phases and, assisted by the TRIALS subroutine, controls processing for the specified number of trials (see Fig. 1). To facilitate processing and to avoid the necessity of searching extremely long time-ordered queues, the primary event structure in TSAR is divided into the nine different sets of events shown at the bottom of Fig. 1. The tenth category—scheduled and periodic activities—is a heterogeneous set of events and simulation housekeeping tasks that occur on a scheduled or periodic basis.

Each of the three phases uses various subroutines to carry out the required computations. The subroutine names are all printed in boldface capital letters in the figures of this section, and a brief description of each subroutine's function will be found in App. A, Vol. III.

1. INPUT AND INITIALIZATION

Figure 2 indicates the interactions among the subroutines that are used to input data and to initiate the various data arrays according to user-specified instructions. Subroutines INIT, INIT0, and INTT1 zero out the storage space for the labeled common statements and then subroutine INPUT (and subordinate routines) enter data that define the several control variables and describe the resource requirements for the different types of tasks, mission characteristics of the aircraft, on-base resource stock levels,

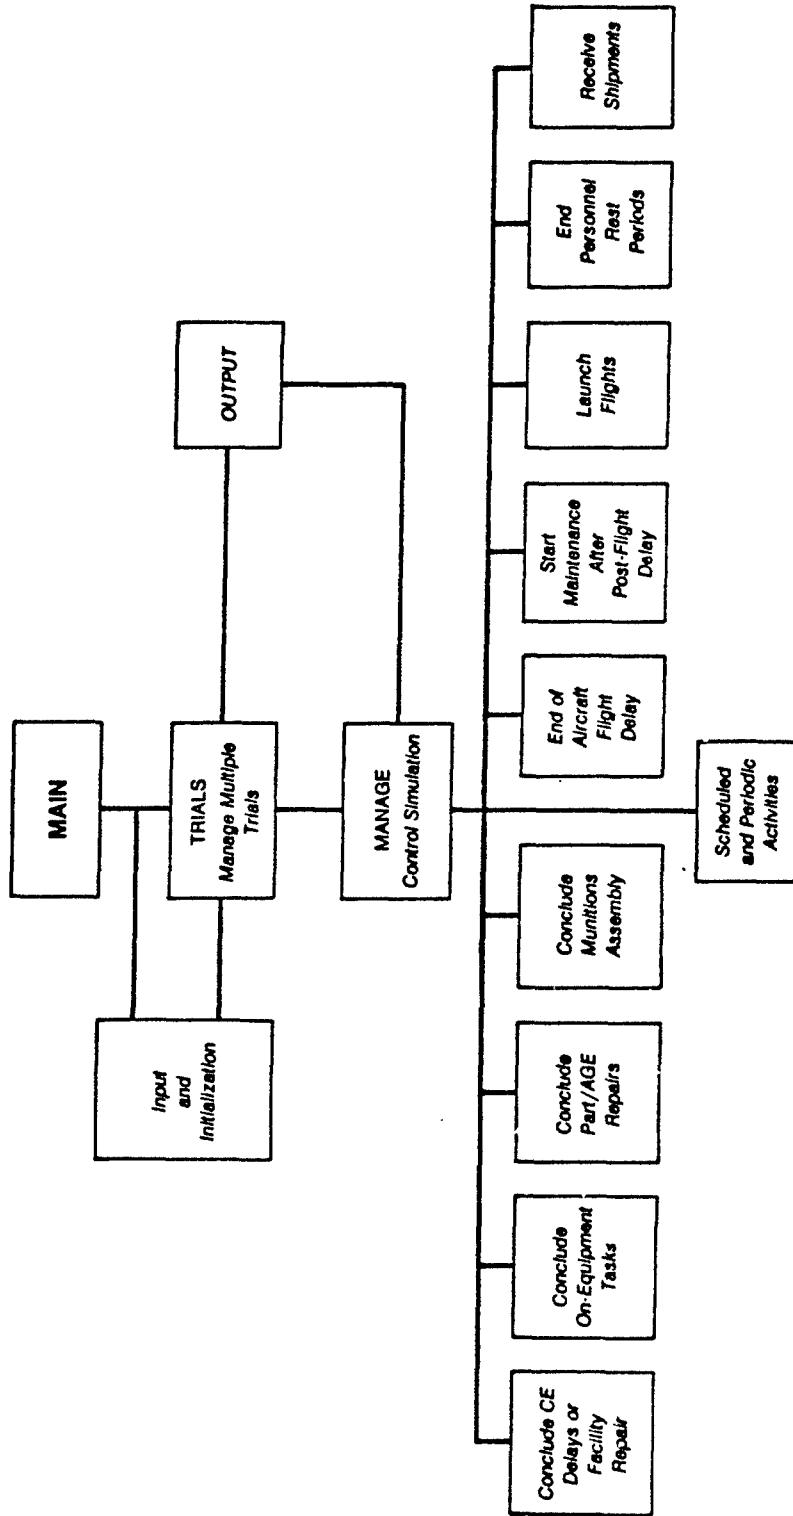


Fig. 1—Basic structure of the TSAR simulation

descriptions of the intratheater transportation and communication systems, and so on. Subroutines REVIEW, AUDIT, and WRAPUP then manage a series of data checks and auxiliary computations that generate a variety of derivative data used during the simulation. After subroutine INITIZ has initialized the dynamic storage arrays, subroutine TRIALS takes over control until the simulation is completed and the final outputs are prepared and printed.

Before transferring control to subroutine MANAGE, which manages the simulation proper, subroutine TRIALS completes several initialization tasks. Subroutines ICHECK and RREQTS compute derivative control data, and subroutines INLIST, HEADER, and CWLIST organize and print a record of the key control variables. Subroutines COMPRT, BASCAP, STATUS, SCSHIP, and ZSHOPS are called as specified by the user; under some conditions they are called to establish a different set of initial conditions for each trial. Subroutine AVGTME computes the "standard time" for on-equipment and off-equipment tasks in each backshop, and subroutine DAYONE reads the first set of sortie demand data.

When these initialization tasks are completed, TRIALS passes control to subroutine MANAGE, which carries out each simulated event in its appropriate time sequence.

2. SIMULATION

Basically the TSAR computer simulation processes one event after another in the order in which the events occur in simulated time, initiating whatever subsequent actions are dictated by the prescribed behavior logic for each type of event. Each of the ten sets of events shown in Fig. 1 is organized so that the next earliest event in each set is always known. The basic task performed by the subroutine MANAGE is to examine the earliest event that will occur in each of the ten sets of events and to determine which of the ten is to occur next. Simulation time is then advanced to that time, and control is transferred to the subroutine processing that event. When that event is completed, control is returned to subroutine MANAGE, which again examines the ten sets of events to determine which event should occur next.

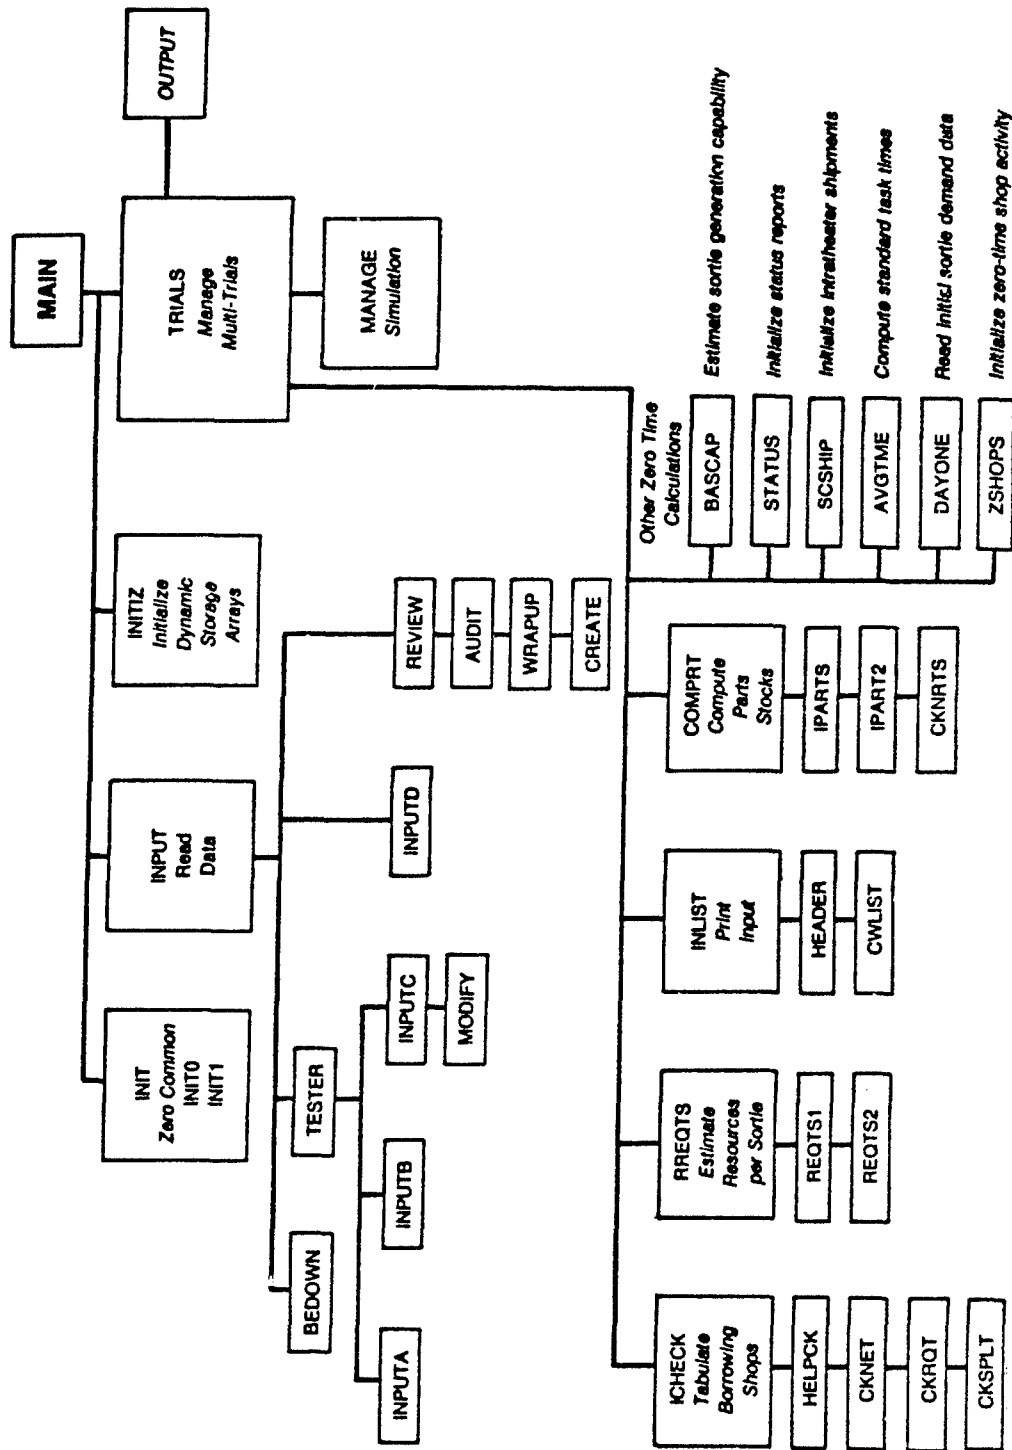


Fig. 2--TSAR input and initialization

The order in which the ten sets of events are examined is as follows:

1. Civil engineering reconstruction job completion times
2. On-equipment aircraft maintenance completion times
3. Parts repair and equipment repair completion times
4. Periodic and scheduled events
5. Flight delay completion times
6. Maintenance initiation when postflight delays end
7. Launch flights
8. Munitions assembly task completion times
9. Release resting personnel
10. Resource resupply arrival times

If two or more of these events occur at the same simulated time, they are processed in the order listed. The order chosen tends to limit the processing requirements. If two events of the same type occur at the same time, there is no way to determine which will be processed first when they are stored in a heap (see discussion below).

The nature of these events varies substantially; all except the fourth and seventh are groupings of event completion times for similar types of events that occur irregularly. The fourth set—the periodic and scheduled events—is a heterogeneous set of events and simulation housekeeping tasks; the seventh stores the times when groups of aircraft (flights) should be launched on various missions.

Each of the main tasks and the scheduled and periodic activities is performed by a cluster of subroutines supported by a set of storage arrays. Although there is substantial interaction between these subroutine clusters, much of the discussion in this volume will examine one group at a time, noting the interactions in passing.¹ Figures 3 and 4 show the clusters of subroutines used to control irregularly occurring airbase activities, for example:

- To launch aircraft or to process them when they land and maintenance is defined.

¹The names and functions of all the subroutines shown on the figures in this section can be found in App. A of Vol. III.

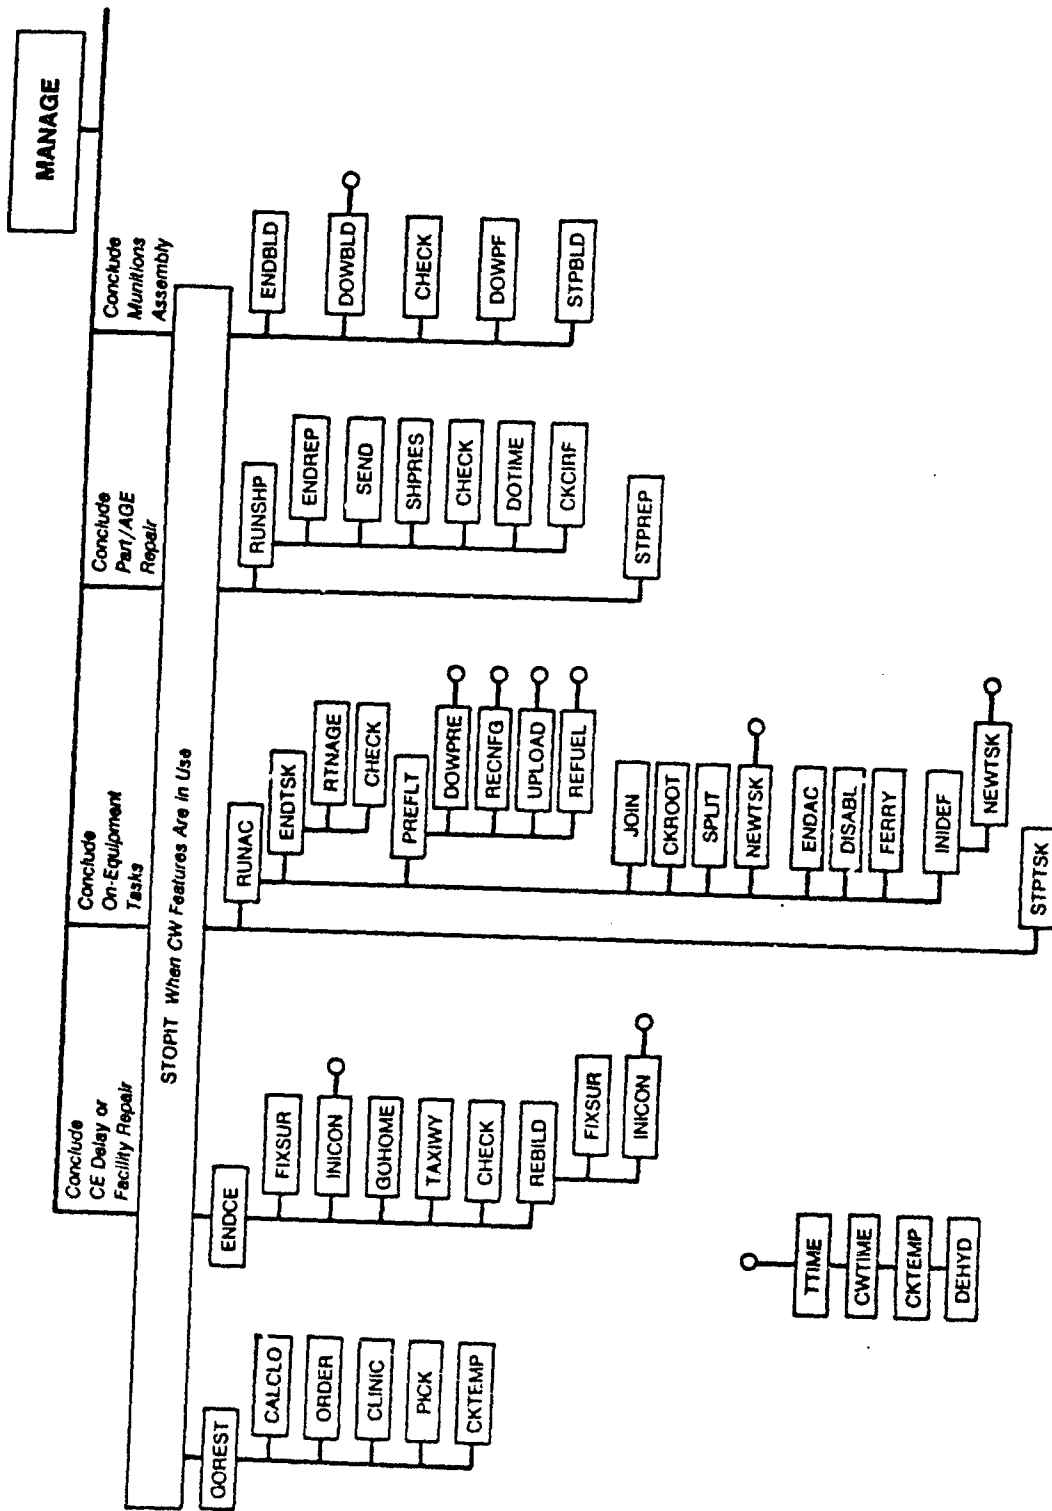


Fig. 3—Subroutines used to conclude tasks




Fig. 4—Subroutines used to manage irregularly occurring airbase activities

- To release personnel who are resting.
- To receive shipments, to conclude postattack delays, and to release resources when on-equipment tasks, parts and equipment repairs, munitions assembly jobs, and facility repairs are completed. As Fig. 3 shows, subroutine STOPIT is called first for these last four activities when the chemical warfare features are being used; assisted by subroutine GOREST et al., subroutine STOPIT checks the condition of the work teams before they are released.

Figure 5 shows the subroutine clusters used to process the scheduled and periodic activities necessary during the simulation.

Critical to the model is the handling of losses due to airbase attack. Subroutine BOMB, shown in Fig. 6, controls the various computations needed to account for these losses. As Fig. 6 shows, airbase attacks require a substantial number of actions to fully assess the various effects on all resources and activities. The computations that are carried out at the time of an attack start at the left on Fig. 6 and proceed to the right.

3. OUTPUT

A great variety of data describing the status of the ongoing simulation can be printed. The user's choices among the output options define which of these many possible results are actually listed, as outlined in Sec. XV. Some results are generated periodically during each simulated day in subroutine MANAGE, and others are printed at the end of each day by subroutine OUTPUT, which is called by MANAGE.

Other results include statistical data that define which resource constraints caused delays in on-equipment maintenance and backshop repairs; these are printed at a user-specified frequency by subroutines DELAYS and TIMES. At the end of each trial, subroutine SUMUP provides a summary of the aircraft sorties flown and the work that was accomplished during the trial. And when the specified number of trials of the simulation have been completed, subroutine TRIALS calls upon subroutines SUMUP and SUMMRY to provide average results based on the several trials. Section XV explains the output options in more detail, and most are illustrated in Sec. XXI in Vol. II with the results of a sample problem.

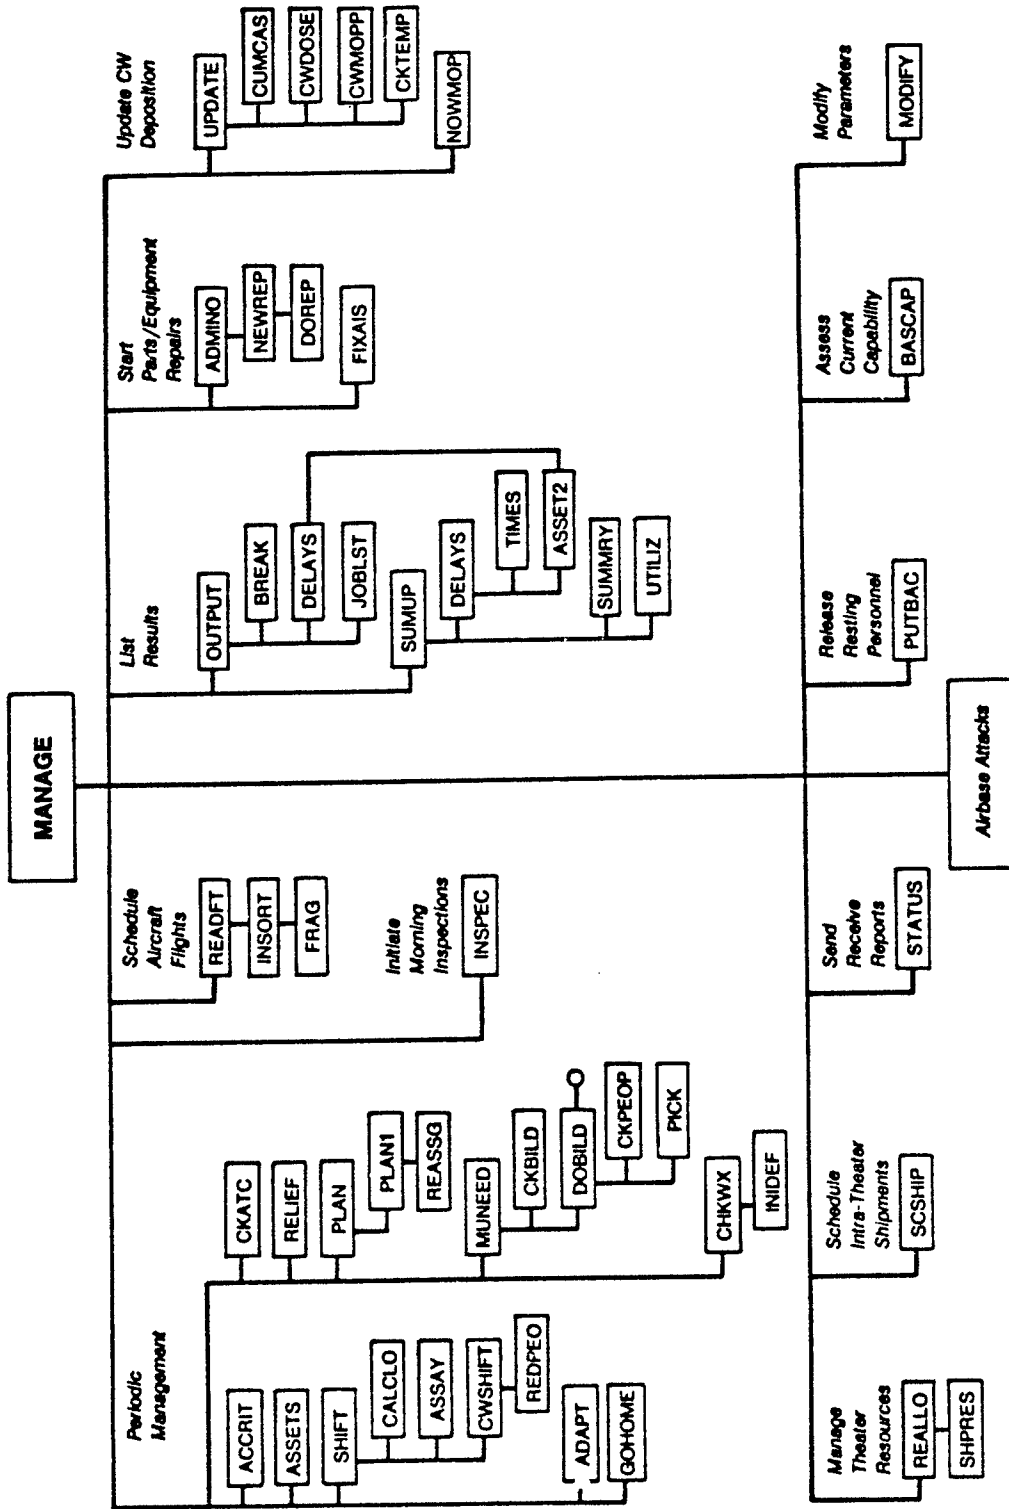


Fig. 5—Subroutines used to manage scheduled and periodic activities

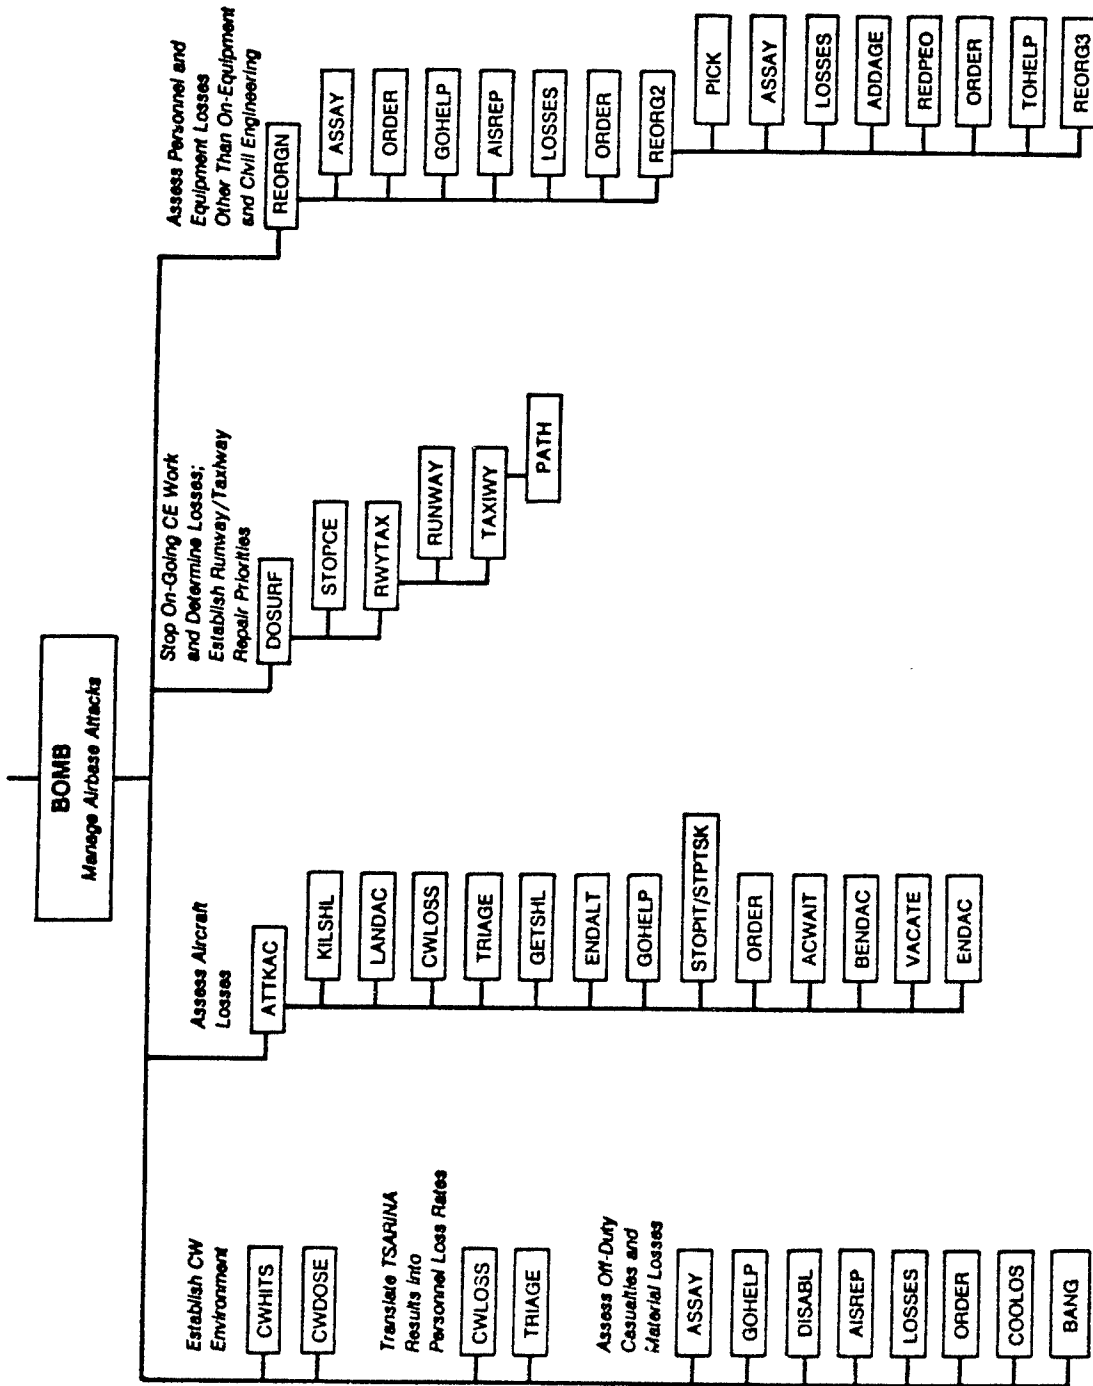


Fig. 6—Subroutines used for airbase attack assessments

4. DATA STORAGE

Each of the ten event sets is maintained in a distinct data storage array that also stores other information that must be associated with the event. For seven of the sets, the data are organized into what has been called a "heap," which is a data structure that permits more efficient processing than a time-ordered queue when only the next earliest event must be known; the flight demand data are queued in the order in which the flights will be demanded. In one instance—the periodic and scheduled events—several of the elements in that event set are themselves the earliest elements from several subordinate "heaps" and time-ordered queues.

In many instances it is not possible to initiate events as soon as they are defined or to pursue them without interruption, so it is necessary to store the relevant information until the resources required to pursue the tasks are available. Aircraft maintenance tasks, parts repair jobs, and equipment repairs are stored in special storage arrays² if they must wait or when they are interrupted (i.e., the WAITSK and INTTSK arrays); the munitions assembly tasks are stored in the BACKLG array when resources are not available for their initiation and in the INTTSK array when they must be interrupted.

The tasks that relate to each aircraft and to each of the work centers or shops that will perform the work are tied together with a system of pointers (or storage location addresses). Each aircraft and each shop at each base maintains pointers to the first and the last of each of these sets of events (in the ACN and SHOPS arrays), and the several events in the storage arrays that are associated with a particular aircraft and with a particular shop are themselves interconnected with a system of pointers. With these pointers the activities associated with any particular aircraft or shop that are ongoing, waiting, interrupted, or deferred can be readily located by examining a short trail of pointers.

The earliest times for each of the periodic and scheduled events are stored as a heap in the array PERIOD, which is maintained in subroutine MANAGE. At this time there are 21 sets of these events as shown in Table 1.

The frequency of some of these activities is predetermined by the program. Other schedules are specified by user input—for example, the transmission of reports, the updating of the chemical warfare (CW) deposition estimates, the modification of control parameters, and the occurrence of airbase attacks. Whenever the processing associated

²All data storage arrays are defined in App. C of Vol. III.

Table 1

EVENTS STORED IN THE PERIOD HEAP

Heap Position		Activity
1	Periodic — even hours	Change shifts for ground personnel Relieve aircrews Project aircraft supply and demand Reassign aircraft missions Check munitions requirements and initiate assembly Initiate deferred maintenance as allowed List stocks of parts, munitions, and POL (every 6 hours)
2	Scheduled — N days	Read new flight demand data and regenerate the demand queue
3	Scheduled — N days	Regenerate flight demand data queue
4	Scheduled	Initiate early morning aircraft inspection
5	Periodic — N hours	Redistribute theater resources
6	Periodic — N days	Regenerate intratheater shipping schedule heap
7	Scheduled (queue)	Receive intertheater shipments
8	Scheduled (heap)	Initiate intratheater shipments
9	Scheduled (heap)	Receive intratheater shipments
10	Scheduled (heap)	Send and receive intratheater resource status reports
11	Periodic — Hourly	List numbers of aircraft waiting by shop, numbers of aircraft with "holes"
12	Periodic — even hours	Conclude administrative delays and process faulty parts for repair
13	Periodic — Daily	Estimate sortie generation capabilities for next 24 hours
14	Scheduled (heap)	Modify TSAR operating characteristics at a previously specified time
15	Scheduled (heap)	Airbase attacks
17	Scheduled (heap)	Release personnel performing buddy care
18	Periodic (specified)	Update of chemical contamination
19	Periodic (odd hours)	Store personnel availability data
20	Scheduled	Activate aircraft transfer directives
21	Scheduled	Assess UXO detonations
22	Periodic (12 hours)	Reprioritize reparable (0530 and 1730)

with any one of these activities is completed, the next earliest activity rises to the top of the PERIOD heap and is considered in concert with the nine other basic sets of events.

A full description of the TSAR simulation entails (1) a complete description of the steps followed for each of the sets of events; (2) specification of the algorithms used to decide when follow-on actions are initiated, and (3) disposition of the various resources that are being accounted for in the simulation. These descriptions and specifications are introduced in Secs. IV through XI; Secs. XII through XV provide an introductory discussion of input-output procedures and other aspects of simulation management.

The following descriptions treat all of TSAR's features and operating modes. But the reader should be aware that TSAR can function usefully in many less complex modes, when that is appropriate. A great many of the features can be dispensed with by simply not entering the pertinent data. At its least complex, TSAR would function with one aircraft, one airbase, one mission, a flight duration, a turnaround time, and a single periodic sortie demand. No resource other than the aircraft would need to be identified.

IV. ON-EQUIPMENT AIRCRAFT MAINTENANCE

The only constraints on the continuous recycling of aircraft in wartime are the requirements for adequate launching surfaces; the availability of aircrews, munitions, and fuel; and the necessary maintenance to permit the aircraft to fly militarily useful sorties. Of these constraints the last is clearly the most complicated since it involves complex interdependencies among a variety of resources. Without maintenance constraints, estimation of an airbase's sortie potential would be straightforward and would require little or no complex analysis. A basic reason for the level of detail in TSAR's formulation was to gain an understanding of the effect of high levels of sortie demand and battle damage on the complex processes that are needed to ready aircraft for combat and that depend on several other actions and resources.

On-equipment aircraft maintenance activities can be divided into scheduled and unscheduled tasks. The *scheduled* requirements include (1) phased (or periodic) maintenance performed at specified intervals of flying time; (2) certain essential ground tasks, including refueling, early morning inspections, mission-dependent postflight inspections, through-flight inspections, and possibly aircraft decontamination; (3) reloading basic munitions, and (4) loading mission-dependent munitions before each flight. TSAR permits each of these types of scheduled maintenance to be simulated, although phased maintenance may be postponed during the more critical phases of conflict, at the user's discretion. When simulated, phased maintenance is performed at night after the hourly flight intervals that are defined by the user for each type of aircraft.

The other problems that develop and demand attention constitute *unscheduled maintenance*. Within TSAR, unscheduled maintenance tasks develop at random or are generated in battle; such tasks can be categorized as required or deferrable, on a mission-by-mission basis. Deferrable tasks may be completed after some number of sorties or before the next day's flying, or they may be deferred indefinitely if mission requirements do not require their completion and the necessary time is not available. When aircraft are operating from austere, dispersed operating locations, many tasks may have to be accomplished at the aircraft's host base; such aircraft may need to be ferried to the host after recovery at a DOB or they may recover directly at the host. Other tasks may require that the aircraft be ferried to a major support base, presumably located

further to the rear. Some unscheduled maintenance tasks may require a check flight when they have been completed to validate the adequacy of the maintenance action.

Each of the several types of aircraft on-equipment maintenance tasks that may be simulated with TSAR are listed in Table 2. Two key items of information are listed for each task type: (1) how TSAR input formats are used to indicate that a particular type of task is to be simulated, and (2) how TSAR determines whether a particular task is required when a flight is complete. The remainder of this section will expand and explain these entries and their distinctions in considerable detail; without a clear understanding of these relationships, a user will be unable to employ the TSAR simulation effectively. The numerous data input formats used with TSAR and explained at length in Sec. XIX of Vol. II use numbered Card Types, each with its unique interpretation. For convenience, these Card Types are referred to in Table 2, as well as throughout this *User's Manual*, by a special notation. For example, CT5 and CT17/9 are used to refer to Card Type #5 and Card Type #17, subtype 9. Although these format references may prove invaluable in actually using TSAR, they are best ignored when this volume is first reviewed.

1. TASK REPRESENTATION

TSAR permits the user to define the requirements for each on-equipment maintenance task (other than loading mission-dependent munitions) as either a one-step procedure, a multistep network of subtasks, or a sequence of multistep task networks. The user also specifies the probability that a problem arises that generates each unscheduled maintenance task during a single sortie and, for aircraft that are to be operated from dispersed operating bases, the probability that the task is not detected before the aircraft lands. The requirements in the one-step procedure—i.e., a simple task—may include a number of each of two types of personnel, one or two pieces of support equipment, a part,¹ an undamaged shop, and an amount of time (specified by a mean and distribution if desired). More complex tasks that involve differing groups of

¹When a part is used in more than one location on an aircraft (e.g., a left and right tire), the user assigns a different part number to each location, and identifies that all such part numbers refer to the same entity (with CT35/4). The part in one of these locations is identified as the "prime" part, and is the one considered for procurement, repair, and distribution activities.

Table 2

ON-EQUIPMENT MAINTENANCE ACTIVITIES


On-Equipment Task Type	Task Specification	Basis for Selection
Postflight inspection, aircraft sheltering, through-flight inspection, fuel tank loading, and other scheduled tasks	Task entered in shop/task sequence (CT29)	Task probability defined on CT5
Mission dependent postflight inspection	Task number entered for each aircraft type and each mission on CT15/5	Scheduled by location of Shop 25 on CT29; probability defined on CT5
Phased maintenance	Inspection frequency and task number on CT15/4 and DOPHAS control on CT17/3	Required at night when cumulative aircraft flight time exceeds hourly inspection periods, and DOPHAS = 1
Early morning daily inspection	Task defined on CT15/5; time on CT17/3	Whenever specified
Aircraft decontamination	Task defined on CT15/3	Whenever there is on-base decontamination or the special decontamination switch (CT17/9) is on
Basic (regular) munitions uploading	Task number(s) in shop/task sequence (CT29), munitions on CT15/1, and task on CT5	Retention probability defined on CT16
Mission-dependent munitions uploading	Loadings (SCLs) and configuration defined with CT12, CT13, and CT14	Retention probability defined on CT16
Refueling		
In shelter	Task defined on CT15/1	After every sortie unless refueled at hot-pit
Hot-pit	Task defined on CT15/3	Code 1 and Code 2 aircraft are refueled at hot-pit hydrant when no wait is required
Unscheduled maintenance		
Frequent	Task entered in shop/task sequence (CT29)	Task probability defined on CT5
Infrequent	Task entered on CT7, and shop listed in shop/task sequence (CT29)	Task probability defined on CT7 entry
Battle damage tasks	Task 30000 entered in shop/task sequence (CT29), and task lists a CT15/2 and CT15/3	Damaged aircraft selection based on attrition rate and damage/kill ratio (CT16)
Check-flight to validate maintenance action	Tasks that can require flight on CT15/88	Probability flight not required on CT15/88 for aircraft types flagged on CT15/5

personnel, equipment, and parts are represented with a task network, or sequence of networks.

To portray these options graphically let us represent a simple task, or root segment, as:

P	Personnel	S
A	AGE	H
R	Time	O
T		P
	Probability	
	Likelihood undetected	


and the other segments of a task network as:



The option of specifying that a particular facility (shop) is available (undamaged) is unique with the first task segment and is therefore not shown as a property of the subsequent segments of a network. Using these block figures, on-equipment maintenance tasks may be represented either as a simple task



or as a complex network of subtasks:




The first task segment in a task network is referred to in TSAR as the *root segment*, and it is this segment that must be referred to when the work represented by the network is to be specified.

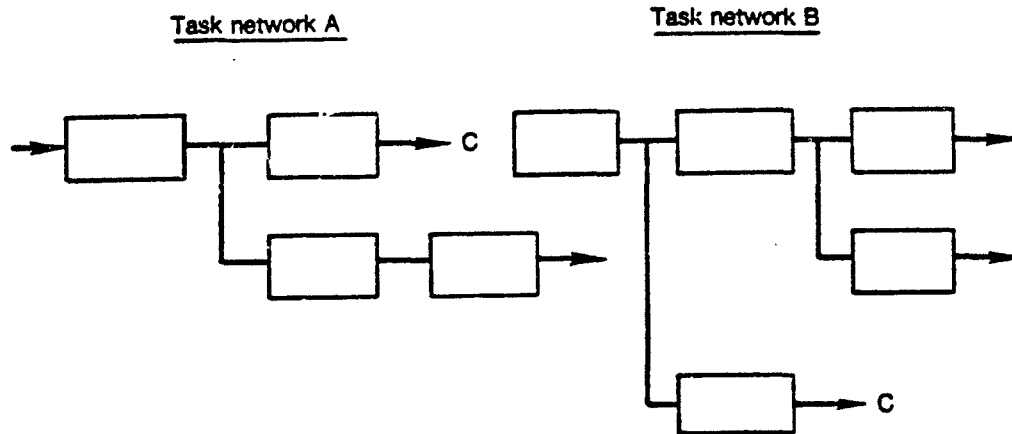
In the network shown above, when the initial, or root segment, portion of the overall task has been completed, three other subtasks are specified for follow-on activity, followed by yet other activities. The three follow-on activities may all be required, or they may each be required on a probabilistic basis. When a task element is found to be unnecessary, subsequent tasks are not considered. Parallel segments may all be required or may be defined as being mutually exclusive. If two or more of such parallel paths of activity must be completed before yet another follow-on activity is initiated, this can be represented by those paths rejoining before that activity. Furthermore, it is permissible to represent nested sets of parallel paths that rejoin as illustrated. However, all paths that split and ultimately rejoin must all rejoin at the same place.

The segments of a task network are initiated whenever the resources for the segment are available, without reference to the availability of resources for other segments, unless the "incompatibility" conditions for the segment (see Sec. XIX, CT19) prohibit task initiation. Although TSAR cannot require the time-coincidence of two or more parallel segments of a network that may actually be performed simultaneously in real life, it can require that all of the segments be complete before any follow-on action is begun, as was illustrated above.


The task segment that immediately follows a segment that includes a part may be made conditional on whether the part was required; thus, in the following networks the task T and the tasks X, Y, and Z may be made conditional on a part having been required in segments A and B.



Task specification and storage may be somewhat simplified when the work procedures associated with several paths have common elements. In schematic terms the two tasks A and B can be defined as:



where the C segment



is common to both tasks.

To be able to represent a situation that sometimes occurs in the field, any segment of a network may also specify the root segment of another network as a subsequent task; this simulates the situation where, after work is accomplished on one job, it is discovered that the actual problem is different than initially thought. The only caution to be observed when task networks are "chained" in this manner is that no two networks may each point to the other, either directly or through intermediate chained networks; otherwise an infinite work loop could be created.


The user may also examine the situation in which certain specialists at specified bases have received cross-training so as to be able to assist or replace another specialist on a specified subset of the latter's normal activities. Cross-trained personnel are assumed to perform the tasks for which they are qualified in the same time as the specialists who normally do those tasks.

The user may also specify alternative sets of personnel and equipment for any of the segments of a task, and these alternatives will be considered whenever insufficient resources are available to perform the task with the normal procedures. If available, the task is done with the alternative resources, without reference to the subsequent

availability of the normal resources. There may be as many alternative sets of personnel, equipment, and time specified for each task segment as the user's knowledge and available data permit.

2. TASK SEQUENCING

The organization and sequencing of the various tasks and task networks that are required on each aircraft after it lands and before it can be launched on another sortie are fully controllable² by the user for each aircraft type and for each airbase. Some tasks may be pursued simultaneously, some may have to be done in a specified order, and others may occur in any order, but not at the same time. These options can be illustrated as follows:




In this instance, Task T₁ is done first; Tasks T₂, T₃, and T₄ are done next, as available resources permit, except that if tasks T₂ and T₃ are both required, they may not be done simultaneously, and if tasks T₃ and T₄ are both required, T₄ must be completed before T₃ may begin. Then, when these tasks are all completed, tasks T₅ and T₆ may begin; the aircraft may be launched when all tasks are completed. Any or all of these tasks actually may be a task-network that must be completed before the task can be considered to be complete. Furthermore, each may occur only with a specified probability. If the aircraft has received battle damage, that work will be done at the point where Task 30000 is specified.

The unscheduled on-equipment aircraft tasks are normally grouped together with the other tasks performed by the same work center or shop for reasons to be described

²Except in the special circumstance when an aircraft must be ferried to a rear base for some portion of the required maintenance (see Sec. IV-8, below).

shortly. Reference to these collections of on-equipment aircraft maintenance tasks simplifies the specification of task organization as illustrated in the following sketch.



Here S_i , S_j , and S_k are the collections³ of unscheduled on-equipment tasks associated with shops i , j , and k . Following an aircraft sortie, each collection is checked to see whether any task associated with that shop is required. Because the majority of the unscheduled on-equipment aircraft maintenance tasks are individually low probability events, TSAR groups together those tasks performed by the same work center or shop and selects at most one following each flight. Processing is speeded up by this simplification (of at most one task per shop per sortie) without a serious loss in fidelity, because the joint occurrence of two or more individually low probability events would be quite unlikely. If some of the tasks associated with particular shops must be carried out frequently (i.e., are not low probability events), they should *not* be included in the shop task collections but should be treated as the "other on-equipment tasks" discussed below.

3. TASK DESCRIPTIONS

As summarized in Table 2, on-equipment aircraft maintenance tasks fall into several categories. With the exception of the mission-dependent preflight tasks (to be discussed in Sec. V), the data defining the personnel, equipment, parts, munitions, TRAP, and time required for each task (and for each segment of a task network), along with data defining the network structure and parts requirements and the shop to which it is assigned, are stored in the TSKRQT array (CT5). If special damage repair personnel (RAM teams) are to be used for repair of battle-damaged aircraft, that requirement can be imposed simply by identifying such personnel as a unique type.

³Up to NOTASK tasks may be grouped together in each of these collections.

The locations of the input data defining the likelihood of the various kinds of on-equipment tasks are outlined in Table 2. For most scheduled on-equipment tasks, such as Tasks #41, #42, and #43 in the last sketch, and the more frequent unscheduled tasks, the probability that the task must be performed when a sortie has been completed is also stored in the TSKRQT array. However, refueling is mandatory, early morning inspections are required each day, and aircraft decontamination depends on the level of airbase chemical contamination when the aircraft lands. The requirement for phased maintenance depends on an aircraft's cumulative flight time, and the need for reloading basic munitions is determined by the probability that these weapons were not expended on the previous mission (CT16).

The numerous low probability unscheduled on-equipment maintenance tasks are handled somewhat differently. The likelihood that each of these tasks will be required after a sortie has been flown is specified with CT7, and these data for the tasks in each work center or shop are collected together in the SHPRQT (shop requirement) array. If desired, the breakrate data in these shop collections may be varied statistically from the input values for use in the simulation—to represent uncertain wartime breakrates—or they may be varied with aircraft sortie rate for specified shops and types of aircraft (CT18/2).

The aircraft repair requirements imposed by battle damage are handled in a unique manner. Following each mission, a random number is compared with the probability that that particular type of aircraft will be damaged on that type of mission (as specified by the user using CT16). For aircraft that are determined to be damaged, each of the list of battle-damage tasks specified for that aircraft type on CT15/2, or for the mission type on CT15/3, is checked. The likelihood that each battle damage task is required is specified by the task probability in the TSKRQT (task requirements) array (entered with CT5); these tasks are treated as mutually exclusive when the task probability is negative. Those battle damage tasks that cannot be deferred will be conducted at the point in the task sequence where Task #30000 is placed. Aircraft repair requirements imposed by damage inflicted during airbase attacks are handled in a similar fashion, except that the tasks are added to whatever other on-equipment work is going on at the time of the attack.

Phased maintenance tasks are carried out whenever the total flying hours for an aircraft exceed any of the phased maintenance time thresholds, unless the user has

specified that these maintenance requirements be ignored for certain times during the simulation; such work is carried out at night, as with deferred maintenance.

With only four exceptions, the requirements for on-equipment aircraft tasks are treated as independent activities. Two of the exceptions concern support equipment, and another involves munitions load crews. For each aircraft type, the user may specify (on CT15/1) one or two types of support equipment for which multiple demands can be satisfied with a single item. The auxiliary power cart and the hydraulic mule might be treated in this manner. The user may also prevent a single aircraft from being assigned more than one munitions load crew by specifying the type and number of personnel that make up a load crew on the CT15/1. Another quite different exception to treating tasks as independent is discussed below in Sec. IV-7.


4. SHOPS AND UNSCHEDULED TASK COLLECTIONS

TSAR provides for a total of 30 shops. All aircraft maintenance personnel, equipment, and parts are "assigned" to one or another of these shops, and TSAR maintains lists of the tasks and repairs that are underway, interrupted, waiting, and deferred separately for each shop. The first 24 shops are to be used for those collections of unscheduled maintenance tasks performed by specialists associated with each of the aircraft maintenance work centers. Generally, the personnel, equipment, and parts used in connection with the tasks assigned to a particular shop (work center) should also be assigned to the same shop, although in the case of personnel and equipment they may be "borrowed" for work on a task assigned to another shop. Thus, the resources (personnel, equipment, parts) associated with unscheduled tasks will be assigned to one of the first 24 shops, most scheduled flight line resources will be assigned to Shop #25, and the munitions-related resources to Shops #27 and #28.

If desired, the personnel and equipment of each shop may also be assigned (with CT45/1) to 1, 2, or 3 separate groups (or AMUs) to support separate subsets of aircraft, and to a wing-level organization for backshop parts repair (CRS and/or EMS) as in the COMO (AFR66-5) maintenance concept. The user should reserve Shops #27, #28, and #29 for the preflight tasks, that is for reconfiguration, munitions loading, and refueling, respectively, as outlined in Sec. VI. Shop #25, the "flight line" shop, is to be used for those tasks other than the preflight tasks that involve munitions and TRAP resources; Shop #25 should also be used for user-specified mission-dependent postflight inspections.

(Shop #26 is used by the program for storing references to aircraft whose mission assignment and weapons loading has been delayed, and Shop #30 is not associated with aircraft maintenance; it is used in connection with munitions assembly and civil engineering.)

In practice, there is only a limited likelihood that the specialists from any given shop will be required for an unscheduled maintenance task after any particular flight, and a much smaller chance that they will be required for two or more distinct tasks from the same shop, so the TSAR data structure for the shop-task collections has been designed such that at most one task from each collection will be selected after a particular sortie. With this restriction, only a single random number need be drawn and compared with the cumulative sum of the probabilities of the several tasks in each collection. If the number is greater than the sum, no task is required; if it is less, the task is determined by the random number's position within the set of cumulative probabilities. This process may be visualized as follows:



In the situation shown, there are only three possible unscheduled on-equipment tasks that the shop may be called upon to perform: T_1 , T_j , and T_k . The probabilities that the individual tasks (entered with CT7) are required after any given sortie are P_1 , P_j , and P_k . After each sortie a random number, RN, is drawn for each shop. If the value corresponds to the shaded region for a shop, no task is required; if the value is less, the task to be accomplished is the one corresponding to RN. When the user has specified that the nominal task probabilities, or breakrates, for certain shops and aircraft types should be modified in some way (as controlled by CT18/2 data), the random number is adjusted before being compared with the shaded region. Also, as will be explained

below in Sec. IV-7, when the user has specified that only a limited percentage of the aircraft land with unscheduled maintenance, the random number is adjusted such that the overall maintenance demands are preserved. As noted earlier, tasks with a fairly high probability of being required following a flight should not be included in these task collections.

5. ILLUSTRATIVE SHOP-TASK SEQUENCE

The various features for representing the organization and processing of aircraft maintenance tasks will permit the user to rapidly define and test a wide variety of different base maintenance structures. An example of an actual structure that might be defined (with CT29) is shown in Fig. 7.

Immediately upon landing, the user may specify either a taxi time and/or a specific postflight delay to account for taxiing, debriefing, etc. If "hot-" refueling or decontamination tasks have been specified and one is selected, it is performed immediately after any taxi-time specified. The aircraft is then presumed to move to an aircraft shelter or parking ramp, where the remainder of the maintenance work is carried out. In the example, it is specified that Task #45, removing hung ordnance, occurs after 4 percent of the sorties. When that is completed, any unscheduled maintenance that is required by Shops #1, #17, #19, #2, #4, #9, and #24 may be initiated. TSAR determines which tasks are required, which tasks may be deferred for the next mission, and what the tentative mission assignment should be for the aircraft when it first lands. Two scheduled tasks as well as any battle damage work are also specified in the example: The requirement to reload guns, Task #62, is controlled by the CT16 entry for the basic-munitions retention probability; the task to hang fuel tanks, Task #47, is not mission dependent, a CT5 specifies that it must be accomplished after 60 percent of the sorties. These tasks are different in character from most other tasks, in that some of the required resources are consumed. Tasks that may consume TRAP or munitions *must* be associated with the special "flight line" Shop #25 when they are not part of the mission-dependent munitions activities, or refueling activity in Shops #28 and #29. The battle damage tasks are implied by specifying Task #30000.

When all of the first set of possible tasks has been completed, shop activity by Shops #8, #3, #21, and #12 may begin, along with Task #51. And when those jobs have been completed, the preflight preparations may begin. These preparations, discussed at




Fig. 7—Structural representation of on-equipment maintenance

length in Sec. VI, may involve a possible delay (CT15/1), followed by the final mission determination, aircraft reconfiguration if required, loading of the mission-dependent munitions, and refueling. As indicated, the delay, reconfiguration, and uploading always occur in sequence and are specified by entering Shop #26 in the shop-task sequence on the CT29 cards. Task #52 is also indicated as being concurrent with the preflight preparations. This task as well as the other individual tasks must themselves be associated with some shop (\leq #25) and use resources assigned to one or another of the shops. The munitions and TRAP tasks that *must be* placed in Shop #25 probably would use personnel and support equipment from Shops #27 and #28, while the other tasks could call on resources normally associated with any of the first 25 shops.⁴ Mission-dependent postflight inspection tasks, that may be imposed using CT15/5, should also be assigned to Shop #25, and will be scheduled where Shop #25 appears in the shop-task sequence.

To specify the shop-task sequence that is to be followed, the user enters a string of numbers using CT29; a different string may be entered for each type of aircraft and for each airbase. These data are stored in the SHPCRD (shop order) array. As the reader may have noticed, individual tasks must always be identified with a number larger than 30 so that they will be distinguished from a shop. Furthermore, Task #30000 is to be used to designate when the battle damage tasks are to be performed. For this example, the shop-task sequence is entered by listing the numbers 45, 0, 30000, 62, 47, 1, 17, 19, 2, 4, 9, 24, 0, 8, 51, 3, 21, 12, 0, 26, 52, 29, 0, 0 in order on CT29 cards. Note that each set of activities that may be carried out concurrently is followed by a zero to designate the end of the set.

Whenever any of the required maintenance tasks is one of those that must be accomplished at a rear base, or whenever unscheduled aircraft maintenance is estimated to take longer than a user-specified length of time, the user-specified task sequence is replaced by three sequential sets of maintenance tasks. The first set, to be accomplished at the operating base, includes refueling and all tasks that would prevent the aircraft from being ferried (i.e., task criticality of 33). The second set includes refueling at the rear base, all tasks that must be performed at the rear base, and other of the aircraft's required

⁴Because of the logic used for checking on tasks that are waiting and that may need the resources that are being released from a previous task, the munitions and TRAP-related jobs—those that use personnel or equipment from Shops #27, #28, or #29—must be associated with Shop #25.

and deferred tasks as are specified by the variable JOBCON (job control) on CT3/2. The third task set, those that are to be done when the aircraft returns to the operating base, includes all remaining tasks that are required.

6. TSAR DETERMINATION OF AIRCRAFT MAINTENANCE REQUIREMENTS

Whenever an aircraft completes a flight (and has been removed from the delay heap in the ACN array), subroutine MANAGE transfers control to entry point LAND in subroutine FERRY, where checks are made to see whether the aircraft was lost on the mission or has received battle damage. Subroutine LANDIT then is called to check whether runway conditions permit recovery and if not what alternate should be used; for aircraft operating from DOBs, subroutine LANDIT also identifies the host base, and an alternate if the host's runway is closed. Subroutine FERRY next calls CKMAIN to determine what maintenance will be required, and whether a DOB aircraft should recover at its host. Subroutine CKMAIN establishes which of the individual tasks and which of the collections of individually low probability unscheduled tasks are required. The tasks and shop task collections are checked in the specified order as illustrated in the preceding subsection; a random number is drawn for each task and shop to determine which if any of the tasks require attention. The malfunctions noted are stored in a temporary array that is processed in subroutine PSTFLT after the aircraft has recovered. Aircraft operating from a MOB or a COB recover as runway conditions permit. For aircraft operating from a DOB, a check is made to see if any of the work must be done at the host, and if that work has been detected. If none of that work has been detected, the aircraft will recover at the DOB and subsequently be ferried to the host; otherwise, it will fly directly to the host. If there are no requirements for work at the host, the aircraft will recover at the DOB to be readied for the next combat mission.

After the aircraft has recovered, the flight crew is released (see Sec. VII for a discussion of that process) and, if battle damage is so severe that repair is not practical, the aircraft is written off and the various parts are salvaged to the extent specified (see CT15/2). Otherwise, subroutine PSTFLT (postflight) is called. (For aircraft that have not survived or are salvaged, the existing records are eliminated with subroutine KILLAC and, if filler aircraft are available or if the user specifies replacements, another aircraft is requested using subroutine ORDER.)

The basic functions performed within subroutine PSTFLT are (1) to initiate any user-defined postflight delay (to account for taxiing, inspection, etc.); (2) to schedule hot-pit refueling, or aircraft decontamination, when appropriate; (3) to determine if any previously deferred tasks must be done at this time; (4) to estimate the expected time that will be required to carry out the maintenance that is essential for various missions; (5) to schedule aircraft refueling and transfer when any of the required maintenance must be accomplished at a host base or a rear base; (6) to establish a tentative mission assignment for the aircraft; and (7) to categorize the newly defined tasks as essential or deferrable for that mission.

Subroutine PSTFLT establishes what the expected time will be for carrying out the essential maintenance, including battle damage. The individual tasks determined in CKMAIN are checked in the appropriate order; the expected time for each task is estimated, based on the nominal task time specified for that task, plus approximate time allowances for (1) whatever inefficiencies are expected because of chemical warfare effects; (2) aircraft that are already waiting at the various shops; (3) parts repair, when parts are known to be required and none are available; and (4) repair of the maintenance facility itself, when the task specifications prescribe that facility and it is damaged.

When these time estimates for the required tasks and for the preflight tasks have been determined, the time at which the aircraft could be ready to fly is established for each possible mission type. If any of the previously deferred tasks are now required or are now essential for any of the missions, the time to accomplish them is included on the assumption that they would be processed simultaneously, before the other tasks are done. These ready-to-fly estimates take into account the user's specifications as to which shops may perform on-equipment tasks simultaneously and which groups of shops must follow other groups. They also take into account only those tasks that may not be deferred for each particular mission. By making the estimates in this manner, the nominal times at which the aircraft could be readied for the various missions will typically differ for the different missions, and these times will also include at least a rough accounting of the inefficiencies, queues, parts shortages, or facility damage that might interfere with the preparations for one mission, but not another.

Unless the aircraft is scheduled to be ferried to its host base or to the rear maintenance base, as discussed below, the next step is to determine the highest priority mission that has insufficient aircraft to meet the known demand, between the times that the aircraft could be ready for various missions and the time horizon for planning. If the

deficient priority is no higher than that for the aircraft's previous mission and occurs no earlier, the aircraft is committed to the same type of mission that it just completed. Otherwise, the aircraft is tentatively committed to that mission with the earliest, highest deficient priority.

The maintenance tasks that are essential for the designated mission are stored in the RQDTSK (required task) array and the others are placed in the DEFTSK (deferred task) array. Final bookkeeping in the PSTFLT subroutine includes updating the aircraft's criticality index (i.e., ACN(-,17)), which maintains a record of which missions may be flown despite the maintenance that has been deferred.

7. LIMITING UNSCHEDULED MAINTENANCE

Normally, TSAR assumes that the probability that unscheduled maintenance will be required on an aircraft after each sortie is determined only by the probabilities specified for each of the individual unscheduled tasks. Thus, the likelihood that any task is required is presumed to be independent of all other tasks. An alternative TSAR mode of operation permits the user to limit the fraction of aircraft that require unscheduled maintenance after a sortie, but also to maintain the total amount of work accomplished consistent with the overall task statistics, by increasing the amount of work that must be done on those aircraft that do experience a "break."

In the past, when TSAR has been run assuming task independence (using unscheduled maintenance data derived from data bases prepared for the Air Force LCOM model), it has been found that the percentage of aircraft that land and require unscheduled maintenance is substantially greater than that reported by flying units. Possible reasons for these discrepancies could include (1) the observed tendency for aircraft breakrates to be reduced at higher sortie rates, (2) an apparent tendency for LCOM data to reflect a conservative view of what must be repaired for an effective combat sortie, and (3) the strong possibility that unscheduled maintenance tasks are correlated, i.e., should not be treated as independent events. All of these possible mechanisms are poorly understood at this time; for TSAR simulations, one can modify the nominal breakrates by shop (with CT18/2 cards) to adjust for observations of the first kind. The user can also introduce a less conservative notion of how an actual wartime MESL (minimum essential subsystem list) would affect the maintenance requirements with his specification of task criticality (on the CT5 card). And if none of the

unscheduled maintenance tasks for an aircraft have a nominal task time in excess of GRACE (CT4/2), the recovery will be designated Code 1. Finally, the user can reflect the possibility that when one task network requires work, it may be more likely than implied by the independence assumption that other networks will also require work. This apparently not uncommon situation in which maintenance starts with one task but leads to another and another can be simulated with TSAR by "chaining" a series of task networks, but that procedure requires a more detailed knowledge of the interdependencies among tasks than exists at this time.

The user may also simulate that a disproportionate amount of the overall work load is incurred by a subset of the aircraft, by specifying (on CT15/3) a value for the total "percentage of aircraft that land with deferrable or nondeferrable unscheduled maintenance" (that have Code 2 or Code 3 "breaks"). When that is done, that percentage of the aircraft are selected at random in subroutine CKMAIN and are checked for unscheduled maintenance; the likelihood that unscheduled maintenance is required in each shop is increased in the proportion necessary so that the total amount of work required for a large number of sorties will be the same as when the tasks are treated as independent. The division of unscheduled maintenance between Code 2 and Code 3 is based on the several TSAR mechanisms for specifying task criticality and deferrability.

8. AIRCRAFT OPERATIONS AT DISPERSED OPERATING BASES

In addition to MOBs and COBs, one may also designate that certain airbases will be used in support of austere operations at "dispersed operating bases" (DOBs). Such bases are envisioned as having extremely limited capabilities for unscheduled maintenance; refueling and rearming would be the main functions, although the TSAR simulation allows the user complete freedom in defining "austerity" for DOBs. Since much if not all unscheduled maintenance will be required to be done at the DOB's host base (typically a MOB), it is determined before the aircraft recovers from a combat mission whether to return to the DOB or go directly to the host. Uncertainty may be introduced into the aircrew's ability to detect such conditions.

Each on-equipment task may be coded to indicate whether repair at a DOB is practical, or whether the aircraft should recover at a base with more extensive repair capabilities; the user also specifies which of the tasks that can't be handled at a DOB will not be detected by the pilot some percentage of the time (individually by task). Aircraft

that recover at a "host," either directly from combat, or after being ferried from their DOB, lose their association with the DOB and become a part of the "host's" complement of aircraft. However, whenever an aircraft at a "host" completes maintenance, the aircraft-transfer-directives (discussed below) are checked and sufficient aircraft are transferred to maintain the specified complement at the host's DOBs. If aircraft at a DOB are found to require work at a rear maintenance base (see the next section), they are ferried directly to that base, rather than the host.

In the evening, when flying tapers off and tasks that have been deferred during the day must be worked off, aircraft at the DOBs are ferried to their host base for the required maintenance (or to a rear maintenance base for specified tasks). Aircraft returning to a DOB from combat late in the day⁵ that don't have any immediate maintenance that requires they go to the rear but do have deferred maintenance that shortly will need attention, are flown directly to their host base. When aircraft at a DOB should be ferried to the host base, but the host has been closed by air attack, the needed maintenance will be further deferred and the aircraft will continue to operate from the DOB, unless ALTDEF (CT4/3) is unity; in the latter case, an alternate host will be sought that operates the same type of aircraft. Tallies are maintained of the cumulative numbers of aircraft transferred among MOBs, COBs, DOBs, and rear maintenance bases.

Aircraft operating from MOBs and COBs will recover at a DOB when no MOBs or COBs have open runways and when USEMER (CT4/3) is zero. Such aircraft receive maintenance and tasking while at the DOB, to the extent practical. If USEMER is unity, the EMERG base is used for recovery rather than a DOB.

To facilitate use of DOBs, TSAR permits the user to issue "aircraft transfer directives" that direct a particular number of aircraft of a specified type to be moved from one base to another, beginning at a designated time. Up to eight such directives may be outstanding for each base at the same time. Each directive specifies the time that the transfer is to begin, the type and number of aircraft, the base the aircraft are to be taken from, the base that they are to fly to, and for aircraft to be transferred to a DOB, the desired mission configuration.

Aircraft may be moved from a MOB to each of several DOBs whenever specified, and then directed to be flown back to the MOB at some later time. If aircraft are ready to be transferred when the directive is issued, they are transferred at that time; if a sufficient

⁵Whenever recovery at the DOB is to be as late as TBEFOR (CT4/3) TTU time units before ENDAY.

number are not ready at that time, the remainder are transferred as they complete maintenance and they become ready to fly. If aircraft are returned to a MOB from a DOB for maintenance, other aircraft will be transferred to maintain the designated number at the DOB. If over eight directives are issued for the same base, the directive that has been operative longest is replaced with the most recent.

A transfer directive may be issued that increases or decreases the number of aircraft to be maintained at a forward location. When such a directive reverses the direction of transfer between two particular bases for a particular aircraft type, the earlier directive is canceled and the later one is activated. Up to 50 (currently) such transfers may be specified, and the code has been structured so that such directives could also be issued later, during the simulation, if an adaptive logic is introduced to generate such commands.

The directives are initially entered into the MOVEAC heap during program initialization and are subsequently placed in the BASDIR array for the appropriate base at the time that they become effective. The actual transfers are initiated in subroutine SENDAC with a call to subroutine FERRY, for aircraft that are ready to transfer when the directive is issued, or in subroutine RUNAC when aircraft have completed maintenance. If the runways are not open, or pilots are not available when an aircraft is to be returned from a DOB, subroutine MANAGE calls GOHOME every two hours to locate and return any such aircraft when conditions are more favorable.

9. REAR MAINTENANCE BASE

Aircraft may be sent to a rear maintenance base either for specially designated tasks or when the estimated completion time for the required maintenance exceeds a user-specified time (MNTLMT on CT3/2). Both regular maintenance tasks and battle-damage tasks may be specially designated as requiring action at a rear maintenance base; these task designations may apply to all aircraft, or only to aircraft at a COB. Naturally, the criticality of any such task must be such that the aircraft may be ferried. If the user has specified a time limit for maintenance at the operating bases (MNTLMT) and the estimated maintenance time exceeds that value, a check is first made that there is at least as much time needed for the maintenance that would be accomplished in the rear as for what must be done at the operating base. If so, and if the time required for the maintenance that must be done at the operating base is as great as MNTF (CT3/2)

percent of MNTLMT, another check is made to see that the time for the maintenance that may be done in the rear is at least equal to MNTR percent of MNTLMT. This final check provides the user with a somewhat more realistic control over which aircraft are sent to the rear and which are not.

After an aircraft has landed and its ready-to-fly time has been estimated in subroutine PSTFLT a check is made to see if aircraft maintenance is to be done in the rear. If (1) tasks that must be done in the rear are outstanding, or (2) the ready-to-fly time exceeds MNTLMT, etc., the required tasks are regrouped into three sets. The first includes a refueling and all tasks that prohibit the aircraft from being ferried. The second set includes a refueling and all tasks that must be done in the rear, as well as whatever other tasks are to be done as defined by the value of the control variable JOBCON (see CT3/2 in Sec. XIX, Vol. II); these tasks are scheduled for the rear base. The third set of tasks includes a refueling and all other tasks that are outstanding and the necessary munitions loading tasks; these are scheduled for carrying out on return to the operating base. When the aircraft has flown directly to the rear maintenance base from a DOB, the aircraft is sent forward to the host base, not the DOB, when maintenance is complete at the rear base. No additional maintenance requirements are assumed to be generated on the return flight to the operating base. To the extent practical, the ordered structure for maintenance tasks that is prescribed in the shop-task sequence with the CT29 cards is maintained within each of the three task sets.

Aircraft spare parts for rear maintenance bases are either individually stocked or, when the automatic parts generation feature is being used, they are acquired by redistributing the spares that are calculated for the operating bases. For tasks that must be done in the rear, all parts are placed in the rear. An estimate is also made of the fraction of the other tasks that will be accomplished at the rear base at the same time that the mandatory work is underway, and a like fraction of all parts is placed in the rear. If aircraft are also sent to the rear whenever the ready-to-fly time exceeds MNTLMT, etc., the fraction of the parts placed in the rear can be increased by the user's specification of RPARTS (CT3/2).

10. AIRCRAFT MAINTENANCE MANAGEMENT

After an aircraft's next mission has been tentatively selected and the various scheduled and unscheduled tasks have been defined in subroutine PSTFLT, subroutine MANAGE transfers control to subroutine STARTM (start maintenance), which manages the initiation of on-equipment maintenance tasks.

When STARTM is called, following the optional postflight delay, TSAR immediately attempts to initiate the required work on each of the first set of required tasks stored in the RQDTSK array by calling subroutine INTSK (initiate task) through entry point NEWTSK. If a task is not incompatible with a task already underway, and the required resources are available to initiate the task, the resources are withdrawn from stock, and the task completion time is determined (using subroutine TTIME); the activity then is placed in the TASKQ heap. When the effects of chemical protection are to be taken into account (i.e., when USECW > 0), the task "heat factor" also passes to TTIME where it acts as a switch to call CWTIME that establishes how much longer the task will take in the chemical ensemble that is appropriate at the work place, and how long the work crew can be permitted to work before they must be allowed to cool off. If the requisite resources are not available, the task is placed in the WAITSK (waiting task) queue of the appropriate shop. (The operation of subroutine INTSK will be discussed more fully in the next subsection.) When all of the tasks that may be performed simultaneously have been processed, control is returned⁶ to MANAGE for other operations.

Subroutine RUNAC is called whenever an on-equipment task has been completed. If the effects of chemical protection are being considered, subroutine MANAGE first calls subroutine STOPIT, which uses subroutines GOREST and CKTEMP to determine how the work crew are to be disposed of. The first step in RUNAC is to call subroutine ENDTSK to release whatever resources⁷ are available, and to assign them to tasks that may have been interrupted or are waiting (by using subroutines CHECK and DOWPRE

⁶When late takeoffs are permitted, each aircraft is also checked to see whether its estimated ready-to-fly time is sufficiently close for it to be considered. If all tasks have been started and the estimated ready-to-fly time is within two hours, the flag—ACN(-,21)—is set so that the aircraft could be considered for a possible late takeoff. The flag is also set when only one task remains that has been initiated and is expected to be completed within three hours.

⁷Except for specific types of equipment that may need to be retained for use on other ongoing tasks.

for unscheduled and preflight tasks, respectively). When ENDTSK returns control to RUNAC, the next step depends upon the nature of the task. Except for preflight tasks, a check is made to see if the task is an element in a task network, and if it is, resources are checked in subroutine INTTSK to start any subsequent task or set of parallel tasks. A check is next made for any tasks that may have been forced to await the completion of the just concluded task, because of an incompatibility (as defined with CT19), and any such tasks are initiated if resources permit.

If at this point on-equipment tasks are in process, control is returned to MANAGE; if no tasks are in process but tasks are still waiting for the appropriate resources, a new estimate is made of the ready-to-fly time before returning control to MANAGE. If there are no tasks in process or waiting, but tasks remain in the RQDTSK array, a new estimate of the ready-to-fly time is computed, and STARTM is called where the next set of required tasks are checked and initiated as resources permit. If no tasks are in process or are waiting, and there are no further tasks required, a check is made to see if conditions permit deferred maintenance to be done at this time; operations in this circumstance are discussed in Sec. IV.15. If no deferred maintenance is to be accomplished, the aircraft is ready to fly and will be launched as required, except when it sustains a ground abort. In the latter instance a task is picked at random from the shop task collections and must be handled before the aircraft is again ready to fly.

When subroutine RUNAC is called at the completion of a preflight task, operation is somewhat different than with other maintenance tasks. After the resources that have been in use are released (and an attempt to reuse them made with subroutine DOWPRE), subroutine RUNAC calls subroutine PREFLT (preflight), which manages the unique task structure used with preflight tasks; these operations are described in Sec. V. When control is subsequently returned to subroutine RUNAC, processing continues in much the same manner as for unscheduled maintenance tasks, except for the task network tests.

One other important feature of the management operation performed in subroutine RUNAC permits the preflight tasks to be deferred in certain circumstances, so that the final decisions regarding mission assignment and munitions may be delayed until further information has been received regarding sortie demand. When these conditions (as discussed in Sec. V) have been met and the preflight delay flag DELYPF has been set to unity, the mission assignment and weapons loading tasks (i.e., Shops #26, #27, and #28) are allowed to wait while the other tasks are processed in accord with the specified

shop-task structure (i.e., the shop sequence data from CT29). When all required tasks are complete, deferred tasks will be initiated if it is estimated that they can be completed before the user-specified last allowable hour for commencing the weapons loading procedures (i.e., LSTTOD). If none can be started, or if all deferred tasks are completed before LOADTM (the earliest hour for commencing to upload munitions), a preflight delay is computed such that it will just be completed at LOADTM, and the aircraft is placed in the delay heap in the ACN array.

11. ON-EQUIPMENT TASK INITIATION

On-equipment maintenance tasks, except for the preflight tasks, are initiated with a call to subroutine INTTSK. This subroutine is initially called from STARTM or RUNAC; if tasks must wait or are interrupted after they are initiated, subroutine CHECK subsequently calls to recheck the availability of the required resources.

When subroutine INTTSK is entered to check for tasks that have been waiting or interrupted, a rough check is made of the existing ready-to-fly time estimate for the aircraft; if it is outdated a crude update is calculated. When a part will be required, base stocks are checked to see if one is available. The task is then checked to see if it must be delayed because of other work in process that is incompatible. If there is no problem, the program next checks for the availability of any facility that may be required and for the personnel and equipment specified for the task. If it has been specified that only one item of the required type of support equipment is needed for several tasks, and one is already assigned to the aircraft, the additional requirement is ignored; similarly, if an aircraft is not to be assigned two or more munitions load crews, and one is already at work, the task is delayed. If the aircraft is assigned to its own squadron with its own personnel and equipment, the required resources are drawn from the appropriate group. If a facility is needed and it is unavailable, or if insufficient personnel or equipment is available, the shortage is noted and the program transfers to check any alternative procedures that the user may have stipulated for this task.

Subroutine GETPEO is called to check on the availability of the required personnel, and subroutine CKAGE establishes the availability of any equipment that is required. If insufficient personnel are available, but on-base personnel have received cross-training, checks are made to see whether such personnel can be used on this task, and, if so, subroutine CKPEOP is called to see if sufficient cross-trained or task-assist-

qualified personnel are available. If there are not, and if the base has not been subjected to a chemical attack, a check is made to see if the required number of specified personnel are involved in parts repair activities; if they are, those repairs are interrupted to acquire the personnel needed for the on-equipment task,⁸ when the needed part is in stock. The time remaining to complete the interrupted repair is stored with the other repair data in the INTTSK array. If the required maintenance specialists cannot be obtained by these procedures, and if chemical protection ensembles are not being worn (i.e., USECW = 0), the last option is to stop an ongoing maintenance task on another aircraft. This will be done only if the ongoing task has at least two hours remaining until completion and if the aircraft has a projected ready-to-fly time at least four hours later than the aircraft for which the personnel are sought.

If sufficient personnel are available but a needed part is not, a check is made to see if it may be obtained from another aircraft by cannibalization. The various options for cannibalization will be discussed in a later subsection. If a part is not located, the fact that the aircraft has a "hole" is filed by calling subroutine RPTNOR (report not operationally ready); if the rules prescribed by the user permit (see Sec. XI), an attempt is made to locate the needed part at another location in the theater and to have it shipped.

If all resources are available the task is initiated with a call to subroutine DOTASK that places the task in the TASKQ heap and also places pointers defining its location in the in-process queues associated with the aircraft and with the shop that is doing the work. The duration of the job is determined on the basis of the mean task time and the distribution specified by the user. When chemical protection ensembles are worn, the "heat factor" is also passed to subroutine TTIME where it acts as a switch to call subroutine CWTIME, which estimates the additional time that will be required because of the ensemble, and how long the personnel will be able to work in light of the ambient temperature, humidity, and chemical contamination at the work place. (See the discussions of subroutines TTIME, CWTIME, and CKTEMP in Sec. XIV.)

The DOTASK subroutine is also used when it is necessary to stop an on-equipment task. Since on-equipment tasks receive priority over parts repair tasks, the only times that on-equipment tasks are interrupted is (1) when a task is stopped for a higher priority on-equipment task, (2) when the number of personnel at a shop is reduced

⁸This could occur in a 66-1 type organization but not in a COMO (66-5) organization, since the parts repair personnel at wing level in COMO are differentiated within TSAR by means of a different identity.

because of a shift change, (3) when the work crew must stop to rest and cool off, or (4) when the airbase is attacked and shop personnel are lost. At those times the subroutine is entered at the entry point STPTSK, and the needed bookkeeping is done on the pointer systems used with the aircraft, the shops, and the INTTSK and TASKQ arrays. When personnel are reduced because of a shift change, the last task that was initiated is the first to be interrupted.

If a part is available, but some other resource prevents the task from being initiated, any alternative procedure (set of resources) for accomplishing the task that has been supplied by the user is checked to see if those resources are available. If they are, the task is initiated using the alternative procedure; if they are not, the task must wait in the appropriate shop's wait queue. If the task had already been waiting, processing is complete. If it is being checked for the first time, subroutine ACWAIT is called to store the relevant data in the WAITSK array; the resource for which a shortage prevented the primary procedure from being initiated is taken to be the reason for the delay. When the task is placed in the shop's wait queue it is placed last in line if ORDWT = 0; if ORDWT = 1,⁹ subroutine INWAIT is called and the task is placed in the shop's wait queue such that the aircraft with the least time remaining before it had been estimated to be ready to fly is placed first.

The last step for a task that is being checked for the first time is to dispose of any part that must be removed from the aircraft. If the part is not normally repairable in the theater, it is eliminated and another may be requisitioned from CONUS. If it is not normally repairable at the base where it was removed, it is shipped to whatever location has been specified for its repair (using the SHIPTO array data input from CT34). It may be shipped directly after removal from the aircraft, or it may first have to be checked on base.¹⁰ If the part can be repaired on base, it is sent to the appropriate repair facility. If the repair facility has been closed by damage from an air attack and it is not projected to be repaired before the part could be sent to another base for repair and returned, or if the AIS stations needed for repair have been lost in an air attack, the bases specified for lateral repair are checked to see if they have an open shop and/or the needed AIS station; if so, the part is shipped to another base.

⁹See CT3/1.

¹⁰Any part, LRU, or SRU with a NRTS rate of 101 is shipped directly after removal from an aircraft (or from an LRU); if the NRTS rate is from 1 to 100, the part must undergo the administration delay before being checked for NRTS action.

The part is removed from the aircraft when the task is first checked, even though the resources were not available to start the on-equipment task at that time; it is assumed that the overall resource demands for the task are adequately approximated by the task's resource requirements whether they are used then or later. The repair of the part is delayed by a time equal to the sum of the nominal task time and the backshop administrative delay time (see Sec. VI).

If the task started in INITSK is a task that had already been started but had been interrupted, or is a task that had already been checked but had had to wait, the necessary bookkeeping for the pointers is accomplished before control is returned to MANAGE.

Of the many data maintained for each aircraft, two are flags used to rapidly identify each aircraft's current status; the first flag (stored in the 12th position—ACN(-,12)) of the aircraft array defines the aircraft's location within the overall sortie cycle, while the second flag (ACN(-,16)) defines the degree to which the aircraft has progressed through the several steps in the preflight process. The states corresponding to various values of the first flag are:

ACN(-,12)	Aircraft Status
1	In flight
2	Inactive for the postflight delay
3	Unscheduled maintenance before final mission assignment
4	Inactive for the preflight delay and final mission assignment
5	Maintenance following final mission assignment
6	Ready for flight
7	Undergoing deferred maintenance tasks

The several preflight states defined by the second flag are outlined in Sec. V.

12. CANNIBALIZATION

When a part must be replaced on an aircraft and a replacement is not immediately available, TSAR may be directed to cannibalize the needed part from another aircraft in certain circumstances,¹¹ and the part that is cannibalized may itself be broken in the

¹¹Parts cannibalization may be selectively prohibited by entering -1 for a specific part type in the CANNM array using CT35/1. If the value is less than -1, the part may only be cannibalized when at least DOCANN aircraft at that base already need that type of part.

process.¹² The rules governing cannibalization are managed by the user with his setting of the control variables CANMOD (cannibalization mode), MXHOLE, DOCANN, DOWNTM, and CDELAY. The basic user choices are (1) whether a part may be cannibalized when there are reparables on base, and if so (2) which of the aircraft may be considered. The aircraft that may be considered must be of the same type and must also be undergoing unscheduled maintenance. Four possible categories are defined: (1) aircraft with parts missing, whose criticality for the designated mission would not be affected; (2) all aircraft that have parts missing; (3) aircraft without holes, if the criticality would not affect the designated mission; and (4) all other aircraft. If cannibalization is selectively restricted to aircraft in either of the first two categories, the donor aircraft must have at least as many missing parts (i.e., "holes") as the recipient aircraft. No matter which category is chosen, aircraft that already have a part missing are checked before the others are checked. When an aircraft has two or more parts of the same kind at different locations on the aircraft, each may be cannibalized unless cannibalization is restricted for certain of the locations; when two or more may be cannibalized, priority is given to the part with the shortest cannibalization times. Parts not normally cannibalized can sometimes be cannibalized if sufficient aircraft already need the same part (see footnote 11 above). The user may also prohibit cannibalization of the part from any aircraft that already has had MXHOLE parts removed, or whose estimated ready-to-fly time is within DOWNTM hours; for aircraft without "holes" TSAR has a built-in minimum constraint of 90 minutes for this time.

These optional constraints are defined for various values of the control variable CANMOD as shown in Table 3.

Cannibalization is done by subroutine CANNIB. When an aircraft is checked for the needed part, the waiting tasks, required tasks, and deferred tasks are checked first in that order. If the same task is found to be required on the aircraft but the part is not required (is not broken), it is assumed that the part can be removed; if the part may be required in a subsequent segment of the task network, it is assumed that the part is not available. If the task is not found in any of those categories, the in-process tasks are checked; if the same task is not being processed, the aircraft is considered suitable for cannibalization.

¹²The probability that a part is broken when cannibalized is read into the BADCAN array using CT35/2; no part will be cannibalized if the probability that it will be broken is greater than NOCANN (see CT4/2).

Table 3

CANNIBALIZATION CONSTRAINTS*

CANMOD	Cannibalization Permitted with On-base Repairables	Eligible Aircraft (none with ready-to-fly time less than DOWNTM hours → 90 minutes)
0		None
1	No	Aircraft with parts missing whose criticality for the designated mission would not be affected
2	Yes	Ditto
3	No	Aircraft with other missing parts
4	Yes	Ditto
5	No	Aircraft whose designated mission is not affected by part
6	Yes	Ditto
7	No	Any aircraft
8	Yes	Ditto

*Parts that may be cannibalized only when the DCCANN constraint is satisfied are distinguished by an entry in the CANNTM array that is less than -1.

When a part is removed from an aircraft, data regarding the new "hole" is stored in the NORQ array using the NORRPT subroutine (see Sec. XIV) and is recorded with the task status flag associated with the task (see below). And when the part is removed from an aircraft for which the related task was not already outstanding, a notice to replace the part must be added to that aircraft's list of required or deferred tasks. If the task criticality of the root segment of the network in which the part is located is negative (i.e., has some probability of being deferred until night), the part is assumed to be mission critical for all missions and the task is added to the aircraft's required tasks.

A task status flag is used to keep track of the state of an aircraft's tasks and is carried along with all references to each task. The values of this flag define task status as follows:

Task Status Flag	Status
0	No part required
1	Part required, not yet recorded in NORQ array
2	Part required, recorded in NORQ array
3	Part required, part removed, not yet recorded
4	Part required, part removed and recorded
5	Replace part only, ignore network, part removed, recorded

When a part is cannibalized from one aircraft to permit work to be carried out on another aircraft, the time required to get the part and to complete the basic task on the receiving aircraft is the sum of the time normally required for that task, plus either the time for cannibalizing a part of that specific kind (from the CANNTM array), or the default cannibalization time; the latter is equal to one-half the true time selected for the task plus CDELAY minutes, as defined by the user with the control variable CDELAY.

13. CHECK FLIGHTS TO VALIDATE MAINTENANCE ACTIONS

Under some conditions it is necessary to test the adequacy of maintenance actions by flight-testing an aircraft. It is possible to simulate such flights in TSAR and the demand they place on aircrews and aircraft. The CT15/88 card is used to specify the numbers of the simple tasks and task-network root-segments for which such flights may be required, and the probability (x10000) that the flight is *not* required when the specified task is scheduled. Such requirements can be entered with the CT15/88 cards for unscheduled maintenance networks specified with CT7 cards and in the shop-task sequence lists entered with CT29 cards; they may also be specified for any of the battle damage task networks specified with CT15/2 and CT15/3 cards. This check flight option will *not* function for *subsequent* task elements in a task network. These features are activated for those aircraft types for which a "1" has been entered for the control variable in column 45 on CT15/5.

This feature is mechanized as follows: After an aircraft has landed and an initial determination of the next mission has been made, the tasks are divided between "deferrable" and "required before the next combat sortie"; the required tasks are then checked to determine if a check flight will be required when that maintenance has been completed. If a flight will be required, a special aircraft flag is set. Subsequently, as aircraft maintenance is carried out, munitions tasks are omitted. When all other required

maintenance is complete, an aircrew is sought to fly the aircraft; if available, and the runway is open, etc., the aircraft is flown (for 45 minutes). If the aircraft cannot be flown at that time, it waits until conditions permit the flight; conditions are checked every two hours at the same time that checks are made for aircraft transfers that have been delayed for lack of aircrew or runway, etc. If a check flight is required when a previously deferred task is carried out, any munitions that have already been loaded are first downloaded. When the check flight is accomplished, additional maintenance may be generated; if it isn't, the aircraft is loaded with the appropriate munitions and is ready for a combat mission.

For aircraft that must be moved to another airbase (either a host base for DOB aircraft, or a rear maintenance base) the test flight is flown from the airbase where the task that required the flight was carried out. For check flights on aircraft that are to be transferred to another base, no additional maintenance is generated by the check flight.

14. POSTFLIGHT INSPECTIONS AND MORNING PREFLIGHT INSPECTIONS

With TSAR, the user may specify distinct *postflight* inspections (tasks) that depend upon the mission just completed. The number of the root element of the inspection (task) for each of up to five missions is entered for each type of aircraft with the CT15/5 card. The time during the maintenance cycle at which this work is to be performed is defined by the location of Shop #25 in the CT29 cards. Thus, postflight inspections may be designated by mission type, for each type of aircraft, at whichever bases the user chooses.

In addition to the other *preflight* activities that may be simulated with TSAR, the user may also require that a scheduled inspection be imposed early in the morning before the aircraft are to be launched. The nature of the task is specified for each aircraft type by entering the task number (or root-segment number) on the CT15/5 card, and by specifying the times¹³ in the morning (hour and minute) at which the inspections are to be started at specific bases on the appropriate CT17/3 cards. Thus, one can impose a unique task for each type of aircraft, and can initiate the inspections at different times at

¹³The inspections will be started on day 1 if an hour less than 24 is specified. If the user wishes to start these inspections on a later day, the hour counting from time zero should be entered: e.g., to start at 5:15 AM on day 3, the user enters 53 hours and 15 minutes, or 5315.

different bases, if desired. And if a task number, or a time, are not specified, inspections of that aircraft type, or at that base, will not be imposed. If the user wishes the inspections to begin at midnight, 24:00 must be entered rather than 0:00.

When it is time to start the early morning inspection, subroutine INSPEC is called, by base, from MANAGE and each aircraft is checked. The inspection is not required for all aircraft, only those that have had a final mission assignment and have been refueled, or are being refueled, and those that have had their mission-dependent munitions uploaded, or are having those munitions loaded. Furthermore, an inspection is not imposed if the aircraft ready-to-fly time is more than two-and-a-half hours hence.

For those aircraft that are ready-to-fly, an attempt is made to initiate the inspection immediately; if sufficient resources are not available, the task waits. For aircraft not yet available but that meet the criteria noted above, the inspection is added as a task requirement to be carried out after the current work is completed. An aircraft will not be launched on a mission until the inspection has been completed.

15. DEFERRED MAINTENANCE

On-equipment aircraft maintenance that has been deferred as nonessential for an aircraft's designated mission may be taken care of in four different ways, determined by the task criticality (CT5) of the deferred task. The first possibility (mentioned in the first subsection) is that a different mission will be chosen for the aircraft for a subsequent flight and the deferred task will become mission essential and be transferred from the DEFTSK array to the required tasks in the RQDTSK array.

The second possibility is that a deferred task may be deferrable *only* for some number of sorties (LTHDEF sorties) or—a third possibility—until the end of the nominal flying day, independent of mission essentiality. In the first instance the task will be redefined as a required task after the LTHDEF sortie, and in the other it will be redefined when subroutine INIDEF (initiate deferred maintenance) is called after the end of the flying day, as discussed next. In both instances, however, the maintenance will be required at any time the task is essential for a new mission assignment.

All deferred tasks are reviewed each evening after the end of what the user has designated as the "flying day," i.e., after ENDAY. At those times, subroutine RUNAC calls subroutine INIDEF when all other tasks outstanding for the aircraft have been completed except, perhaps, the preflight task set that may have been delayed until the

early morning hours. Subroutine INIDEF is also called at 2000, 2200, and 2400 during the night by subroutine PLAN to check for needed resources that may have been released by other tasks.

At these times, subroutine INIDEF redefines as required the deferred tasks that must be done at night and also attempts to initiate each of the aircraft's deferred tasks that are optional, if the nominal time for the task will permit it to be completed no later than the hour specified by the user as the last time at which the rearming process must begin (LSTTOD—last time of day). After checking that tasks aren't already waiting at the task's designated shop, the INTSK subroutine is called to check on the availability of necessary resources. If available, the task is begun; if not, it is left as a deferred task rather than being redefined as a waiting task, because that status would prevent further actions with that aircraft. When the task data have been filed in the TASKQ array, the mission-capable status of the aircraft is updated.

The fourth possibility for working off deferred maintenance tasks occurs on those days for which the user has specified that the weather will not permit operations at a particular base for specified aircraft types. Subroutine MANAGE calls subroutine DODEF (do deferred tasks) periodically, and the weather status is checked for each base and each aircraft type at four-hour intervals starting at 0400, when it is presumed that the day's weather conditions will be known. For all aircraft that are otherwise ready to fly, subroutine INIDEF is called and checks whether available resources will permit that aircraft's deferred tasks to be completed by the LSTTOD on the following day. This processing follows the same rules as were described above.

16. AIRCRAFT STATUS PROJECTION

A simulation of airbase operations must emulate, at least in a limited way, the scheduling and control activities that are carried out by the job-control shop at each airbase in order to utilize the available resources efficiently. Choices must be made as to the tasks to be performed, repairs to be done, munitions to be assembled and aircraft assignments. In the real world these choices are made in the context of a much richer body of knowledge regarding assets, capabilities, and requirements than is possible (or at least practical) in a simulation. Furthermore, the procedures used and results obtained in the real situation are, at least in part, dependent upon the skill, knowledge, and experience of the particular job control managers available and therefore vary from one

circumstance to another. All that reasonably can be expected of a simulation are mechanisms to allow the user to define broadly differing policies for managing aircraft maintenance and repair jobs and to achieve a degree of efficiency in the utilization of the available resources.

TSAR incorporates a variety of features for these reasons, a key one of which is the periodic development of what might best be called the projection of aircraft supply and demand. These projections provide the context within which decisions are made regarding aircraft assignments, unscheduled maintenance, and munitions buildup for the subsequent two-hour period.

As is outlined at greater length in Secs. VII and XIX (CT50), the sortie demand data specify the airbase, the aircraft type, the mission, the number of aircraft, the mission priority, the receipt time of the demand, and the desired launch time. Provisions are included that also permit the user to stipulate that a number of aircraft of a particular type be maintained in an alert (cocked) status, so that they may be launched whenever they are needed for a specific mission. These data provide the information with which the pattern of sortie demands is projected.

Since the current status of each aircraft assigned to a base is known at any particular time, one may also make a projection of when sorties of various types might be launched. These projections are also made every two hours for each base, each aircraft type, and each mission for each of the several priority levels. The projections of sortie demand and aircraft availability cover a period of several hours, out to an internally adjusted time horizon. By comparison of these two projections, aircraft assignments are made so as to give priority to the more urgent, higher priority demands.

These projections are developed in subroutines PLAN and PLAN1 and the essence of the supply and demand comparison is stored in the SORDEF (sortie deficiency) array for use as decisions are required. The time horizon for these projections is controlled internally and may be made a function of the time of day; typically the time horizon is relatively long during periods of limited activity and shorter during periods of more intense flying.¹⁴ The projections of supply and demand within the time horizon are divided into 16 equal time blocks for each time horizon; each sortie

¹⁴The time horizon is controlled either by the user or by the default conditions; as currently written, the default conditions are: a planning horizon of 12 hours from midnight to 0400, 8 hours from 0401 to 1600, 20 hours from 1601 to 2000, and 16 hours from 2001 to 2359.

demand time and estimated aircraft ready time is associated with the appropriate time block.

17. PREPARING THE PROJECTIONS OF SUPPLY AND DEMAND

Subroutine PLAN is called by MANAGE on even-numbered hours, and the first step is to estimate aircraft supply. Each base and each aircraft is checked and the estimated ready-to-fly time determines which time block is credited with an available aircraft. The ready-to-fly time is determined either by the value that was estimated when the maintenance requirements are determined in subroutine PSTFLT (and occasionally updated) or, for those aircraft that are currently in flight, the ready-to-fly time is projected on the assumption that the aircraft will be reassigned to the same mission and will spend a nominal amount of time in unscheduled maintenance (as specified by the user with CT15/1). These data are collected for each mission and for each aircraft type and stored temporarily in the SUPPLY array; they are then converted to cumulative distributions over time, and subroutine PLAN1 is called to project the demand and derive the needed comparisons. The ACA (aircraft assignment) array is updated at the same time for the aircraft that are currently on base and have already been assigned to specific flights.

Subroutine PLAN1 is called separately for each type of mission, each type of aircraft, and each base. The demands for each such subset are first collected for the highest priority demand—Priority #1—in array DEMAND and converted to a cumulative record in array SUM. The aircraft supply for that mission and aircraft type (that was stored in the SUPPLY array) is then projected ahead on the assumption that available aircraft will be launched when required for the first priority flights and will return, and be turned around in the nominal sortie cycle time specified by the user with CT15/1. The projected surplus or deficiency during each time interval for first priority flights is then stored temporarily in a local array. This entire procedure is then repeated for each of the lower priorities, with the continuing assumption that all higher priority flights are also flown and subsequently turned around, when sufficient aircraft are available.

Three data are then stored in the SORDEF (sortie deficiency) array for each of the 16 time blocks. The highest deficient priority and the total demand at all priority levels during and subsequent to each time interval are stored in the first position; the deficiency at the highest deficient priority is stored in the second position; and the number of sorties

expected to be available at the highest deficient priority level is stored in the third position. These data are used in assigning aircraft during the subsequent two-hour period.

When these data have been prepared for all bases, aircraft types, and missions, four final actions are carried out in PLAN. The first is to check whether the flag that will delay the preflight procedure should be set. If the nominal flying day is complete—i.e., it is ENDAY or later—the DELYPF flag is set to permit the preflight process to be delayed and deferred maintenance to be initiated.

The next action in PLAN is to collect the total number of known sortie demands for each type of aircraft and mission at all bases and to store that information (in ACMDTA(12,-,-)) for use in the CIRF repair algorithms. Then, subroutine REASSG (reassign) is called to check whether more aircraft have been readied for a mission than are needed; if so, they are reassigned to a mission that is deficient. The last activity, conducted at 2200 and at midnight, is to check that all maintenance that had been deferred until night will receive attention.

V. PREFLIGHT TASKS AND MUNITIONS BUILDUP

The preflight events dealt with by TSAR include a preflight delay, final mission assignment, aircraft reconfiguration, loading of mission-dependent munitions, and refueling. Additional munitions—i.e., the basic munitions that are always to be carried—will normally be entered separately as individual tasks, as explained in Sec. IV. The other tasks to be discussed in this section in connection with the preflight tasks are the munitions buildup tasks. The procedures and resources associated with these events are sufficiently different in detail that nine special subroutines were developed. When the basic control for aircraft maintenance is passed to subroutines STARTM and RUNAC by subroutine MANAGE, the management of the preflight events is further delegated by subroutine PREFLT, as was mentioned in Sec. IV; for the munitions buildup tasks, MANAGE transfers control directly to MUNEEED or DOBILD.

The preflight delay was envisioned as a period of dead time that the user might wish to specify (CT15/1) before the munitions-related events and (typically) subsequent to the completion of the unscheduled maintenance tasks. When it is necessary to delay the preflight events until after the expected receipt of sortie demand information, the length of this delay is modified endogenously. Immediately following this delay a final determination is made as to the next mission that the aircraft will fly and a tentative assignment is made to a specific flight, alert force, or set of spare aircraft. These selections are based on the most recent projections of aircraft supply and sortie demand and may involve a change of mission from that designated tentatively at the time of postflight "inspection." After TSAR determines the appropriate aircraft configuration required for the most effective munitions available for the next mission, the aircraft is reconfigured as necessary and the weapons are loaded if they were not retained from the prior sortie.

The periodic projections of aircraft supply and sortie demand are also used to generate the demands for munitions buildup. The munitions demands imposed by the sorties that are expected to be flown are compared with the available and in-process munitions, and work is initiated to offset any apparent shortfall. The prescribed procedures give priority to the earliest high-priority sorties that have been demanded.

Several TSAR work centers, or shops, are set aside exclusively for use with the preflight events. Shop #26 is associated with the preflight delay and assignment, Shop #27 with reconfiguration, Shop #28 with mission-dependent munitions loading, and Shop #29 with refueling. Shop #30 is responsible for all munitions buildup tasks. As discussed in Sec. IV, the "flight line" shop, Shop #25, also can be used in connection with the basic munitions and certain TRAP.

When the preflight events, or tasks, are listed in the user-supplied task-shop sequence data (CT29), as described in Sec. IV, Shop #26 and Shop #29 may be listed in any sequence with the individual tasks and other shop numbers. However, the most logical arrangement would be to list Shop #26 (which implies mission assignment, reconfiguration, and mission-dependent munitions uploading) as, or with, the last group of shops. Thus, if one had designated only four maintenance shops, and listed the shop sequence as 1, 2, 3, 4, 29, 0, 26, 0, 0, all tasks would be processed as quickly as resources and task incompatibilities permitted, except for the mission-dependent munitions tasks that would be accomplished last. If, however, the sequence were listed as 1, 2, 0, 26, 0, 29, 3, 4, 0, 0, the work required by Shops #1 and #2 would be completed first, the final mission assignment and weapons loading would be done next, and the work by Shops #3 and #4 and the refueling would be done last. In general, it is advisable to defer final mission assignment and munitions selection as much as practical, in order to permit those decisions to be made with the most current information.

A special control variable is provided to facilitate a separation between fueling and the rearming operations. If NOFUEL is initialized as unity, these operations will not be done at the same time; this constraint overrides any contradictory rule implied by the shop sequence listing.

Management of the preflight maintenance tasks is facilitated by a flag that is maintained for each aircraft in the 16th position of aircraft array—i.e., in ACN(-,16). The flag can be set to 13 different values defined as follows:

Preflight Flag ACN(-,16)	Value When Refueling Is:	
	Not in Process	In Process
Preflight tasks (Shop #26) have not been initiated	1	8
Preflight delay is in process	2	Not permitted
Delay (Shop #26) complete; awaiting assignment, or assigned but awaiting 2nd of Shop #27 subtasks	3	10
Reconfiguration (Shop #27) is in process (one or two subtasks)	4	11
Reconfiguration (Shop #27) complete; one subtask of Shop #28 may be complete	5	12
Munitions loading (Shop #28) is in process (one or two subtasks)	6	13
Preflight tasks complete	7	14

NOTE: As can be seen, refueling (or any other task) may not be carried out during the preflight delay.

1. MANAGEMENT OF PREFLIGHT TASKS

Preflight tasks are managed by subroutine PREFLT in much the same manner as subroutine STARTM manages unscheduled maintenance tasks. When tasks for Shop #26 or Shop #29 are first identified in STARTM, control is immediately transferred to the entry point PRFLT in subroutine PREFLT. Unless the munitions related tasks (Shop #26) are to be delayed, or another maintenance task is in process, the preflight delay is initiated and the preflight flag is updated before control is returned to STARTM. If the delay may not be initiated, the task is stored in the wait queue associated with Shop #26.

When the preflight delay is concluded, MANAGE transfers control to RUNAC at RUNAC2, and control is immediately passed to the beginning of the PREFLT subroutine, where the termination of preflight tasks is managed. The procedures for terminating the other preflight tasks parallel those used for other tasks: Subroutine STOPIT is called when USECW > 0 to check on the condition of personnel with GOREST and CKTEMP, and then RUNAC is called to release any remaining resources; when USECW = 0, subroutine RUNAC is called directly. In both cases personnel that have been designated as a load team are retained rather than released, until all jobs on the aircraft that are waiting for a load team have been completed. The only exception is

when the load team's temperature would exceed the allowable limit, and the team must stop and cool off. RUNAC calls ENDTSK to release the resources before calling subroutine PREFLT to check for additional preflight tasks. If the personnel are available, ENDTSK attempts to reassign them using subroutine DOWPRE. (DOWPRE fills much the same function as subroutine CHECK does for unscheduled maintenance; the primary differences between DOWPRE and CHECK are that the former first checks to see that both subtasks of the reconfiguration and uploading tasks are complete before reassigning personnel and equipment, and it does not have any equivalent to the parts repair sections of CHECK.)

When control is returned to subroutine PREFLT, an attempt is made to initiate the next preflight task unless the preceding task has not been fully completed. Four distinct subroutines are used to handle aircraft assignment (ASSIGN), reconfiguration (RECNFG), munitions loading (UPLOAD), and refueling (REFUEL) tasks, because of the distinctive characteristics associated with each task. These subroutines are called in the appropriate order by subroutine PREFLT and by subroutine DOWPRE when preflight tasks have had to wait.

2. MISSION ASSIGNMENT

As soon as the preflight delay is completed, the final mission assignment for the aircraft is made using subroutine ASSIGN. The scheduled ready-to-fly time is first interpreted in terms of the 16 time blocks into which the periodic estimates of aircraft supply and sortie demand are divided. The highest outstanding priority demand for the mission for which the aircraft had been designated at the time of the postflight inspection is then identified. The process by which this is done is first to identify the aircraft's lowest permissible assignment priority, the maximum number of aircraft that are expected to be ready, and the maximum number of aircraft to be assigned at that priority level using data generated by the look-ahead planning process described in Sec. IV. Next, the requirements for alert aircraft and then the requirements for scheduled flights are each checked from the highest priority level to the lowest permissible level. The aircraft is assigned to the highest priority demand that has not already been filled.

If the aircraft is not assigned by this procedure to the mission for which it was designated, a check is made to see which other missions the aircraft could be readied for, taking into account whatever maintenance has been deferred. The procedure just

described is followed for whatever other missions the aircraft is able to fly, until the aircraft is assigned. If it still has not been assigned to an alert force or a scheduled flight, it is committed to the mission to which it was tentatively assigned during the postflight inspection and is associated with the other spare aircraft configured for that mission.

In the event the aircraft had returned from its previous mission with its munitions on board and it is assigned to a different mission, the munitions are returned to stock without any specific delay or requirement for personnel or equipment. Since the new mission will probably require that the aircraft be reconfigured, it is assumed, in effect, that the munitions downloading is a part of the reconfiguration.

3. AIRCRAFT RECONFIGURATION

After an aircraft has had its next mission assigned, subroutine RECENFG (reconfigure) is called to check whether the various racks, pylons, etc. (TRAP) with which the aircraft was equipped for the previous mission are appropriate for the next mission. If not, they must be removed and the aircraft must be reconfigured.

Before explaining those procedures, we will first review how the appropriate weapons load is determined. For each aircraft-mission combination, the user may specify up to ten different standard combat loadings (SCLs); these should be ordered with the most desired munitions first. The characteristics of an SCL include an aircraft configuration (a number corresponding to the entries in the CONFIG requirements array), a heat factor, and one or two sets of munitions, each with a specified requirement for personnel, equipment, and time for uploading. Each configuration, in turn, is characterized by a heat factor and one or two sets of TRAP; the requirements for mounting each set of TRAP include personnel, equipment, and time. As with such descriptors in the other kinds of tasks, any of these requirements may be satisfied with a null entry; if, for example, the same crew using the same equipment is to load both sets of TRAP, those requirements and the total time could be specified for the first task, and the descriptors for the second reconfiguration task could be limited to the TRAP, with null entries for personnel, equipment, and time.

In determining whether a reconfiguration is required and what the new configuration should be, subroutine RECENFG checks first on the configuration of the SCL that is preferred for the assigned mission. A check is first made on the status of the munitions shop if that facility has been specified as an essential resource. If that

constraint is satisfied, the munitions stocks are checked next. Only then is a check made to see whether the specified configuration is the same as or different from the aircraft's current configuration. If it is different, a check is made to see if either of the two sets of TRAP is common to the two configurations; if so, it is presumed that they will not need to be changed. When the new TRAP requirements are established, a check is made of their on-base stock levels. If either the munitions or the TRAP required for reconfiguration are not available, the next best SCL is checked. If these resources are insufficient for all SCLs, the task must wait. The task must also wait when there are sufficient of these resources for an SCL, but insufficient personnel and equipment. Cross-trained personnel may be substituted for the normal personnel requirement on those tasks and bases that are specified. When all resources are available the appropriate munitions and TRAP are withdrawn from stock. The times for the reconfiguration tasks are computed taking into account the specified heat factor when $USECW > 0$; if the task involves a designated load team that can work on a sequence of tasks before needing to rest and cool off (when $USECW > 0$), the task temperature of the team at the conclusion of any prior task is also taken into account. When TRAP must be downloaded it is assumed that it will take the same amount of time to download a set of TRAP as is required to upload it, but that the personnel and equipment associated with the new TRAP will perform the job.

4. MUNITIONS LOADING

When reconfiguration is complete, subroutine `UPLOAD` is called to initiate the munitions loading tasks. Because the required munitions were set aside when the requirements for reconfiguration were checked, all that needs to be done is to check on the facility itself, when specified, and the personnel and equipment required for the loading subtasks. If they are available (substitute personnel may be used when specified), a call to `ADDTSK` places them in the `TASKQ`. If $USECW > 0$, the call to `TTIME` and `CWTIME` determines the appropriate task time for the protective ensemble that is being worn and the expected temperature rise for the work crew. If a load team is being used and is too hot from earlier work on the aircraft to be able to do any more, they are sent to cool off. If personnel are not available, the tasks are placed in the wait queue for the munitions shop; that queue is checked by subroutine `DOWPRE` whenever resources from that shop become available. If only one of the subtasks may be initiated, the other is placed in the wait queue.

5. REFUELING

Refueling is included among the preflight tasks but does not have a rigid relationship to the other preflight tasks, as they do with each other. Refueling is accomplished by Shop #29, whose position in the shop sequence list is under the user's control, as discussed earlier. Thus, this task may be placed where desired in the task sequence list. Furthermore, the refueling task may have its own list of incompatible tasks, as does an unscheduled maintenance task. In addition, the user controls the special variable NOFUEL, which prevents fueling when any of the munitions-related tasks are in process if it is initialized to unity.

Management of these restraints is handled by the PREFLT subroutine and, when necessary, by the DOWPRE subroutine. When conditions pennit, subroutine REFUEL is called to process the fueling task. The only feature unique to this task is the requirement for a quantity of POL. The amount of fuel required is taken to be a characteristic of the aircraft type; the other resources required for refueling are stored in the TSKRQT array, along with those for the unscheduled maintenance tasks. When subroutine REFUEL is called, the required POL is withdrawn from stocks and the necessary personnel and equipment are assigned; if the resources are insufficient for the basic refueling procedure and for any alternative procedures that are listed, the task is placed in the refueling shops' wait queue. Control is returned to subroutine PREFLT.

When the user has specified that aircraft are to be hot-pit refueled immediately after landing and before the aircraft has taxied to the intended parking location, the normal fueling task is omitted if the hot-pit hydrant is used.

6. MUNITIONS BUILDUP

Although munitions buildup is discussed here in connection with the other munitions-related activities, it constitutes a completely distinct set of off-equipment functions that are managed independently from the aircraft-related tasks in a separate set of subroutines. Resource requirements for the buildup of each type of munition are specified in much the same manner as simple parts repair jobs, but the procedures used to schedule and prioritize these assembly activities are unique to these tasks. Nonready munitions may be categorized simply as "unassembled" or may be represented by a list of the individual components required to assemble a single round. In the latter instance, stockage records are maintained for the individual types of components, and the

limitations imposed by component shortages and by the use of a particular type of component (e.g., a laser guidance package) in two or more weapon types can be represented.

The periodic aircraft supply and sortie demand projections provide the basic "operations" data that drive the weapons buildup selection and prioritization logic. Immediately following that projection, subroutine MANAGE transfers control to subroutine MUNEEED (munitions needed) to determine munition needs (when the control variable BUILD (CT1) has been initialized to 1). For each type of munitions assembly personnel, a tally is first prepared at each base of the number of munitions assembly tasks that are expected to be completed within the next two hours, and are waiting. Another tally is prepared of the number of each type of munition that could be assembled, based on the numbers of components that are available, and not already committed to ongoing or waiting assembly jobs. Finally, a tally is made of all on-base munitions that are loaded, assembled, being assembled, or are already waiting to be assembled. Subroutine MUNEEED then tabulates the sorties that are projected to be flown in terms of launch time, priority, mission, and aircraft type, on a base-by-base basis. Flight times within the planning time-horizon are divided into five time blocks. Demands for alert aircraft are presumed to generate equivalent munitions demands in the first and third time blocks.

With these demand data, control is then transferred to CKBILD (check buildup requirements). This subroutine first converts the sortie demands into the munitions demanded by the preferred SCL for each particular mission and aircraft type and then checks whether sufficient munitions are available or committed to be built for the various demands. The checks are made first for the highest priority missions in the second time period, then for the next priority, etc. Next, the demands in the third time block are checked, etc. (Because time would not permit the buildup of munitions to meet the demands for the first time block, they are not considered.) Whenever sufficient munitions are not available or have not been scheduled to be built, a weapons buildup task is defined—if sufficient unassembled munitions, or munition components, are available—and control is transferred to subroutine DOBILD where the required personnel and equipment are checked (substitute personnel types may be designated). If sufficient resources are available a location for assembly is selected and the task is stored in the BUILDQ array; distinct sets of facilities may be defined for assembling guided and unguided munitions. If tasks cannot be initiated they are placed in the wait queue in the

BACKLG array, until the number waiting equals the number of tasks that are expected to be completed before the munitions requirements are checked again. If sufficient unassembled munitions are not available, or the necessary personnel are not available, the adequacy of munitions for the next lower priority SCL (for that particular mission and aircraft type) is then checked.¹ If no munitions tasks can be started, the demand is dropped. This process continues for all priority levels and time blocks for each base in turn.

If the munitions assembly resources are not fully committed to the immediate demands, they may be used to build up a reserve;² the choice of the munitions to be assembled is based on the existing supplies and the history of the demands for munitions (as generated during the simulation).

When a munitions buildup task has been completed, subroutine MANAGE (or STOPIT) transfers control to subroutine ENDBLD where the task is removed from the BUILDQ heap, the shop pointers are updated, and the personnel and AGE are returned to stock.

When control is returned to MANAGE, it is immediately transferred to the DOWBLD (do waiting buildup) entry point in subroutine DOBILD (if the CW features are not activated), where a check is made to see if the released resources can be used for another weapons assembly task; when job termination is handled by subroutine STOPIT (to deal with the CW effects), that subroutine calls DOWBLD if the personnel were released.

¹ Advantage may be taken of this logic to avoid committing all personnel on abnormally long assembly jobs and to give priority to more rapidly assembled though less-desirable munitions, thus obtaining some kind of ready munition for the aircraft. To take advantage of this option, the assembly personnel must be divided into two groups: some appropriate for the hard-to-assemble munitions and the remainder capable only of assembling other kinds. Furthermore, the first group of personnel should be specified to be cross-trained to do the work of the second group, and the assembly tasks for the more readily assembled munitions should be flagged so that the cross-trained personnel may be used. When these things are done the personnel that are used to assemble the better munitions will be used on the less effective munitions when there are no longer any of the better munitions to be assembled. The resource against which aircraft waiting times are charged will not be a munition type that is not available for assembly, but rather will be the first type of munition that could have been assembled, given the other resources.

² See columns 31-35 on CT17/1.

VI. OFF-EQUIPMENT MAINTENANCE—PARTS AND EQUIPMENT REPAIR

TSAR provides the user with features that permit the examination of a great many questions related to parts stockage and parts repair policies. Indeed, various questions concerning autonomous and consolidated parts repair capabilities within the theater were central in shaping TSAR's theater characteristics. In its present form, TSAR may be used without any consideration of aircraft parts, with autonomous airbase parts repair facilities, with repair in whole or part at other operating bases, with a centralized parts repair facility in the theater, or with a combination of the last three modes. The constraints imposed by faulty support equipment may also be reflected.

A specialized set of subroutines handles the several elements of the parts and equipment repair procedures. The first three of these subroutines can be used to initialize the parts stockage data and the spare-parts pipelines from CONUS to the theater, and, when there is a CIRF, between the CIRF and the operating locations. The first subroutine used for parts and equipment repair determines the appropriate administrative delay to simulate before initiating the repair process. Following that delay, other subroutines check on the availability of resources, store the repair jobs that are initiated, and conclude the repairs; another special subroutine is available to disassemble LRUs to obtain SRUs. When parts repair is done at a CIRF, other subroutines come into play. These procedures will be outlined briefly later in this section and discussed more completely in Secs. X and XI.

1. INITIALIZATION OF PARTS INVENTORY AND PIPELINE DATA

Although the initial parts inventory and pipeline data may be entered for each base, much as for the other classes of resources, the user instead may elect (by initializing OUTFIT) on CT3/3 to have those data generated as an integral part of the input and initialization process. When this option is elected (for some or all bases), the nominal quantities of parts that should be procured for each base are determined according to the standard computational procedures outlined in Chap. 11 of *Air Force Supply Manual 67-1* [9], or, for WRSK kits, with an approximation to the cost-sensitive

DO-29 [10] procedures. For in-theater units, both POS and BLSS are assessed.¹ In their most basic form, those procedures estimate the number to be procured as the sum of (1) the expected number being repaired on the base, (2) the expected number undergoing repair off base, and (3) an additional number to hedge against stochastic variations in the demand.

After all data have been entered, subroutines COMPRT (compute parts) and IPARTS (initialize parts) are called by subroutine TRIALS to carry out these computations if the control variable OUTFIT is not zero. The estimates are made on the basis of (1) the parts-procurement-policy planning factors that the user enters using CT23/70 and CT23/72; (2) the expected daily demand rate for each part based on the task and parts-repair probability data entered with CT5, CT7, and CT8; (3) the NRTS data entered for each part² with CT23/20x (and CT23/30x); and (4) parts cost data entered with CT23/66. If desired, the user may specify different safety stock factors for LRUs and SRUs, and for those tasks that may be deferred indefinitely and those that may not.

For units that are deployed to the theater, the nominal parts allowance or WRSK may be computed by either of two procedures. In the first procedure, the allowance is computed on the basis of 30-days supply at the planned wartime sortie rate for the RR (remove and replace) items, and the same as for BLSS for the RRR (remove, repair, and replace) items. A 30-day supply of SRUs that are not repairable is included for LRUs that are RRR; stock levels for RRR SRUs are computed in a manner analogous to the LRU computation. With the second procedure, used when the control variable PMODE (CT3/3) is greater than zero, the WRSK allowance is computed in accordance with an empirical algorithm that approximates the cost optimization procedures used in the AFLCM 171-46 [10].

¹The user may modify these computed stock levels to reflect stock shortages or expected battle damage, etc., by entering the additional stock with the basic CT23. As now structured, 500 part types may be modified in this manner at each base. The NRTS rate specified with these cards will override any value entered using the CT23/20x or CT23/30x cards if the control variable CHNRTS on CT3/3 is initialized to unity; a null entry on the basic CT23 cards will be interpreted as a zero NRTS rate.

²When the same type of part is used in more than one location on an aircraft (e.g., a left and right tire, a left and right engine), a different part number is assigned to each location. Parts in the several locations are identified as the same part with the CT35/4 cards; data pertaining to procurement and repair actions need only be provided for the part identified as the "prime" part. Equivalent procedures are used when the same SRU is used in multiple locations in one or more LRUs.

If the user desires to define parts shortfalls over and above those that are in the pipelines, three options are provided. In the first option, the actual number of each type of part that is procured for a base can be reduced by a fixed percentage that the user specifies with the control variable SHORT. A second option permits the user to simulate shortfalls differentially for various part types, and the third option is to employ the two options together. Either or both types of shortage may be used to simulate the parts environment that the user judges to be most realistic. The actual shortfall for each type of part will be the expected value of the shortage if RANDM is zero, or will be drawn from a Poisson distribution if RANDM is unity. If NEWPRT is initialized, the parts initialization computations, including these considerations of shortages, are redone each trial.

The number of serviceable items on base for each part type is set equal to the number procured, minus the nominal number that would be expected to be in the pipeline. In other words, it is assumed that there are no on-base reparable. The number in the pipeline (i.e., being repaired off base) is the largest whole integer in the value developed in the prior computation or, if RANDM is unity, a number drawn from a Poisson distribution with a mean equal to that value. If the number estimated for the pipeline is larger than the number that had been procured (taking shortages into account), the pipeline number is either reduced to the number available or (when ZNORS = 1) the difference is made up by removing the parts from on-base aircraft at zero time (thereby generating NMCS aircraft).

The actual formatting of the parts stockage data generated by subroutine COMPRT is that for CT23 and CT31—i.e., as for user-specified parts inventory data and for shipments from CONUS.

Under some circumstances a user may not want to use the automatic parts generation routines for all runs but may wish to take the results of those computations and reuse them as input for the model. This could be useful either to avoid repeating that calculation for a large number of runs or to permit him to combine parts computed in two different ways at a single base.³ Whenever PPRINT is set to 30 or more, model execution

³On some occasions the user may wish to represent a base that has one unit of aircraft stationed there in peacetime and an additional unit deployed there before hostilities start; the former would have POS and BLSS stocks and the latter would have brought along their WRSK. Normally, the automatic parts generation routines will not accommodate the required calculations if both aircraft are of the same type, but it can be done if a CIRF

is terminated at the end of the parts initialization computations just after the results of those computations have been organized in the format specified for the basic CT23 and CT31 cards. By storing the card-image copies of the appropriate part of the output, the user can have the needed CT23 and CT31 cards available for subsequent use.

As discussed in Sec. V.9, aircraft spare parts for rear maintenance bases are either entered directly (with the basic CT23 cards) or, when the automatic parts generation feature is being used, they are provisioned by redistributing the spares that have been calculated for the operating bases. For tasks that must be done in the rear, all parts are placed in the rear. An estimate is also made of the fraction of the other tasks that will be accomplished at the rear base at the same time the mandatory work is underway, and a like fraction of all parts is placed in the rear. If aircraft are also sent to the rear whenever the ready-to-fly time exceeds MNTLMT, etc., the fraction of the parts placed in the rear can be increased by the user's specification of RPARTS.

When the user is examining CIRF operations, other considerations affect the parts initialization process. For the procurement computation the user may (1) neglect the effect of the CIRF on NRTS rates and (2) ignore any advantages of scale in the SRU computation, or may take one or both into account. These choices are controlled by the value of the control variable OUTFIT. If OUTFIT is unity, the NRTS rates that are used for computing the number of parts to be procured for each base are those that would apply if there were no CIRF; and the number of SRUs is the sum of those computed for the individual bases, even though all the LRUs may be repaired at the CIRF. This mode (OUTFIT = 1) permits the user to stock a set of bases at levels identical to those that would be estimated if there were no CIRF. If OUTFIT is set equal to 3 or 4, the procurement computation presumes those NRTS rates that apply with a CIRF (the data entered with CT23/30x); if it is set equal to 2 or 4, the safety factors in the SRU procurement computations reflect the scale advantages to be expected when the demands for several bases are consolidated at a CIRF.

is not being represented and if the TSAR storage area has been dimensioned for at least one more aircraft type than is being used in the simulation. When these conditions are met, the second area of the POLICY array (i.e., the CT23/3xx cards) may be used to store the required NRTS data for the second of these units; in addition, a CT23/70 card should be entered that specifies the base "kind" to be 3, and an aircraft type that is otherwise not used. The user must then also duplicate the CT7 for the aircraft type that is to be stationed at the base, but mark them as though they were for the "unused" type specified on the CT23/70. If the user wishes the nominal breakrates to be modified according to entries on CT44 cards, the CT44 must be entered for both the nominal aircraft type and for the "unused" type.

The authorized level of stock computed for each base assumes that all serviceable LRUs are at the operating locations. SRUs, however, are allocated in the same proportions that in-theater work is accomplished on their parent LRU. Thus, without a CIRF, all parts are at the operating bases, but when a CIRF is introduced, some of the SRUs will be at the base and some at the CIRF for LRUs that are partly repaired on base and partly at the CIRF. When certain aircraft maintenance tasks must be carried out at a base in the rear, any parts used with those tasks are emplaced at the rear base; furthermore, if the user's choice of JOBCON indicates that other tasks are to be done in the rear whenever the aircraft is there, the portion of the parts that are appropriate for the tasks expected to be done in the rear are also retained at the rear base.

After the nominal parts level and the available number of serviceable parts have been computed and stored for each type of part at each base, the parts pipelines are initialized. When there is no CIRF, the parts that are in the pipeline are scheduled for delivery within the user-specified (CT23/70) order-and-ship time, with the actual day picked at random for each item. When a CIRF is assumed to be present, there will be some items in the base-CIRF-base pipelines and others in the CIRF-CONUS-CIRF pipeline. The mean numbers in each pipeline for each type of part are estimated on the basis of the user-supplied data regarding the various times and the daily demands generated at the operating bases. Items are then positioned in both pipelines for delivery after the simulation is begun.

2. INITIALIZATION OF STOCKS FOR BATTLE DAMAGE REPAIRS

Parts also may be stocked automatically for repairing battle damage sustained in air operations. The quantities stocked at each base are based on a specified number of sorties for each of a specified number of aircraft, and on the battle damage rate expected on the average during the first 30 days of conflict (assuming the various mission types are flown equally). The number of aircraft is the initial number on base or, when OUTFIT is not zero, the number of aircraft specified for the spares stockage algorithms. The number of sorties is entered with CT15/2. The condemnation rate for parts removed in connection with battle damage repair tasks is assumed to be 100 percent.

The stocks of these battle-damage spares allocated to the various operating bases take into account any task specifications that mandate the task be accomplished at a rear base. The allocation also takes into account (at least approximately) the likelihood that

some tasks normally done at the operating base will actually be cleaned up when an aircraft is in the rear for mandatory rear-area maintenance.

3. ON-BASE PARTS REPAIR

Whenever an attempt is made to initiate an on-equipment task and a faulty part is found, or a faulty SRU is found during the repair of an LRU, parts that are never repaired on base may be NRTSed immediately; otherwise, parts are set aside for a delay time before the actual repair process may be initiated.⁴ The delay is determined in subroutine ADMIN and is equal to the sum of the mean time for the on-equipment task (to simulate the time for removal) and an administrative delay. (The user specifies the mean and distribution for this delay by shop and by base, using CT47.) When that delay is completed, the NEWREP entry point in the INIREP subroutine is called. (If the variable EXPEDite is initialized on CT4/1, and there are no serviceable parts of the required type, the administrative delay is reduced to 1/EXPED of the nominal value, or to zero if EXPED exceeds 10. This feature permits the user to simulate an organization in which the time required to process a reparable can be expedited when necessary.) However, if the shop has been closed by an air attack or if the needed AIS stations have been lost in an air attack, alternate repair locations are checked; if appropriate, the faulty part is shipped to another location.

When the entry NEWREP is called, a check is first made to see whether the part will have to be repaired elsewhere (is to be NRTSed), or whether it can be repaired on base; this is done by comparing a random number with the NRTS rate. The resources required for the repair process are determined next. One or more procedures may be specified for each type of part and each procedure can be composed of one or more sequential steps. The first procedure is assumed to apply when it is determined that the part is to be NRTSed, unless it was NRTSed immediately on removal from the aircraft. If the part is to be repaired on base and has two or more possible repair procedures, the identity of the required procedure is determined with a random number using the data provided on the CT8 cards as to the relative likelihood that one or another of the procedures is required. One of these repair procedures could be used to represent the

⁴if the NRTS rate for a part, LRU, or SRU is 101, it is shipped immediately upon removal from the aircraft (or from the LRU); if the rate is from 1 to 100, any decision to ship the unit is made after an administrative delay.

checks that are required even when the apparent fault is not found (i.e., CND—could not duplicate—parts). Each step of each parts repair procedure can specify requirements for a number of one type of specialist, one or two types of equipment (including a particular AIS station), and time; if the part is an LRU that may have a defective SRU, each SRU is specified by including it as an additional requirement in the first step of an LRU repair procedure. The specifications for each step of a parts repair procedure may also include an indication that cross-trained and/or task-assist qualified personnel may be employed, reference to a substitute procedure, and a heat factor.

The next step is to check whether the shop has been closed by air attack and, if not, whether the necessary personnel⁵ and equipments are available. If the normal resources for the repair are not available, and the resources for any alternate means of accomplishing the job are not available, the repair must wait; when the resources are available, parts that are to be NRTSed are consigned for shipment with a call to subroutine NRTSIT, and the required personnel and equipment are committed for the specified time (the timing error in dispatching the part before the time has expired is neglected for convenience in coding). When the effects of chemical protection are being considered, the length of time that the repair personnel may work before stopping to cool off is determined with the call to TTIME and CWTIME; this determination takes into account the workers' MOPP (dictated by the agent vapor pressure within the facility) and the temperature and humidity at the work place.

If the part is to be repaired on base and an SRU is defective, the faulty SRU is withdrawn during the first step of the LRU repair procedure and is placed in a two-hour administrative delay. Then checks are made to see if a serviceable SRU is in stock. If none are available and an aircraft is NORS for the LRU, the user may specify (by setting CANSRU > 0) that another LRU of the same type may be sought in the wait queue and disassembled to obtain its serviceable SRUs if it doesn't require the same SRU—i.e., it may be cross-cannibalized. Subroutine SALVAG searches the wait queue and carries out the cross-cannibalization. If the repair job still cannot be started, because of the shortage of an SRU, it is placed in the wait queue of the appropriate shop. If the user has specified that jobs that must wait are to be prioritized (by initializing ORDWT = 1), the

⁵If the data base differentiates between flight-line specialists and back-shop repair personnel, and repairs are to be conducted at a base where the personnel are not organized in that manner, the personnel requirements are interpreted in terms of the equivalent flight-line specialist.

repair job is placed in the wait queue (using subroutine WAIT) according to the magnitude of the variable RANK.

To determine RANK, one first computes the current value of

$$\text{VALUE} = 2 \times (\text{HOLES} - \text{SERVICABLES}) - \text{ENROUT}$$

using the on-base values for the part. He also computes IMPORT, a function of the total number of mission types for which the part is critical. For positive values of VALUE,

$$\text{RANK} = - (100 + \text{IMPORT}) \times \text{VALUE}$$

and

$$\text{RANK} = - \text{VALUE} \times \text{MTBF}$$

for zero or negative values, where MTBF is the average number of sorties before failure of the part.

When parts repairs are ranked in this manner, the parts needed for the most aircraft on the most missions have the most negative number, and the parts least likely to be needed have the most positive number. Equipment repairs are given a RANK = 0, on the assumption they are less necessary than parts needed to release an aircraft, but more necessary than a part that is not yet needed. When resources are available to begin a repair, the queue is searched from the smallest value (most negative valued rank) to the largest (most positive).

This procedure is followed, with two exceptions. First, if the required resources are not on-base, the part is placed in the queue with a RANK of 32750, at the end of the queue. Second, if the repair requires an AIS tray that is not functioning, and it is the only station on base of that kind, the RANK is set to 32600. Then, when repairs are checked in subroutine CHECK, the search through the queue is stopped if the RANK is 32600, or after 100 parts have been checked if the RANK is as great as 1000.

When the necessary resources are at hand to initiate the job, subroutine DOREP is called and the repair job is entered in the time heap associated with the REPQ array. If the part for which the resources have been committed is NRTS, the repair job is flagged by specifying the negative value of the repair procedure. The DOREP subroutine is also used when it is necessary to interrupt an on-going repair job. When that occurs, the job is transferred from the REPQ array to the INTTSK array, the SHOPS array pointers are updated, and the personnel and equipment that had been engaged are released. A special

provision is included to deal with the problem of terminating a repair for which the part itself was destroyed during an airbase attack.

The INIREP subroutine is also used when resources are released and an attempt is made to start parts repair jobs that have been interrupted or are waiting. The resource requirements are checked, and if the job can now be started, the INIREP subroutine updates the various pointer systems related with the INTTSK, WAITSK, and SHOPS arrays.

When the administrative delay for an SRU is completed, entry NEWREP is called and checks are made to see whether it is to be NRTSed or may be repaired on base, much as for an LRU. Checks are next made to see if the required personnel and equipment, or the personnel and equipment for substitute procedures, are available to start the repair procedure. If they are not, the repair must wait; if they are, then the SRU is NRTSed when appropriate and the personnel and equipment committed for the specified time, again much as for the LRU.

When a step of a parts repair job has been completed, control is transferred from subroutine MANAGE to either subroutine RUNSHP (run shop), or, if USECW > 0, first to STOPIT so that the condition and needed rest for the personnel may be determined. If the repair procedure has additional steps, a random number is compared with the probability that the subsequent step is required. Then, after a call to subroutine ENDREP to release personnel and equipment and to update the pointer systems used with the REPQ and SHOPS arrays, subroutine INIREP is called to initiate any subsequent work. If none is required, the part or rebuilt LRU is put into stock. Unless the special parts disposition logic is applicable (i.e., unless SHPREP > 1, see Sec. XI), the repaired part is retained if it was removed on base or returned to the base where it was removed. When it is retained on base, and when there are aircraft on base that require a part, subroutine CHECK is called when control returns to RUNSHP. If an aircraft is still waiting for the part, the appropriate on-equipment task is initiated. When the part was not removed on base, or if the special parts disposition logic selects another base, the part is shipped to the appropriate base. Similarly, when an SRU repair is completed, resources are sought to repair an LRU requiring that SRU. When control again returns to RUNSHP, subroutine CHECK is called again with the shop number to be sure that the newly released personnel and AGE are reassigned if they are needed.

The RANK of each waiting repair is reevaluated twice each day in subroutine REPRTY. As each waiting part is checked and its position in the queue adjusted, the number of HOLES that is used in the computation for that part type is reduced by one, so that all parts of the same type will not all be at the top of the queue. The periodic ranking of a part at a CIRF sums "VALUE" for that part type across all airbases and ranks the repair in a manner analogous to that described above. Intratheater reports of resource status are used in these algorithms when they are generated.

4. OFF-BASE PARTS REPAIR

When a faulty part is found to be NRTS, a check is made as to where it is to be shipped for repair. Based on the data supplied by the user with CT34, different destinations may be specified for each type of part, subject to the data limitations outlined for that Card Type in Sec. XIX. If there is a central repair facility in the theater, TSAR assigns it the base number MAXB. If a part is to be NRTSed to a depot outside the theater, the destination should be entered as (MAXB + 1)—i.e., one greater than the largest numbered base.

For RR items (an item with NRTS \leq 100), an option has been provided to permit the nominal shipping instructions to be overridden when the number of serviceable LRUs falls below a specified percentage (ADAPTR) of the base's initial number of LRUs. When this occurs, the list of bases specified for lateral resupply is checked to find a location that is able to repair the item (NRTS < 100) and has an undamaged shop. This option can be used, for example, to simulate an adaptive parts repair doctrine that discontinues reparable shipments to the depot and attempts to accomplish the repair in theater, when parts stocks are low.

A faulty part may also be shipped to another operating base, even though it would not normally be NRTSed, when the shop in which the repair must be done has been closed by damage from airbase attack. When this occurs the lateral resupply base list is checked for a base with the shop open and a lower NRTS rate for the part in question; if a base is found, the part is shipped to that base if the two-way shipment time is within one day of the reconstitution time for the damaged shop.

If the part is shipped to another operating base for repair, the part is treated just like any other job generated at that base and begins by undergoing an administrative delay. The number of the originating base is preserved so that the part may be returned

when repairs have been completed if the special parts disposition logic is inoperative. Depending upon the NRTS rate for that type of part at the receiving base, the part could be shipped to yet another base; if it is repaired at that base, it will be shipped back directly to the originating base when repairs are completed, unless the disposition logic is operative and selects a different destination. It is left to the user to design the CT34 inputs such that a faulty part will not be NRTSed from one base to another until it arrives at the originating base.

If a part is condemned or is shipped out of the theater, its replacement, when one is specified, is consigned for delivery directly to the base of origin even though a CIRF may be operating, unless the control variable CONSIG is initialized to unity. In the latter case, all parts returned from CONUS are consigned to the CIRF for transshipment according to the user-specified theater resource management algorithms.

When a part is shipped to a centralized intermediate repair facility in the theater, it is subjected to an administrative delay but is then managed by a different set of rules that govern the priority it receives and its disposition when the repair station is completed. These will be outlined fully in Sec. XI after the properties of the transportation and information nets used in connection with these operations are explained in Sec. X. Parts repairs cannot be expedited at the CIRF, but the user can control administrative delay times and parts repair times on a shop-by-shop basis to account for the different working conditions at a CIRF, using CT47 and CT48.

5. SUPPORT EQUIPMENT REPAIR

Many special kinds of support equipment are needed for the specialized jobs that must be conducted on a modern military airbase. And most of these equipments are both complex and expensive; malfunctions are fairly frequent, and their maintenance and repair constitute an essential set of activities. Such malfunctions and the repair of faulty equipment may also be simulated in TSAR.

Support equipment repairs are handled in much the same way as spare part repairs, and with many of the same subroutines and procedures. However, TSAR provides two quite different representations of equipment failure and repair. The simpler representation is used for all equipments other than the AIS—Avionics Intermediate Shops—those complex test equipments that are used to test and repair avionics on late model aircraft. The basic distinction is that in the simpler representation, equipments are

either serviceable or they are not; AIS equipment may be partly mission capable as well. Both representations are described below.

Equipment Repairs Other Than AIS Sets

Whenever a task that has used support equipment (other than an AIS set) has been completed, each item of equipment is checked to see if it needs maintenance by comparing a random number with the probability that that type of equipment will require maintenance following a job. If maintenance is required, the equipment first undergoes an administrative delay, much as for spare parts although the length of delay is different. When that administrative delay is completed, the attempt to initiate the repair is processed in the same subroutines as a faulty aircraft part. Each type of equipment is associated with a particular shop, and the repair procedure may either be specific or chosen at random from among a set of alternative procedures. And as with spare parts, each repair procedure may consist of a sequence of steps. Each step of an equipment repair procedure may specify a type and number of personnel, one or two pieces of repair equipment, and a duration; substitute personnel and/or alternative procedures may be specified for consideration when the normal resources are not available. But these specifications do not include the spare parts that might be needed to repair the equipment; such problems can be approximated, however, by specifying that equipment repairs can be carried out without delay for parts on some occasions, while on other occasions they are subjected to a delay equivalent to the order and ship time for spares.

If resources are available when an equipment repair is first attempted, the resources are assigned to the repair, the completion time is established, and the job is placed in the repair queue, REPQ; if resources are not available, the job must wait. Equipment repairs that must wait currently are entered in the shop wait heap with RANK = 0; if equipment and parts are competing for the same repair personnel and equipment, the equipment repairs are given priority over spare parts for which serviceables are available but they will follow the repairs for parts needed for work on aircraft. As currently structured, all equipment repairs are performed on base; equipments are not NRTSed to other bases.

Simulation of AIS Maintenance and Repair

The specialized support equipment used for testing and repairing avionics on late model aircraft—the AIS—also may be simulated in TSAR. A full "string" of AIS will normally have several different complex electronic test equipments, or "stations," and each type of station is used for testing several different LRUs. Each station is composed of many hundreds (thousands) of submodules, and these stations are themselves subject to various malfunctions that can require substantial maintenance. Furthermore, when any of the numerous low failure rate (and therefore unstocked) AIS parts fails, it is necessary to order one from another location, and that station will then be able to test only some portion of its normal LRUs. Thus, a station will be in one of three states: fully mission capable, partly mission capable, or inoperative. If two or more stations of the same type are available, partial mission capability generally can be minimized by consolidating all missing parts at one station.

The manner in which these characteristics are modeled in TSAR is adapted from the work of Jean Gebman and Hyman Shulman at RAND. Whenever an AIS station is used to repair an LRU or SRU, the nominal part repair time is increased to allow for maintenance of the station itself. Because such maintenance may actually occur either before or after, or even during, the repair of the part, it is assumed that the part is not released until the job is completed. At that time, the LRU is released for use and a check is made to see if any piece part needed for maintenance on the AIS was not in stock. If so, the station's residual capability to repair LRUs is estimated on the basis of statistics that indicate how frequently each LRU repair capability is lost on the average when an AIS part is back-ordered. To do this, we imagine that each station is divided into several sections or "trays," with one tray for each type of LRU; when a part is back-ordered the mission capability of each tray is determined on the basis of the statistical experience.

During the simulation, a check is made following each LRU repair to see whether during maintenance on the AIS station it was found to need a part that is not in stock. If one is needed but there are two or more stations of that type on the base, it is assumed that the needed part will be cannibalized from another station if necessary and that all missing parts are consolidated at one of the stations. Thus, when an AIS part fails at any station, checks are made for each LRU tray associated with that type of station and a list is generated of all LRUs that cannot be repaired until the needed part is obtained. A sample is then drawn from the user specified order-and-ship-time distribution, and the

appropriate receipt time is entered in the LIMBO array; not until that time occurs is the capability restored for repairing those LRUs.

As will be noted, there are no specific repair procedures or specific personnel or equipment used to repair AIS equipments. Instead, the repair time of each part that is processed is increased to account for AIS maintenance, and AIS repair capabilities are probabilistically curtailed to simulate a shortage of parts to repair the AIS.

VII. AIRCRAFT SORTIE DEMAND AND AIRCREW MANAGEMENT

The ultimate objective of an airbase is to provide combat-capable aircraft when they are required, and a base's capability for meeting that objective can depend importantly upon the pattern of the demand. In TSAR, that demand pattern is controlled by the user's input data (CT50), and the user is provided sufficient options that most plausible requirements should be readily simulated.

A demand for a flight of aircraft specifies the type of aircraft, the mission, the mission's priority, and, normally, the base; it also specifies the number of aircraft to be launched (and the minimum acceptable number), the time they are to be launched, the time that the airbase is informed of the demand, and the recovery base. If desired, the user may also specify that a specific number of aircraft will be maintained on alert at a particular base for unscheduled demands. In addition, he may define a composite flight, made up of several sets of aircraft, or flights, each with a differing configuration, as would be required, for example, for representing coordinated attacks by defense suppression aircraft, CAP, and BAI aircraft.

Except for composite flights and specified alert forces, it is not mandatory that the launch base be specified. If the control variables "STATE" and "SELECT" are both greater than unity, a daily forecast is made of each base's sortie generation capabilities, and these forecasts are used to designate a base for any sortie demands for which a base has not been specified. However, since TSAR does not include geographic concepts, such selections are not constrained by range-to-target considerations.

For user convenience the demand data may be stated either on a day-to-day basis or in terms of demands that recur each day with a stipulated probability (or any combination of these techniques). For the recurring demands (the periodic demands) the launch time may be entered as a precise time or as a time block; when a time block is specified, the program picks a time at random from within the block. Furthermore, several such flights (up to 32) may be specified with the same entry; when this is done, the launch time of each flight is selected at random from the time block. With these features a few entries suffice to represent a rich and varied pattern of flight demands.

1. GENERATING SORTIE DEMAND DATA

The initial day's sortie demands are entered before the simulation begins. They are entered after all other data are input and after subroutine INLIST has provided whatever listings of input data were requested. Subroutine READFT (read flight data) reads and organizes these data with the assistance of subroutine SORT, which orders the flights by their specified launch times (or required shelter departure times, when a taxi time is specified) and manages the pointer system used with the FLTRQT (flight requirements) array. If the launch base has not been specified for any of the sorties demanded, subroutine FRAG is called to select the base best able to fulfill the demand, the one with the lowest current level of demand relative to its estimated sortie generation capabilities. When all data have been entered, flights with common characteristics (launch base, aircraft type, mission, and priority) are interconnected with the pointer system associated with the FTZ array.

The sortie demands for the next day and for subsequent days are also managed in the READFT subroutine. These demands are reexamined each evening at 1945 simulated time when this subroutine is called by subroutine MANAGE. If the user wishes to specify new flights or to change specifications for alert forces or periodic flights, these data are read at this time. If there is no new information, the following day's demands are based on the periodic flight demands or other flight data submitted earlier. As before, any flight demands for which a base has not been specified have a base chosen with the FRAG subroutine, using updated estimates of the bases' sortie generation capabilities, which are created daily at 1930 by subroutine BASCAP (base capabilities).

If, when the sortie demands are organized for the following day, an airbase is out of operation because its runway is closed, the demands on that base may be reassigned. If the runway is projected to remain closed for any part of the following day and other bases have aircraft of the type specified, demands that are required to be met before the runway is to open are reassigned either by subroutine FRAG, just as though the launch base had not been specified, or, if SELECT is zero, in proportion to the numbers of aircraft on base. Demands to be met after the runway is scheduled to be opened are not reassigned.

Provisions have also been made for entering endogenously generated flight demand data. These provisions would be used if and when the resource management

logic is expanded to permit endogenous decisions regarding sortie demands. Such a decision would be communicated by calling the entry point SORTIE in the READFT subroutine where the flight would be entered into the sortie demand pattern. If the base is not specified, subroutine FRAG selects the base best able to fill the demand.

2. LAUNCHING THE AIRCRAFT

When the time specified for launching aircraft occurs, subroutine MANAGE transfers control to subroutine FLIGHT. After checking that the flight need not be canceled because of weather conditions or runway damage, a check is made to see if aircraft have been assigned for a scheduled flight or are available in the alert force when the demand is unannounced. Each aircraft is checked to see if it has actually been readied for flight and has not suffered a ground abort. A check is also made for each aircraft to see that its access to the runway is not prohibited by bomb damage to the taxiway network. If aircrews are to be accounted for, subroutine FLYERS is called to locate a crew that is then tentatively assigned to the aircraft.

If fewer than the required number of aircraft are ready among those assigned to meet the specific demand, and if the demand has a priority at least equal to the highest deficient priority (as defined in Sec. IV.19), the spare forces, later flights of the same or lower priority, and alert forces of lower priority are each checked in turn for a ready aircraft of the appropriate type and mission configuration. If, after all these sources are checked, the number of aircraft available to meet the demand is less than the minimum permissible number, the assigned aircraft and then the spare aircraft are checked to see if aircraft are available that will be ready within whatever time is allowed for late takeoff for aircraft of that type on that mission. If the minimum permissible number of aircraft have still not been located, the flight is canceled. If their number is sufficient, the constraints imposed by air traffic control will be checked when the user has initialized DOATC. If the runway is already scheduled for another flight to take off or to recover at the same times, a check is made (with subroutine USEATC) for a takeoff time slot within the allowable late takeoff time. If one is found, a check is made of runway availability when the flight would be expected to return, or within a user-specified time thereafter. If times are available, the flight is launched with a call to subroutine LAUNCH that updates all the appropriate tallies and pointers. If times cannot be found when the runway will be available for takeoff and recovery of the full flight, a check is made to see if fewer aircraft (but at least the minimum number) can be handled. If not, the flight is canceled.

If an aircraft is designated to recover at a different base, that bookkeeping is also accomplished at the time that the aircraft is launched. When it has been determined that the aircraft may be launched (or depart from its shelter), it is placed in the aircraft delay heap with the appropriate landing time. The flight times for each aircraft in a flight are determined independently, unless recovery as a unit has been specified on CT16 for that type of aircraft and mission, or unless the air traffic control feature is in use.

Composite Flight Requirements

Certain additional complexities will be noted in the FLIGHT and LAUNCH routines as a consequence of the options for composite flights and for late takeoffs. When the minimum forces and takeoff and recovery times must be found for each of several different flight demands to prevent all from being canceled, it is necessary to withhold all launches until all flights have been checked. Furthermore, if, after checking several flights, it is found that at least one cannot be satisfied, it is necessary to modify various aircraft assignments and to release all tentatively assigned air crews. Similarly, when an aircraft is going to be launched late, it is necessary that certain data be retained until that time. To facilitate the latter operation the aircraft is placed in the aircraft delay heap until one TSAR time unit after the aircraft's expected ready-to-fly time; if it is still not ready to fly at that time, the sortie is canceled.

Air Aborts

As each aircraft is launched it is checked for an air abort; if one is to occur, the aircraft is scheduled to land with a full load of munitions at the launching base after a six-minute flight and the base is *not* credited with a successful sortie. It is handled like any other aircraft in the ensuing postflight inspection except that munitions are not required and attrition and battle damage are not assessed.

3. AIRCREW MANAGEMENT

TSAR's provisions for accounting for aircrews are controlled by the control variable CREWS; when initialized to 1 on CT1, these features are activated.

Aircrew members are accounted for on an individual basis, much like aircraft. Each aircrew is qualified for only one type of aircraft. Their assignments are managed so that each crew will receive a specified minimum amount of uninterrupted sleep during

each 24-hour period and at least the specified minimum rest between sorties. These two times are specified with the control variables SLEEP and REST. In addition, the required time between sorties may be increased for particular missions to account for the additional time required to study target data and take intelligence briefings. To avoid unnecessarily long shifts and early exhaustion it is presumed that aircrew assignments can be managed such that they remain off duty until they are needed and will retire early whenever the demand permits.

Aircrews are managed with data stored in the PILOTS and PILOT arrays. PILOTS maintains a record of the number of aircrews on base and pointers to the first and the last of those crew members who are on duty and off duty; these data are maintained separately for each aircraft type on each airbase. The PILOT array maintains status information on individual aircrews and pointers to the other crews with the same duty status.

The several operations required for aircrew management are carried out by different sections of the FLYERS subroutine. An aircrew is located for a tentative assignment by calling the entry point GETPLT (get pilot) or SAVPLT (save pilot) for a late takeoff. When the aircraft is launched, the assignment is finalized with a call to the entry point FLYAC. When the sortie has been completed, crew data are updated with a call to LANDAC; if the aircraft is recovered at a different base, the pilot "billeting" is rearranged at that time. If the crew is due for a sleep period, they are placed off duty; if the aircraft was lost but the aircrew was not, it is assumed that the crew cannot be reassigned for a minimum of four days.

In addition to these operations, entry point RELIEF is called at two-hour intervals by MANAGE to check the on-duty crews and to relieve them as required. If the control variable RELIEV is zero (CT4/1), aircrews are available for the full shift once they are called for duty; if RELIEV = 1, aircrews are presumed to go off duty immediately after their last flight of the day. When the airbase has been attacked and the user has specified the location of air crews in the TSARINA data base, subroutine DISABL is called by subroutine BOMB to inflict whatever losses were inflicted and to update the aircrew information. If the aircrew are located by TSAR default location assumptions (see Sec. VIII.2), subroutine REORGN handles the necessary accounting at the time of an attack.

VIII. AIRBASE ATTACK AND RECOVERY

The most serious disruption that an airbase can experience is undoubtedly that associated with a heavy airbase attack. The lack of any generally agreed-upon estimate of the effects of attacks on a base's capabilities to recover and generate useful aircraft sorties were prime motivations for TSAR's development. And the highly variable damage patterns and chemical deposition patterns that would be experienced on bases that are subjected to attack contributed importantly to the decision to create a model with sufficient detail that the critical effects of these highly variable patterns could be captured. Without considering the alternative procedures that could be used and the bottlenecks that could arise when resources are lost, one could hardly hope to represent the probable behavior of an airbase during the crisis following an attack

1. ATTACK EFFECTIVENESS ESTIMATED WITH TSARINA

In TSAR, airbases are attacked and resources are damaged or destroyed based on the results generated by TSARINA calculations. The user is free to schedule attacks at whatever times and at whichever bases he chooses, and TSAR is structured to accept fairly detailed input damage data.

The companion TSARINA airbase damage assessment model uses detailed descriptions of the location, contents, and vulnerability to conventional weapons of various airbase facilities, as well as detailed specifications of enemy attacks and weapon characteristics.[3] This Monte Carlo model estimates damage and casualties due to conventional weapons and computes the surface contamination and vapor concentration that result from chemical weapons. The current version of TSARINA permits assessments of damage to an airbase complex composed of up to 1000 target elements (runways, taxiways, shelter buildings, etc.) with up to 2500 packets of resources. The targets are grouped into as many as 30 different vulnerability categories; and the locations of the different types of personnel, equipment, munitions, spare parts, TRAP, building materials, and POL can be distributed among the target elements.

A single attack may involve as many as 100 weapon-delivery vehicles and up to 20 different types of weapons. Point-impact weapons (such as GP bombs, precision-

guided munitions, and BKEPs), UXO, and area weapons (such as cluster bomb units and mine dispensers) can be represented, as well as chemical bombs and warheads delivered by surface-to-surface missiles or aircraft, and chemical sprays released from aircraft.

TSARINA determines the actual impact points of point-impact weapons and the centroid of area weapons by Monte Carlo procedures—random selections from the appropriate delivery error distributions. Conventional weapons that impact within specified distances of each target are classed as hits, and estimates of the damage to the structures and to the various classes of support resources are assessed using "cookie-cutter" weapon-effects approximations.

For chemical weapons, TSARINA estimates the time of arrival of the first chemicals at each target, as well as the distribution of surface contamination and vapor concentration for each of up to three chemical agents that may be used in a particular simulation. These estimates also use Monte Carlo procedures to reflect uncertainties in the direction and velocity of the wind, as well as in the delivery accuracy.

2. REPRESENTATION OF AN AIRBASE IN TSAR

An airbase is represented in considerably more detail in TSAR than in the Early-TSAR formulation. Although neither version of TSAR incorporates airbase geometry explicitly, TSAR now captures many of the location-dependent effects. And the interdependence of TSAR and TSARINA is substantially greater than with the earlier versions.

A primary difference is that the network of taxiways that interconnect the runways with the aircraft shelters and parking ramps on an airbase is now treated explicitly in TSAR. User-supplied input information specifies the arc-node structure of the taxiway network (CT17/4), and the nodal location of each aircraft shelter (CT17/5) and parking ramp (CT17/6). Each segment (arc) of the taxiway network, each aircraft shelter, and each parking ramp is entered as a target in the TSARINA data base, and the user also specifies a set of monitoring points in the TSARINA data base that are used in conjunction with the chemical attacks.

To simplify creation of the needed input data and to minimize the possibility of errors, a detailed sketch should be drawn for each airbase. The taxiway arcs, aircraft shelters, and aircraft parking ramps should be numbered consecutively and the corresponding TSARINA target cards should be entered into the TSARINA data base in

the same order; Arc #1 would be the first taxiway segment in the data base; Shelter #1, the first shelter; and Ramp #1, the first parking ramp, etc. The nodes *must* be numbered consecutively, and each taxiway intersection and the end of each taxiway stub should be assigned a unique node number, starting with #1, but the order in which they are numbered is arbitrary. One target type applies to all taxiway segments and another target type applies to all aircraft parking ramps; three different types of aircraft shelters may be designated at each base.

The airbase sketch for the sample problem outlined in Sec. XX of Vol. II is presented in Fig. 8 to illustrate these conventions. If the effects of chemical attacks are to be evaluated, a set of monitoring points must be selected, remembering that the ambient conditions at a monitoring point will serve as the ambient conditions for all targets closest to that monitoring point for assessments of the persistent effects of chemicals between attacks. Clearly the monitoring points should be selected near those on-base locations of importance to the evaluation; a uniform grid will generally be inappropriate and wasteful of the substantial processing associated with updating the chemical conditions at each monitoring point.

The new data requirements imposed on TSARINA are the consistent ordering of the target cards, as just described, and the specification of the monitoring points when CW attacks are to be evaluated. Other new data input options include specifications that define the delivery of UXO and mines, the delays before the UXOs are to detonate, and the characteristics of chemical weapons. Another option permits specified attackers, in attacks subsequent to the first, to be targeted at the MOS selected in TSAR after the prior attack, thereby simulating effective post-attack intelligence gathering by the attacker.

The on-base layout of the target elements is specified for TSAR by entering TSAR card types that define (1) the numbers of the nodes at each end of each arc and the arc length (CT17/4), (2) the number of the node that each aircraft shelter (CT17/5) and parking ramp (CT17/8) is nearest to, (3) the numbers of the arcs that jointly constitute each runway (CT17/6), and (4) the proportions of unsheltered aircraft that would be parked on each ramp (CT17/8).

Location of Aircraft

At the beginning of a TSAR simulation, each aircraft will be assigned a shelter, if sufficient shelters are available. This assignment takes into account (1) the number of user-designated "alert" shelters, for which "alert" aircraft are given priority, (2) the




Fig. 8—Layout of runways and taxiway network

average number of aircraft permitted per shelter, and (3) whatever types of aircraft the user has designated as not able (or not to be permitted) to occupy a shelter. Aircraft not assigned shelters are assigned to a parking ramp, based on the capacities of the parking ramps. When DOSHEL has been input as 1, aircraft vacate their shelter when they are launched and are reassigned a shelter, if available, when they return. Unsheltered aircraft are moved to an empty shelter when it is vacated if they are not a prohibited aircraft type.¹ If an aircraft is parked at a hot-pit fuel hydrant after recovery, that location is also known (if designated by the user). Thus, at the time of an attack, TSAR knows where each on-base aircraft is located (unless it is taxiing to or from the runway).

Locations of Activities and Resources

The on-base locations of the various sortie generation activities and of the supporting resources are used to determine damage at the time of an attack and, when chemical protective ensembles are worn, to define the working environment for the activities. For most resources, their locations can be specified only by describing the percentage of each type of resource associated with each target in the TSARINA data base. But for personnel and equipment, resources that move about an airbase in doing various tasks, the user has the option of letting TSAR determine where these resources are located at the time of an attack (and subsequently determine the damage to these resources), based on TSAR's knowledge of the ongoing activities. This latter option—to use the TSAR "default locations"—is used whenever the on-base locations for certain types of personnel and equipment are *not* specified in the TSARINA data base.

When the locations of various percentages of particular types of resources are specified in TSARINA, TSARINA is able to assess the total fraction of those types of resources that have been lost to conventional effects in an air attack and to transfer that information to TSAR, where that percentage is applied to the then-existing stocks of those types of resources. This procedure for locating resources is the only option available for assessing losses to munitions, TRAP, serviceable spare parts, building materials, and POL. Since each type of resource (e.g., MK-82 bombs, AM-2, TERs) may be distributed in several locations, each with its distinct vulnerabilities, the vulnerability of these resources can be represented rather realistically (assembled munitions can even be distinguished from their components).

¹When DOSHEL = 0, aircraft are always assumed to be in their initial location whenever they are not in flight.

Personnel and equipment locations at the time of an air attack can also be specified in the TSARINA data base, instead of TSAR determining their locations as the default condition. But the distribution of personnel or equipment that is specified in a TSARINA data base is assumed to be appropriate at any time of day and at any activity level whenever an attack occurs. There is no way of changing the TSARINA percentile distributions to account for the distinctions between personnel that are working, personnel that are awaiting assignment, or personnel that are resting in a collective-protection shelter. One set of distribution assumptions will be applied for all attacks, unless a completely new set of TSARINA target cards is entered for different attacks. The TSAR default-location assumptions were developed to overcome these constraints for personnel and equipments. And these TSAR-specific default locations are always used in determining the environmental characteristics of the work and resting places when the effects of chemical protection are to be assessed (i.e., USECW > 0). Properly applied, these default locations provide a realistic representation of the on-base distribution of these nonstationary resources. When TSAR uses the default locations for the personnel and equipment, their loss rates are those passed to TSAR from TSARINA for the facility to which they are assigned at attack time.

In brief, the TSAR default location assumptions presume that personnel for each of the several generic types of tasks will be found at their work place when they are working, and in other specific locations when they are off duty, when they are on duty but unassigned, or when they are cooling off. On-equipment tasks occur at the aircraft, parts and equipment repair at the responsible shop, guided munitions assembly at Shop #30, unguided munitions assembly at facility #44, and civil engineering activity at whatever facility or surface is being repaired. Furthermore there is a mechanism (see CT37) with which the user can specify that each nonaircraft activity is actually carried out at a location randomly selected from a set of locations, according to the work load capacity of the locations. This notion of describing a "distributed" capability for an activity can also be used in defining the location assumptions that TSAR should apply for other personnel states (i.e., off duty, cooling off, and awaiting assignment). The next section will define these options in greater detail.

As suggested earlier, the user may choose to specify the locations at attack times of some types of personnel in his TSARINA data base and to depend on the TSAR default locations for the other types of personnel; this is perfectly permissible, but the

user should remain aware that it is the default assumptions that will define the work and resting places between attacks for all personnel types, insofar as defining the environmental conditions for CW assessments.

The user has the same two options for specifying the location (and vulnerability) of support equipment at the time of an attack, except that the means of distinguishing between equipment that is assigned to a task and the equipment that is unassigned are somewhat different. When the location of a type of equipment is indicated in TSARINA, the percentage loss estimate from TSARINA is applied in TSAR to all serviceable pieces (but not to repairable pieces) of that type of equipment, whether they are in use or not. If, however, the user has set ONLYUE = 1 (only unassigned equipment) on CT2/1, the damage calculated for equipments located in TSARINA will be applied in TSAR only to unassigned equipment; equipments that are in use will sustain the equipment loss rate for the location where they are being used. Equipment not located in TSARINA is assumed to be at the shop to which it is assigned when not in use and to be at the work place when it is in use, and damage is evaluated at those locations.

The only additions to these general rules are that the vulnerability of repairable spare parts and equipments and of munitions already committed to an aircraft are determined by the damage sustained at the location of the appropriate tasks, rather than being based on damage generated at locations specified in TSARINA (the damage fraction estimated for "spare parts" at facilities engaged in munitions assembly are applied to the munitions).

TSAR Default Locations and Distributed Capabilities

TSAR assigns all jobs and equipments to specific locations. These locations determine the work environment as well as the location for damage estimation of personnel and equipment (except that the location during an attack may be overridden by a location specified in the TSARINA data). The assigned location may be a specific building or outdoor area, or it may be a set of locations where a particular type of activity would be expected to be carried on. A particular "facility" number is set aside by TSAR for each of various possible activities and functions that the user may wish to represent; all "facilities" from #1 through #50 have TSAR-assigned roles that will be explained in this section. The user may specify that the capability of any of these activities is actually "distributed," in which case the TSAR-assigned facility is only the first of a set of facilities to which that activity will be assigned; the user must define the other members

of these sets (using facility numbers between #51 and NOFAC) and must define the relative capacity of each member of the set (or for backshops an absolute capacity for jobs may be specified—see CT37 in Sec. XIX, Vol. II).

The Early-TSAR capabilities for representing distributed functions still apply but have been expanded. As before, facilities #31, #32, and #33 are assumed to be the locations for on-duty flight-line personnel (from squadrons #1, #2, and #3, respectively) when they are not assigned to work on an aircraft; when assigned, they are either in the appropriate aircraft shelter or on the ramp where the aircraft is parked. Also as before, on-duty backshop personnel who are not working are assumed to be in the shop with which they are associated (Shops #1 to #24); wing-level munitions personnel are in Shops #27 or #28, and guided munitions assembly personnel are in Shop #30. In all cases, the facilities with the same numbers as the shop numbers may each be the parent location of a set of locations into which that capability is "distributed."

Additionally, facilities #40, #41, #42, #43, #44, and #50 are now understood to be the parent locations for the following resource categories:

- Facility #40 Aircrews; either all on-base crews, or if facility #41 is used (see CT17/9) only those aircrews that are on duty awaiting flight assignments
- Facility #41 Off-duty aircrews, when they are to be assumed to be at different locations than on-duty crews
- Facility #42 Unassigned on-duty civil engineers
- Facility #43 Unassigned civil engineering equipment
- Facility #44 Location where unguided munitions are assembled
- Facility #50 All off-duty personnel except aircrews

Table 4 summarizes these default location assumptions that are used by TSAR.

The storage locations in the FACILITY array for facilities #36 through #38 are reserved for specifying the repair procedures for the three types of aircraft shelters. Facilities #46, #47, #48, and #49 are reserved to represent the buildings and equipments used to facilitate the rapid launching and recovery of aircraft. Facilities #34, #35, and #45 have not had a function designated.

Any of the facilities numbered from #1 through #50 (except #36 to #39 and #46 to #49) may be designated as the parent of a distributed set of facilities, each with its own physical and chemical characteristics and its own relative capacity (specified by the user with the #37 Type Cards). The numbers for any of the additional facilities that jointly form a distributed capability must be selected from within the range 51 to NOFAC.

Table 4

TSAR DEFAULT LOCATION ASSUMPTIONS

Personnel Types	Personnel Location			
	Awaiting Assignments	Working	Cooling Off	Off Duty
Aircrews	#40 ...	In flight	NA	#41 ... (optional)
Flight line Squadron #1 Squadron #2 Squadron #3	#31 ... #32 ... #33 ...	Aircraft shelter or 	Work location or #31 ... or CP ... fi	#50 ...
Wing-level load crews	#27 ... #28 ...	Parking ramp		
Backshops and wing-level maintenance	Shop #1 ... through #24 ...	Specific facility	Work location or CP ... bs	#50 ...
Munitions Assembly Guided Unguided	#30 ... #44 ...	Specific facility	Work location or CP ... ma	#50 ...
Civil Engineers	Personnel #42 ... Equipment #43 ...	Specific facility or Taxiway segment	Work location or CP ... ce	#50 ...

NOTES: Locations specified in TSARINA data for specific personnel types override these assumptions at attack time.

Shop #5 ... implies Shop #5 and any locations to which the capability is distributed.

Locations #51 to NOFAC (<399) may be used to designate CPs and distributed capabilities.

Repair of specific part types may be restricted to "parent" shop.

Facilities #46, #47, #48, and #49 are reserved for use in conjunction with air traffic control.

Other facilities in this range may be designated as collective-protection shelters, as discussed later in connection with the CT43/6 card. The user must exercise care that none of these additional facility numbers is defined in more than one way or appears in more than one distributed set, and all such facilities must be present in the TSARINA database so that the relevant "damage" data will be transferred to TSAR.

3. DATA TRANSFER FROM TSARINA TO TSAR

With the new TSARINA input data, along with the data previously required, TSARINA generates and passes to TSAR the following results for each attack:

- An estimate of the conventional damage to all specified facilities and to the personnel, equipment, and spare parts in that facility.
- Estimates of the conventional damage to each aircraft shelter and to the aircraft, personnel, equipment, and spares in the shelter at the attack time.
- Estimates of the fractional damage by conventional weapons to whichever types of personnel, equipment, spares, munitions, TRAP, building materials, and POL have their locations specified in the TSARINA data base (including a damage estimate for fuel trucks that are being refilled at attack time at targets designated as truck refilling locations).
- The numbers of UXO, mines, and craters that prohibit aircraft passage on each segment of the taxiway network.
- Detonation delay times for each type of UXO, and damage estimates for EOD and civil engineering activities at and near the detonations.
- The locations and radii of all craters on each runway and, by virtue of the data noted above, the numbers of UXO and mines on each segment of the taxiway network that is also a segment of a runway.
- The locations, relative to the MOS, and radii of those craters created by attack weapons that were designated as aimed at the MOS determined after the prior attack.
- The time of arrival and the intensity of the initial surface deposition for each of up to three types of chemical agents, and the initial vapor concentration due to that deposition, at each designated facility, aircraft shelter, taxiway segment, aircraft parking ramp, and monitoring point.

- The number of the monitoring point that is closest to each designated facility, aircraft shelter, taxiway arc, and parking ramp.
- A record of any groups of facilities that actually occupy the identical location on the airbase (and hence need to be repaired only once).
- The percentages of each personnel type represented in the TSARINA data base that are associated with each TSARINA target type and with each monitoring point.
- A projection of the future surface and vapor intensities for each chemical agent at each monitoring point.

The transfer of these TSARINA results to TSAR involves a substantially more complex process than did the previous versions of these simulation models. In the original TSAR, the CT40 cards were used for entering all attack damage data, and the formatting rules were straightforward. The data pertained only to conventional attacks and permitted the user to specify damage to each of the several classes of supporting resources (personnel, equipment, parts, munitions, TRAP, FOL, and building materials) and also to specify damage to whatever base facilities were designated. To account for CW attacks as well as conventional attacks, the data transferred with the CT40 cards are now much more complex. Except when a user wishes to represent the simplest kind of conventional damage that does not involve damage to the runways, taxiways, or aircraft, users should not attempt to prepare CT40 cards themselves but rather should depend upon TSARINA to generate the appropriate input for TSAR. The data now called for include not only the conventional damage data mentioned above but also the amount of surface contamination from each chemical agent at each shelter, on each taxiway segment, on each aircraft parking ramp, and at each of a set of monitoring points that the user has specified in defining the TSARINA data base. The CT40 cards are also now used to transfer the fraction of each personnel type that TSARINA associates with each monitoring point and each target type.

The user can no longer use CT40 to specify the number of repairs that must be made to open a runway, because this determination is now made within TSAR, rather than being transferred to TSAR from TSARINA. TSARINA generates an entirely new dataset that is stored on disk and includes location data for all craters on each runway for each attack at each base (or the locations, relative to the prior MOS, for those attackers

designated as attacking the MOS). By storing data for the several attacks and then interpreting them within TSAR, a record can be maintained of all craters, and a crater is eliminated only when it has been repaired. This new dataset stored by TSARINA also contains a compact record of all the chemical deposition data for each of the monitoring points specified in the TSARINA data base. These include the time from the beginning of the attack until the chemical arrives at the monitoring point and the density of contaminants of each agent type that are deposited at each monitoring point. Furthermore, these data include a compact representation of the timewise variation of surface contamination and vapor concentration after the attack at each of the monitoring points. These various chemical data are stored for each component of each chemical burst that is assessed in TSARINA.

Because there is really no practical way for a user to generate the chemical deposition data without the help of TSARINA, instructions on how those data are formatted are not provided. Users who wish to use the new runway/taxiway features of TSARINA and TSAR, or to use the features that permit the analysis of chemical warfare and its effects on operations at airbases, must plan to use the TSARINA model to analyze the attacks and to generate the required TSAR input data.

4. OVERVIEW OF ATTACK TIME COMPUTATIONS

With these data from TSARINA and the knowledge of the taxiway network provided by the new TSAR input data, TSAR generates the following results at the time of an attack:

- Selection of an MOS on one of the flight surfaces, taking into account the UXO, mines, and unrepaired craters that remain from all prior attacks, as well as those added in the present attack.
- Determination of which undamaged aircraft shelters and which aircraft ramps can access the MOS.
- Determination of the preferred taxiway repair schedule to maximize the rate at which undamaged aircraft shelters can gain access to the MOS.
- Casualties due to conventional and chemical effects for personnel or equipment that were working on any of the taxiway segments at the time of the attack.

- Determination of which aircraft have been damaged and the losses to personnel and equipment being used for on-equipment tasks at the time of the attack.
- Determination of damage and casualties, both conventional and chemical, at all of the designated facilities.
- Determination of the priority for repairing damaged shelters by ranking them starting with the least damaged shelter.
- Determination of losses to equipment, munitions, spare parts, TRAP, building materials, and POL.

The procedures and assumptions used for these determinations are discussed in the next several sections.

5. DATA PROCESSING FOR ATTACKS

The TSARINA-generated CT40 cards are read by TSAR during initialization. The scheduled attack times are placed in the ATTACK array heap, the CT40 damage data are stored in the DAMAGE array, and the other data transferred on the CT40 cards are stored in the MPPERS, ARC, SHEL1, SHELOS, RAMPS, UXODTA, EXPLOD, and DUFFAC arrays.

At the time of an attack, subroutine MANAGE calls subroutine BOMB where several temporary arrays must first be zeroed and the preattack numbers of on-base personnel stored temporarily; the attack characteristics then are determined. If the attack includes chemical weapons, the effective warning time is determined, and if it is the first attack to involve chemical weapons, subroutine CWHITS is called to initialize the chemical hit data that are stored on disk Unit 18. The next steps for chemical attacks are to update the residual preattack dosage estimates at each monitoring point with a call to CWDOSE and to estimate the percent casualties for each personnel type associated with one or another monitoring point in the TSARINA data base. These estimates are developed in subroutine CWLOSS and take into account the MOPP that the individuals are wearing at the time of the attack and the warning time; the casualties and the portions that would be hospitalized due to the chemical attack are stored for all personnel types whose locations were specified by TSARINA.

When these several preparatory calculations are completed (or skipped for conventional attacks), subroutine BOMB unpacks whatever conventional attack damage results are stored in the DAMAGE array. Estimates of personnel casualty rates due to conventional effects are added to whatever chemical casualty rates were estimated by calling subroutine TRIAGE, where it is assumed that the causes for the casualties are independent events; if the special CT39/99 cards are used to specify that additional percentage casualties are to be assumed among off-duty personnel from unspecified causes, those losses are also included. Furthermore, if the user has availed himself of the convenience provided in TSARINA of describing an "all other" category of personnel and has also chosen to assume that these personnel are distributed other than uniformly by type in the various locations (i.e., has used the option offered with C17/9), the nominal casualty percentage is adjusted for each personnel type.

When the percentages of the personnel that are to be casualties have been stored, the numbers of aircrews and the numbers of off-duty support personnel that are affected are determined using calls to DISABL (for aircrews) and ASSAY (for others). On-duty support personnel are checked later, and for those personnel types whose locations were specified for TSARINA, losses are based on the percentage casualty data stored in the temporary SHOPEO array (except for those engaged in on-equipment tasks). Whenever casualties are sustained, personnel may be required to provide "buddy care" for a specified period of time. Other members of a work team are used if they are not themselves casualties, or if they are, other members from the same squadron or wing organization are used; off-duty personnel are assigned to assist off-duty casualties.

Subroutine BOMB also applies whatever loss rates were generated in TSARINA for support equipment, aircraft spare parts, munitions, TRAP, building material, and POL. Unassigned equipment losses are assessed immediately, whereas equipment that is assigned to a task is handled later. As with personnel, the user may use TSAR's default location assumptions for equipment as well as for personnel or may specify their locations in the TSARINA input on a type-by-type basis. An additional option can be used for equipment; when the user has set ONLYUE to 1 on CT2/1, the TSARINA generated loss rates for equipment are applied only to the unassigned equipment, and the TSAR default location assumptions control the losses to equipment that is assigned to a task. With this option, all facilities at which TSAR may assume the equipments are assigned will themselves need to be entered into the TSARINA data base, so that the loss estimates for equipment at that facility can be computed and transferred by TSARINA.

Locations *must* be specified in TSARINA for serviceable spares, munitions, TRAP, building, materials, and POL if these types of resources are to sustain losses. As with personnel and equipment, these resources may be treated type by type, or many types may be treated as a group using the "all other types" feature for specifying their storage locations in TSARINA. The option for assigning variable loss rates to each of the "all other" types that is outlined for Card Type #17/9 may be used with any of these classes of resources and is taken into account in subroutine BOMB as these resources are decremented for losses.

For POL both the fuel storage capacity and the residual amount of fuel are decremented by the percentage loss specified by TSARINA, on the assumption that the fuel would be distributed in proportion to the capacity of the various storage tanks. Whenever TSARINA results include POL losses, the POL tankage feature in TSARINA should be activated by initializing the entry in the fourth field on the REDO Card, so that the POL losses transferred to TSAR for a sequence of attacks will be based only on that percentage of the original tankage that was still intact at the time of the attack.

The last data to be acted on in BOMB, before the aircraft are treated, are those for taxiways, aircraft parking ramps, and facilities other than aircraft shelters. TSARINA generates data for each of these entities that are defined in the TSARINA data base and transmits a record of any physical damage or chemical contamination to TSAR (for facilities, data are passed only for those targets that have had their TSAR "facility" number entered on the TSARINA TGT Card). The physical damage and estimates of the combined conventional and chemical casualties among personnel and equipment at each particular facility at attack time are computed and are stored in the temporary FACDAM array; the amount of chemical contamination, the numbers of UXO and mines on each taxiway, and the numbers of craters that must be repaired for aircraft to transit the taxiway are also recorded; and finally, the expected fraction of aircraft that could be damaged and the chemical contamination are stored for each aircraft parking ramp. If detonation delay times have been specified for UXO, a specific time is selected for each UXO with a call to subroutine BANG, and the relevant information is stored in the EXPLOD heap.

The last TSARINA data to be extracted from the DAMAGE array are those that define damage to on-base aircraft and to aircraft shelters. When these data are encountered subroutine ATTKAC is called, where the level of damage to each aircraft

shelter, the loss rates for personnel, equipment, and spares in each shelter with closed or open doors, and the chemical deposition data for the individual shelters are stored in the SHEL, SHELOS, BSHELT, and CWSHEL arrays. If shelter repairs are ongoing at attack time, and the shelter is redamaged, losses to the personnel and equipment are assessed, and the total shelter damage is computed.

A check then is begun for each aircraft assigned to the base. If an aircraft is in flight, it is skipped; if it is on base its location is determined to be either (1) in a particular aircraft shelter, (2) on a particular parking ramp, (3) refueling at a hot pit located on a specific arc, (4) being decontaminated on a taxiway, or (5) taxiing. TSARINA provides data for estimating damage that is specific to each location for the first three items and a single estimate for aircraft in the open or "aircraft taxiing," which is the average of the damage expectancy for an aircraft located at random on the taxiway network.

For those aircraft that are not taxiing, a record is available of all ongoing work, and the maintenance personnel and equipment that are at risk with each aircraft are known; and, because the nature of the activities are known, a determination can be made as to whether an aircraft shelter door would have been left open on the basis of the user's specification of that likelihood on CT18/1. The condition of the door and the TSARINA data permit estimation of the damage to each aircraft and the casualties and losses among the assigned personnel and equipment; both conventional and chemical effects are taken into account in judging casualties. On-equipment tasks are the only instance in which the TSAR default location assumptions are imposed automatically, and any location specifications in TSARINA for these personnel and equipment are ignored. If the aircraft is damaged but repairable, the repair tasks are determined. If the aircraft is beyond repair but parts may be salvaged, a random number is drawn for each part on the aircraft's parts list and compared with the product $FSALVG \times$ "the parts survival rate"; those that may be salvaged are placed in stock. Flight assignments and preflight tasks for damaged aircraft are canceled, and the record of tasks required to ready the aircraft for flight is reorganized. Records of destroyed aircraft are purged from the data base with a call to subroutine ENDAC.

When all the data have been entered that had been stored in the DAMAGE array for an attack, subroutine BOMB calls subroutine DOSURF where several specialized computational activities associated with the base's runways and the taxiway network are handled.

Effects of Damage to Taxiways and Runways

TSARINA provides TSAR with the numbers and locations of bomb craters and the numbers of mines and UXOs that are on or sufficiently close to runways or taxiways to cause damage, or threaten aircraft using the runways or taxiways. For craters on runways, the runway number, the crater coordinates relative to the runway, the crater radius, and the attack number for each crater are kept in a single table. As the craters are repaired they are purged from the table, but unrepaired or partially repaired craters remain to be combined with those from later attacks. For the remaining damage (UXO and mines on runways and craters, mines, and UXOs on taxiways), the runways and taxiways are divided into sections and only the total numbers that are still unrepaired in each damage category in each section are maintained, not their precise coordinates.

After an attack, RWYTAX tests to see if operations are possible from an MOS that is of sufficient size to provide emergency flight operations. To do this, the runways are searched (in RUNWAY) to find if there is a clear area of the prescribed minimum size. The MOS may be restricted such that it is parallel to the sides of the runway, or a location that is skewed may be sought (see CT177). This area may be either rectangular or rectangular with a superimposed triangular area needed for cable clearance with a mobile aircraft arresting barrier. If no such clear area can be found, the area that would require the minimum amount of work to clear is located and designated as the runway strip to be cleared, or MOS.

The user may select one of two alternative metrics to determine the runway MOS to be cleared: the MOS with the smallest number of craters to be repaired, with ties settled by choosing the MOS with the smallest number of manhours required to clear mines and UXOs; or the MOS with the smallest total number of manhours needed to repair all craters and clear all mines and UXOs. The former is used when TSKRWY is set to 0 and the latter when it is set to 1. If both manual pickup and sweeping procedures are available for mine clearance, the procedure actually used to determine the MOS is the one requiring the smaller number of manhours to clear all the mines from the MOS.

Clearance of the designated MOS has priority over other base repairs, and aircraft operations are prohibited until the MOS is cleared. It is assumed that UXOs must be cleared first, then the mines are cleared, and finally craters are repaired. The runways are divided into sections (see the discussion of the runway/taxiway network below) and the repair ordering holds for each section of the designated MOS.

The procedures for removing a UXO depend on the munition type, and the procedures for repairing a runway crater depend upon the crater size; each of these procedures and those for repairing a taxiway crater may involve a sequence of steps, each with a time and its own requirements for personnel, equipment, and material. At night, specific task times may be increased to account for the difficulties of working under nighttime conditions. For crater repairs, a succeeding step may be permitted to proceed at the same time as the preceding step. These requirements are controlled by the user's specifications in CT37/66, CT37/77, CT37/88, and CT37/99. When data have been provided from TSARINA to define the detonation time of an UXO, and a UXO "detonates," the entry point DOBANG in BANG is called. One of the UXO remaining on the appropriate taxiway segment is picked at random, and if any personnel are working on it at that time, a specified fraction are casualties; and a specified fraction of those are killed. If personnel are working on any other task on the same taxiway segment, another specified percentage are assessed as casualties; if personnel are working on adjacent taxiway segments (i.e., segments that share a node), another percentage are casualties. When any of these repair teams sustain casualties, the task is stopped and the injured personnel are placed in the "hospital." The UXO is then removed from the EXPLOD heap.

The repair procedure used to clear mines on a runway section depends upon the types of resources available (personnel, equipment, and materiel) and the magnitude of the job. If both manual and sweeping procedures are specified (see CT37/99), the procedure requiring the smaller number of manhours to clear the section is used if resources are available. If resources are not available (either because the resources are being used for clearing other runway sections or have been damaged by attacks), any alternative procedure specified for the given procedure will be used if resources are available. The alternative procedure must either be of the same type as the primary procedure—sweeping if sweeping, manual if manual) or be the primary procedure of the opposite type. If resources for the primary or alternative procedures are not available, then any alternative specified for the alternative procedure is used if resources are available, etc.

If only sweeping is available, the primary sweeping procedure is used if resources are available. If resources are not available, any alternative procedure specified will be used when the resources are available. If only manual procedures are specified on

CT37/99, then the missions are cleared with manual procedure if resources are available. If resources are not available for the primary manual procedure, any alternative procedure is used when resources are available.

After an attack the system of taxiways that connects the aircraft shelters and parking ramps to the runways may be so damaged that some of the aircraft cannot reach the designated MOS. These aircraft are then unavailable to fly missions until paths are cleared to that MOS.

It is not usually necessary to clear all of the taxiway sections to provide access to the MOS for all of the aircraft. For the purpose of determining which sections need to be repaired and the preferred order to repair them, the system of runways and taxiways is treated in TSAR as a single combined network. The arcs of the network are sections of taxiways or runways. The nodes of the network are the locations where the sections meet. TSARINA reports the damage (numbers of craters, mines, and UXOs) to each arc.

In TSAR the aircraft shelters and parking ramps are associated with the nodes at which they are located. Aircraft are assigned to specific shelters and parking ramps, and repair times (for clearing mines and UXOs and repairing craters) are determined for each arc. As part of the algorithm described below for determining the preferred repair schedule for clearing the arcs (implemented in TAXIWY), the minimum-repair-time path to the designated runway MOS is determined (in PATH) for each occupied shelter and parking area. Those with minimum-repair-time paths equal to zero are able to reach the designated MOS without repairs.

A heuristic rule was developed for determining near-optimal taxiway repair schedules for TSAR, using as a criterion the minimum average time that aircraft do not have access to the designated MOS. The steps of the heuristic rule are:

1. For each node k still without access to the MOS, perform the following calculations:
 - a. By using a shortest path algorithm, determine a minimum-repair-time path from the given node to the MOS.
 - b. Set T_k equal to the total remaining repair time for the arcs on the path of step 1.a and set A_k equal to the total number of aircraft at nodes on this path.
2. Find K , the value of k that maximizes A_k/T_k .
3. Working outward from the nodes at the ends of the MOS to node K along the minimum-repair-time path to node K , find the first arc on the path that has not yet been repaired. This is the next arc to be repaired.

The repair schedules selected by this heuristic rule compared very favorably (an average increase in repair time of less than 1 percent) with the optimal repair schedules in a test involving 100 sample problems on a typical NATO airbase.[4]

For taxiway sections, we assume that UXOs are cleared first, then mines are cleared, and finally craters are repaired. Clearing the MOS has priority on repair resources over taxiway repairs, so that the MOS will generally be cleared first. Then, as the arcs are repaired over time, more and more aircraft will have access to the runway and be able to fly missions.

When the MOS has been cleared to permit flight operations and a sufficient number of taxiway segments have been cleared to permit all shelters and ramps to access the MOS, the user may specify that additional runway clearances be made by specifying a nonzero value for the runway repair mode (RRM) on the CT177 card.² The permissible values of RRM and the runway clearance actions taken are:

- 0 Stop runway clearance after the MOS has been cleared.
- 1 When the designated MOS and taxiway sections have been cleared, continue clearance of the MOS runway by first extending its length to X, then extending its width to Y (X and Y are dimensions input on the CT177 cards for the Extended MOS). Finally, if "Limited Extension" on CT177 is not initialized, the entire runway containing the MOS is closed.
- 1 When the designated MOS and taxiway sections have been cleared and the MOS is not on the main runway, clear an MOS on the main runway and sufficient taxiway sections to provide access to the new MOS for all shelters and ramps. Then extend the MOS on the main runway as when RRM = -1.
- 2 Clear the main runway as when RRM = 1, after lengthening the original MOS to X.
- 3 Clear the main runway as when RRM = 1, after lengthening the original MOS to X and then widening it to Y.
- 4 Clear the main runway as when RRM = 1, after lengthening the original MOS to X, then widening it to Y, and finally (when "Limited Extension" is not initialized) clearing the entire runway containing the original MOS.

The MOS on the main runway and the additional taxiway repair schedule to provide access to this MOS are selected by the same criteria used to select the original MOS and taxiway repair schedule.

²If the MOS is so short that a mobile aircraft arresting system (MAAS) is necessary for recovering aircraft until the clear surface is extended further, the substantially increased intervals required between recovery aircraft can also be simulated (see the discussion for CT177 in Vol. II).

If, when extending the MOS on a runway other than the main runway, less work would be required to clear the same size surface on the main runway, the clearance task is switched to the main runway, and all further MOS extensions take place on the main runway.

When the runway clearance tasks determined by the value of RRM have been completed, civil engineering repairs on the runways stop until a subsequent attack requires further work.

Other Assessments at Attack Time

When the damage to the base's horizontal surfaces has been updated, the location of an MOS selected, and the preferred taxiway repair schedule determined, control is returned to subroutine BOMB. BOMB then calls subroutine REORGN to begin a series of other computational activities.

The first step taken in REORGN is to calculate for each element of each distributed function the fraction of the surviving capability that is associated with each such element. These data are used in subsequent calculations to determine the losses of certain sets of resources. The first set of resources considered are off-duty support personnel whose loss rates were not specified by the TSARINA locations. As described in Sec. VIII.2, the default location of off-duty personnel is Facility #50 and any other facilities the user has joined to this facility as members of a distributed set. The numbers of each type of personnel, the proportions at each of the facilities, the chemical protection characteristics at each of the facilities, and the damage and chemical contamination at the facilities are used to estimate the casualties and fatalities at each facility. Personnel who are to be hospitalized from conventional or chemical effects are placed temporarily in the SICK array, and the numbers of personnel that are casualties are withdrawn from the on-base counts of healthy personnel in the PEOPLE array.

The casualties to on-duty and off-duty aircrews are evaluated next, and then those for unassigned on-duty ground personnel. Each of these categories of personnel has a nominal facility that may be a distributed set of facilities, each with its own relative capacity and chemical protection characteristics.

Unassigned equipments whose locations are not specified in the TSARINA data are treated next. As with personnel, equipments may be associated with a squadron, with a backshop, with munitions assembly, or with the civil engineers; each of these groups has a TSAR default facility (Table 4), which may be the first of a distributed set. The

damage estimates for "equipment," passed from TSARINA for each of these facilities is applied to determine equipment losses.

Parts and equipment repairs going on at the time of the attack are evaluated for losses next. Because a facility is selected when each such task begins, the personnel and equipment loss data for the specific facilities can be applied to estimate the losses. The loss fraction for "parts" at each facility is used to determine if the part being repaired was itself destroyed. If losses are sustained, the task is canceled; if losses are not sustained, the job is interrupted and surviving personnel are taken off the job and sent to GOREST by STOPIT when USECW > 0; if USECW = 0, the length of the job is extended to simulate the specified postattack delay. Repairable parts and equipment in the administrative delay are checked next; a facility is selected for each, based on the capacities of the undamaged shops, and their destruction is determined from the parts loss rate (or equipment loss rate) estimated for that facility (the "parent" shop is assumed for those part types that must be repaired in that shop—see CT35/2).

Parts and equipment repairs that have been waiting or were interrupted sometime before the attack are checked next; their losses are based on logic paralleling that for items in an administrative delay.

Interrupted and ongoing munitions assembly jobs are handled next; task facility locations are assigned when the task is started, so the losses among the associated personnel and equipment are derived from the appropriate damage data for those facilities. For munitions jobs the "damage to parts" datum from TSARINA is applied to estimate destruction of the munitions; the mean area of effectiveness (MAE) information entered into TSARINA for parts obviously must reflect this usage of the TSARINA "damage to parts" output.

When the effects of chemical warfare are not being simulated (i.e., when USECW = 0), the projected completion time for ongoing on-equipment tasks (other than preflight tasks) and munitions assembly tasks are increased to account for the postattack delay discussed earlier. This approximates the need for personnel to take care of emergency postattack jobs for this length of time, before they can return to their work. When USECW > 0, all tasks are interrupted and personnel disposed of with the STOPIT subroutine; jobs are then restarted after the appropriate delay using whatever personnel are then available.

Civil engineering repairs on facilities other than runways and taxiways are treated next. The first step is to interrupt each ongoing job and estimate how much of the facility repair task remained at the time of the attack. Losses among personnel and equipment in use on each job are estimated as for most other types of work: The loss rate associated with the work place is applied, unless personnel and equipment loss rates were passed from TSARINA for the particular types of personnel and equipment involved on the job. It is assumed that building materials are consumed at a uniform rate during construction and the unused materials are returned to stock when a task is interrupted for the attack. If the building has received additional damage, some unused building materials are also lost; the fraction of the unused materials lost is assumed equal to the square root of the new damage fraction. After decrementing the assigned personnel, equipment, and material for any losses, the job is interrupted and the residual personnel are either added to those available (when USECW = 0) or disposed of by STOPIT, and the equipments are returned to stock. New damage to the facility is added to what remained unrepaired from previous attacks assuming that the damage from past and present attacks are independent; that is, the total damage is taken to be $D = 1 - (1 - D1) \times (1 - D2)$ where D1 and D2 are the damage fractions from the previous and latest attack, respectively.

All on-base activities have now been checked and all losses have been established. A call is next made to ADDAGE to check the distribution of the surviving equipments among the wing and squadron organizations; equipments are reassigned so that the numbers possessed by each organization are in the same proportions as the "target" levels for each organization. When USECW = 0, the same action is carried out for personnel with a call to subroutine REDPEO, but when USECW > 0, personnel are only redistributed among the several organizations at the time the shifts change. (Too much processing would be entailed if organizations were reorganized every time personnel numbers rose and fell and it would not be particularly realistic.) A record is maintained in the eighth element of the PEOPLE array to indicate that reorganization will be required when shifts are changed. In either case, personnel that require hospitalization are handled with a call to subroutine CLINIC (if RECUP ≠ 0), and replacements for fatalities are ordered if specified with CT33.


Only a few steps remain to complete the attack analysis. Summary damage data are filed and then three pseudo-civil-engineering tasks are scheduled for completion when (1) the runway can first be used, (2) runway and taxiway repair work can be

restarted, and (3) other on-base tasks may be restarted. The user can introduce these delays, as explained in Vol. II in connection with CT17/9, to account for the various emergency problems that would need to be treated following an airbase attack but have not been simulated, such as fires, emergency utility, and road repair. When these "tasks" have been stored in the CEJOBQ array and control has been returned to BOMB, the dosages at all monitoring points are updated and the MOFPs appropriate for all facilities are determined. Finally, the air traffic control performance factors are updated in light of the current damage at the base.

6. POSTATTACK RECOVERY AND RECONSTRUCTION OF BUILDINGS AND OTHER FACILITIES

The base engineers and other civil engineering resources also are called upon to carry out other emergency repairs essential for restoring critical functions. The procedures and options for handling these repairs are different from those explained above for runways and taxiways.

The options the user has available for representing on-base facilities other than for runways and taxiways will be explained in the context of the following sketch:



The shaded blocks are intended to represent actual facilities and all blocks represent repair procedures. Thus, "A" depicts the situation in which the activities of a given shop are carried out in a single location, and damage to that facility can be repaired using

either a basic repair procedure or—the backup block—an alternative procedure. The nature of the basic procedure is defined with the FACLT Y array data for this shop, and the alternative procedures are defined by entries in the CERQTS (civil engineering requirements) array. The "B" depicts another shop whose activities also are carried in a single location, but three sequential civil engineering procedures are required to restore shop operations. Data defining each of these steps occupies a column in the FACLT Y array.

Situation "C" is a distributed shop whose activities are carried out in three distinct locations, each of different size and capacity; the main location requires a three-step process to restore operations, whereas the auxiliary locations require only two steps. Each of these locations and each of the subsequent procedures occupies a column of the FACLT Y array.

The time for the repair or restoration of these on-base facilities is related to the amount of damage, the type of structure involved, the numbers of civil engineering personnel and equipment that can be brought to bear on the job, and whether the work must be done at night. Each facility is described with the C137 entry by a facility number, a size, and the nature of its construction (or more correctly, the nature of the reconstruction or reconstitution required). The "facility" number is identical to the location of these descriptive data in the FACLT Y array. The first 50 numbers are reserved for the applications noted in Table 4; other locations in the array are to be used for data that describe alternative locations of distributed shops and subsequent steps in a repair process. Thus, when more than one civil engineering repair procedure is required to restore a particular facility to operational status, the descriptive data for the subsequent phases of the process are to be stored in otherwise unused columns of the FACLT Y array. When the facility sustains damage, all steps of such a repair sequence will be assumed to sustain the same percentage damage.

The resource requirements for the procedures used in repairing facilities of the differing types are filed in the CERQTS array. For each procedure some number of each of two types of civil engineering personnel and equipment may be specified. The quantities specified in the basic procedure entered for each type of structure should represent the largest-sized force that can reasonably be put to work on that type of job. For most types of jobs, alternative procedures should also be included (and defined in the CERQTS array) so that they may be adopted when insufficient resources are available for the basic procedure.

The overall magnitude of a repair task is given by the product of the "size" of the building and "percent damage." The time to carry out the repair and the quantities of (up to) two types of building materials that are required for the repair procedure are specified in terms of the requirement for one "unit" of reconstruction, where the "unit" is defined with the same metric the user chose in specifying the "size" of the facilities of that type. For materials, the quantities required for the repair are simply the amount for "one" unit of repair times the overall magnitude of the job. For example, if 10 items of material type #1 are required for a unit of reconstruction, and a building of "size" 80 sustains 50 percent damage, 400 items of material would be required for the repair.

The user is given greater flexibility to specify reconstruction time, however, because of the possibilities for nonlinear relations between repair time and the magnitude of damage. In general, the time to repair a facility is expressed as:

$$\text{Constant} + t \times (\text{Damage Percent} \times \text{Facility Size})^b$$

or

$$T = \text{Constant}(B) + t \times N^b \quad (1)$$

where we define t as the time for one unit of construction and N as the number of units of reconstruction that are required.

On the CT37 card, appreciation of the variation of repair time with damage level is specified with the compact code number C , where C is defined as:

$$C = (B - 1) + 12 \times P \quad (2)$$

By choosing C , the user specifies one of the 12 choices for $\text{Constant}(B)$ ranging from 0 to 48 hours

$$0, 1, 2, 3, 4, 6, 8, 12, 18, 24, 36, 48 \text{ hours} \quad (3)$$

and one of seven choices for the exponent "b," where $b = g(P)$ and

$$g(P) = 0.5, 0.75, 0.9, 1.0, 1.1, 1.25, \text{ and } 1.5 \quad (4)$$

for P from 1 to 7.

The code, C , that designates the components of the repair time in functional form, is interpreted in FTIME as:

$$P = C/12 \text{ the largest integer multiple of 12 in } C, \quad (5)$$

where $1 \leq P \leq 7$ and

$$B = C - 12 \times P + 1, \text{ where } 1 \leq B \leq 12. \quad (6)$$

If, for example,

$$C = 48$$

then from (5) (6)

$$P = 4$$

$$B = 1$$

and from (3) the constant is zero hours, and from (4) the exponent is 1.0, or $T = tN$. Thus, for this example the facility repair time equals the number of construction units times the repair time per unit of construction, that is, a linear relationship between damage and repair time. The number of units of reconstruction required, N , is the product of the TSARENA-generated damage fraction for the specific facility and the facility "SIZE" defined on the CT37 card (for aircraft shelters $SIZE \geq 25$). The user's choice among the seven possible values for the exponent b specifies the shape of the resulting function as illustrated in Fig 9. Likewise, choosing the time for a "unit of repair," t , and choosing among the 12 values for B permits the selection of a repair time distribution appropriate to the user's needs. Table 5 provides illustrations of how facility repair time can be specified to vary with the amount of damage to a facility of size = 100 and for unit repair time, $t = 60$ minutes. Since the time required for certain civil engineering tasks is greater at night, the user may stipulate this additional time by noting the task time at night as a percent of the task time during daylight in the task description on CT38.

When subroutine REBUILD is called after the appropriate postattack delay, all damaged aircraft shelters are ordered from least damaged to most damaged in array DAMSHL, so that the least damaged get priority when shelter repairs are considered. Then, all damaged facilities are checked in the order of priority the user provides for each base. If the civil engineering personnel and equipment that are available are sufficient to initiate repairs of all damaged facilities on the priority list up to the facility designated as the base's CRBLDG, subroutine INICON (initiate construction) is called




Fig. 9—Variations of facility repair time

for each damaged facility to allocate the personnel and equipment and to withdraw sufficient building materials from stock to complete each job. The job completion times are placed in the heap in the CEJOBQ array. This process continues until all damage is under repair or the resources are exhausted.

If the civil engineering resources are insufficient to start all the critical tasks, the allocation starts with the highest priority facility that is damaged and proceeds until resources are exhausted as was just described, except that the first alternative repair procedure is used when one has been identified. Aircraft shelters are treated when facility #45 is encountered in the CEPRTY array (CT39).

When resources are exhausted, and when a central repair facility (CIRF) has been identified, one final task remains. For each shop that was damaged in the attack, a rough check is made to see if the parts could be shipped to and from the CIRF in less time than

it is projected to repair the shop; if so, the faulty parts are shipped. This last task is carried out in the subroutine SHCIRF (ship to the CIRF).

7. COMPLETION OR INTERRUPTION OF CIVIL ENGINEERING REPAIRS

When a civil engineering repair task has been completed or must be interrupted because the task personnel must rest, control is transferred to subroutine ENDCE directly when USECW = 0, or, when USECW > 0, by subroutine STOPIT, where the condition of personnel is checked using GOREST. Those personnel determined to need rest by GOREST are assigned to their rest location for the appropriate rest period. Those personnel determined to be casualties during the work period either to heat prostration or chemical effects are hospitalized if sick, or removed if killed.

For structural repair tasks, the equipment and any remaining personnel are released from the task and a call is made to FIXSUR to see whether resources are sufficient to start any remaining runway or taxiway repair tasks (runway and taxiway repairs have priority over other civil engineering repair work). Then, if the structural repair task has not been completed but resources are available to continue the work, the

Table 5

EXAMPLE REPAIR TIMES FOR A FACILITY WITH SIZE = 100

Hours

Code Number	Code Components				Percent Damage					
	B	P	Delay Hrs	b	1	10	25	50	75	100
12	1	1	0	0.5	1	3.2	5	7.1	8.7	10
48	1	4	0	1.0	1	10	25	50	75	100
84	1	7	0	1.5	1	31.6	125	354	650	1045
16	5	1	4	0.5	5	7.2	9	11.1	12.7	14
52	5	4	4	1.0	5	14	29	54	79	104
88	5	7	4	1.5	5	35.6	129	358	654	1049
21	10	1	24	0.5	25	27.2	29	31.1	32.7	34
57	10	4	24	1.0	25	34	49	74	99	124
93	10	7	24	1.5	25	55.6	149	378	674	1069

(t = Unit Time = 60 min)

Repair Time = Delay(B) + t × (Percentage × SIZE)^b

C = 12 × P + (B-1)

task is continued. If the repair task has been completed but subsequent repair tasks are needed to complete the repair, the next task is started if resources are available. If no additional work is required, the facility repair status is updated and INICON is called to start any remaining structural repair tasks; when the repaired facility is a maintenance shop, subroutine CHECK is called to initiate the various activities that were held up because of the damaged facility; when the repaired facility is an aircraft shelter, it is entered in the available shelter queue.

If the task being completed or interrupted is a runway or taxiway repair, the number of tasks remaining on the runway or taxiway segment is updated (data on the number of mines and the percentage completion for each UXO and crater repair on each taxiway and runway segment is stored to account for tasks that have only been partially completed when they are interrupted). If the task being completed is a runway repair and no runway repair tasks remain, then the runway closed indicator is set to zero (aircraft that have access to the runway can then fly missions). If the task being completed is a taxiway repair and no repair tasks remain for that taxiway segment, the runway nonaccess indicator is set to zero for any nodes of the taxiway network that have access to the runway because this taxiway was cleared (aircraft in shelters or on parking ramps associated with these nodes then have access to the runway).

Whether the taxiway or runway repair task has been interrupted or completed, the personnel and equipment are released in ENDCE, and if any runway or taxiway repairs remain, FIXSUR is called to try to start the remaining work. Then INICON is called to try to start any remaining structural repair jobs with the remaining resources.

IX. AIRBASE OPERATIONS IN A CHEMICAL ENVIRONMENT

1. INTRODUCTION

There is widespread concern that airbases may be attacked with chemical and biological weapons, as well as with conventional weapons. Many initiatives have been taken to limit the effect that such attacks might have on an airbase's sortie generation capability and many more are under consideration. Chemicals delivered to an airbase with surface-to-surface missiles like the Soviet SCUD, or by aircraft with bombs or dispensed from aerosol spray tanks, or (particularly for BW) by covert agents can be expected to very seriously disrupt, if not terminate, sortie generation at an ill-prepared airbase.

Special clothing and related equipment are needed to protect personnel who must move about in such an environment, and specially prepared facilities are needed if personnel are to work and live in a contaminated area without individual protection. Chemicals of the kinds that might be used could remain a threat for several days or, with highly volatile options, could blow away in minutes. Most of the chemicals of concern are highly toxic even in minute quantities and great care must be exercised to avoid liquid droplets or (in some cases) vapor from penetrating into an individual's system through the skin (percutaneous poisoning), and to avoid inhaling the highly toxic vapors.

For persons working in the open, heavy suits with poor heat transfer properties and limited permeability are currently worn to limit the percutaneous hazards, and awkward masks are worn to limit the inhalation hazards. Personnel attempting to carry out the many kinds of tasks required to prepare an aircraft for an effective sortie engage in many strenuous tasks and frequently have to perform tasks that require considerable dexterity. The cumbersome ensembles that must be worn to protect these personnel from chemical hazards inhibit sortie generation because mobility and dexterity are limited and interfere with the worker's vision and ability to communicate effectively with fellow workers. These constraints can have several effects, but the primary measure of these MVDC (mobility, visibility, dexterity, communications) constraints is to lengthen the time that is required to successfully complete a task. Work under these conditions can also lead to a rapid buildup of body temperature, excessive perspiration, and possible dehydration, so that jobs will have to be interrupted before they can be completed to

permit work crews to recuperate. In other cases, the tasks will be completed but the individuals will be severely stressed and will also need a special rest period to recuperate before they can commence another task.

Recuperation can occur by resting at one's work place, by going to an on-base location with conditions more favorable for recuperation, or by going to an off-base location where, possibly, the effects of chemicals are absent. The necessary recuperation will depend on many factors, and a simulation should be able to represent a range of possibilities. Some bases may have many on-base collective protection shelters (CPSs) with a large capacity that permit personnel to enter and leave quickly, but the CPSs on other bases may have a very limited capacity or require a long time for processing personnel that are entering. In the latter instance many of the personnel may either need to go off base or remain at their work place under environmental conditions that may slow recuperation.

In some circumstances personnel may work until they collapse from heat prostration or from excessive dehydration (or other problems) and may need to be hospitalized; other personnel may need to be hospitalized because of the toxic effects of the chemicals, and some may die from such effects. The likelihood of these serious effects will depend in a complicated way on the quality of the protective ensembles, the warning time available to don the protective ensembles, and the ambient environment in which they are worn.

Features have been included in TSARINA and TSAR that provide mechanisms for simulating many of these various complications; these features are activated in TSAR by initializing the control variable USECW (CT3/4). The user may simulate a sequence of conventional and chemical attacks on one or several airbases and compute the surface deposition and vapor concentration at all points of interest at the time of an attack, and account for the effects of agent evaporation and vapor transport after the attack. Coupled with user-prescribed rules regarding the level of personal protection required for various agent concentration levels and data describing the heat retention properties of each of three different personnel ensembles, the work time and required rest time as well as casualties and fatalities are computed for the various tasks under the various environmental conditions.

2. COMPUTATION OF ALLOWABLE WORK TIME AND REQUIRED REST TIME

Body Temperature

The predictive equations for body temperature response to environmental factors that were developed by Dr. R. F. Goldman et al., in the Military Ergonomics Laboratory of the U.S. Army Research Institute of Environmental Medicine, Natick, Massachusetts, are summarized in [11], where the combined effects of metabolic and heat stress on a human's rectal temperature is evaluated as

$$T_F = 36.75 + T_M + T_{RC} + T_E + T_A$$

where

T_F = final equilibrium rectal temperature

36.75°C = basal metabolic rate

T_M = added temperature due metabolic load

T_{RC} = temperature change due radiation and convection

T_E = added temperature due evaporation

T_A = added temperature due a lack of full acclimatization

and the variation of rectal temperature with time was formulated as:

$$T = T_I + (T_F - T_I \{ 1 - \exp(\tau(td - t)) \})$$

where

T = the temperature at time t

T_I = the initial temperature

T_F = final equilibrium temperature

τ = the time constant°C/hr

$$= 0.5 + 1.5 \exp(-0.3(T_F - T_I))$$

t_d = initial temperature time lag

= $58/M$ (M = metabolic heat rate), and

t = duration of work, hours

These equations, and the equations presented in [11] for T_M , T_{RC} , T_E , and T_A , have been implemented in subroutine CWTIME and CKTEMP.

Excessive Dehydration and Sweating

A different group, the Working Group on Thermal Environments of Subcommittee #5, "Ergonomics of the Physical Environment," of the International Standards Organization, proposed new standards for work in hot environments [12]. That paper suggests analytic procedures for computing limits on the permissible working time when (1) the required evaporative heat loss is not achievable, (2) the required cutaneous wetting is excessive, or (3) the required sweating loss involves an excessive dehydration. Although specifically noted in [12] as not yet being applicable to the case in which personnel are wearing special protection clothing, these procedures have been incorporated into TSAR (in subroutine DEHYDR) on the basis of AF/AMRL's suggestion that they represent the best procedure currently available for assessing these effects. The constraints implied by these three criteria under the existing environmental working conditions are each checked in subroutine DEHYDR and the shortest working time and longest required rest times are estimated; these are subsequently compared with those derived from the Goldman formulation of the rates at which the rectal temperature will rise and fall.

The examples in [12] and the results of the simulation carried out to date with TSAR suggest that the wetting and dehydration criteria will be of importance primarily at extreme ambient temperatures as might be expected in Southwest Asia but will seldom be important in NATO's Central Region. Therefore, the user is provided a switch to deactivate the extra processing that will not be relevant in many situations; the name for the switch, NOVOGT, (C:3/5), was chosen to reflect the fact that Dr. J. J. Vogt of the Centre d'Etudes Bioclimatiques du C.N.R.S., in Strasbourg Cedex, France, is credited with being the primary contributor to the formulation incorporated in subroutine DEHYDR.

Since the codes that check the limitations imposed by the excessive sweating and dehydration criteria and by the temperature rise criterion are each confined to a single subroutine, it will probably be quite easy to replace these analytic expressions with improved representations if they become available.

3. APPLICATION OF WORK TIME CRITERIA TO TSAR TASKS

Whenever a task is to be started in TSAR, a call is made to subroutine TTIME with the mean time and the time distribution that is specified for the task; a sample task time is chosen (if there is a time distribution), and when the control variable USECW is 0, control is returned to the calling routine. But when USECW > 0, subroutine CWTIME is called by TTIME with the "heat factor" for the job, the job location, and the total task time. Based on the type of chemical ensemble used at the base and the MOPP that is appropriate for that work location under the prevailing chemical conditions, the task time degradation corresponding to that task's MVDC factor is applied to obtain the total work time.

A call is then made to subroutine CKTEMP to determine how long a crew of average personnel could work before their rectal temperature rose too high. Since the ambient temperature, humidity, and wind velocity are stipulated for each two-hour period, CWTIME actually makes a sequence of calls to CKTEMP for each time period (a step-wise integration) until the task must be stopped because of excessive temperature, or until the temperature has been determined for the end of the job. When the user also wishes to check the limits that may be imposed by excessive cutaneous wetting or excessive dehydration, subroutine DEHYDR is called when subroutine CKTEMP is first entered, where the length of time that the crew could work and the length of time they would need to rest is computed in light of these criteria. Then, when the permissible work time based on temperature has been determined by CKTEMP and CWTIME, that time is compared with the time limit implied by the wetting and dehydration criteria, and the smallest allowable work time is chosen. Similarly, the rest-time requirements implied by the several criteria are compared and the longest time is selected.

The value for the limiting temperature is specified by the user (on CT3/5), either directly as a temperature or indirectly as the allowable probability of collapse; the limiting temperature is derived from the allowable collapse probability on the basis of the straight line relationship proposed by Braddock, Dunn, and McDonald, Inc., in which the

collapse probability is zero below 38.38°C and 100 percent at 41.14°C (101.1 and 106.0°F).

When a task is stopped, either because it has been completed or the work limit has been reached, subroutine MANAGE calls subroutine STOPIT if USECW > 0; this subroutine manages the disposition of the work crew and either restarts an unfinished task with a fresh crew or interrupts the task if personnel are not available. Subroutine GOREST is called to check the released crew; a check is first made to see if they would have suffered from the toxic effects of chemical contamination under the existing working conditions, or if they would have collapsed from an excessive temperature or from excessive wetting or dehydration. If they have, they are hospitalized in accord with the time distributions specified for hospitalization due to these different effects; if not, a check is made to see where they are to rest and cool off, and subroutine CKTEMP is again called to determine how long the crew must rest. As explained below in the subsection on collective protection shelters, the user has a variety of options for describing where work crews will go to cool off; they may remain at their work place, they may go to some building, or they may go to a collective protection shelter. Whichever choice is made, the environmental conditions at the rest location determine the MOPP that must be worn during the rest period and the time it will take before the temperature has fallen sufficiently for the crew to be available for reassignment.

Since the exponential rise and fall of rectal temperature would theoretically lead to an infinite cooling time if the crew's temperature were to be required to fall to the equilibrium temperature at the rest location, the crew is required to wait only long enough for their nominal rectal temperature to fall to within the user-specified DELTA hundredths of degrees Centigrade of that temperature before being reassigned (exclusive of any additional time imposed by the collective protection entry delay logic). If the crew were reassigned immediately, the appropriate initial crew temperature would be DELTA above the equilibrium temperature for that facility, but if they were reassigned later, their temperature would be somewhat lower. By keeping track of how long it has been since the last person of each particular personnel type was assigned and how many are currently available, a rough correction is made to account for this additional cooling whenever a task is begun or restarted.

TSAR is relieved of the requirement to keep a record of the current temperature of all individuals because the temperature at the beginning of a task is not allowed to

depend on each individual's history but only on this equilibrium temperature. This is possible because crews in TSAR are not broken up and reassigned as parts of other groups until they have rested sufficiently for their temperature to fall to the required level. Although this requirement is not completely realistic—crews might be required to work on some tasks until they collapsed—it represents what is believed to be a reasonable compromise between a more realistic treatment and the substantially greater detail that TSAR would have to deal with to maintain a continuous record of the temperatures of all individuals.

Having said all that, there is an important exception. For those groups of people who are designated (on CT15/1) as munitions load crews, who are required to carry out a series of tasks to prepare an aircraft for flight and will be the only such crew present with the aircraft, a series of tasks is permitted. Under these conditions, it is practical to use the temperature of the crew at the completion of one task as the initial temperature for the next, and thereby allow one load crew to carry out a sequence of tasks without an interruption for a mandatory rest period. As explained elsewhere, all such TSAR preflight munitions loading tasks must go to completion and cannot be interrupted, so the load crew's final temperature is estimated when they begin each task. If the temperature will exceed the specified limit, they are not permitted to start unless it is the first such task.

4. SURFACE CONTAMINATION AND VAPOR FROM CHEMICAL ATTACKS

Airfields may be attacked with chemical agents by aircraft or surface-to-surface missiles. Aerial-delivered chemical munitions include modified conventional unitary and cluster bombs, specially designed spray bombs, spray tanks, and air-to-surface missiles. The chemical weapons would generally be airburst (over or upwind of the airfield), creating an immediate vapor and liquid hazard from the falling, evaporating agent droplets and a residual hazard from both the liquid droplets on the ground and vapor from evaporation of those droplets.

Surface deposition patterns of the droplets depend upon agent release parameters (e.g., altitude, shape, size, and orientation of the initial agent droplet cloud), meteorological parameters (e.g., wind velocity and direction as a function of height, ambient air temperature, atmospheric stability parameters), and agent parameters (e.g., density, diffusivity, saturation concentration, droplet diameter distribution). The

contours of surface deposition patterns vary from the nearly rectangular shapes of aircraft spray tanks releasing an agent perpendicular to the wind direction to the cigar-shaped patterns created by ballistic missiles.

Each of the delivery systems has constraints on the delivery parameters (e.g., the nearly vertical trajectory of ballistic missiles or a requirement for low-altitude penetration of air defenses by aircraft). However, within the constraints imposed by the delivery system, chemical munitions will generally be designed to release their agent fill at altitudes that tend to maximize some value of area coverage by agent deposition on the ground (e.g., to maximize the area covered by a density of 100 mg/m^2 of agent).

Deposition patterns can be obtained from test data or from computer models of atmospheric transport, dispersion, and evaporation of liquid droplets. Two such models are NUSSE2 [13] and the model described in [14].

In TSARINA, a deposition pattern is represented by a set of up to 17 rectangular or elliptical "layers" (see Fig. 10). The pattern has a wind velocity associated with it and a reference point (usually the burst point of the chemical weapon) from which the




Fig. 10—Illustrative fit to the surface deposition contours by rectangular layers

rectangles are offset. The rectangles (ellipses) are perpendicular to the wind direction, but otherwise arbitrary relative locations and sizes. Each rectangle (ellipse) represents a constant surface contamination density and the total contamination density at a point is the sum of the densities for the rectangles (ellipses) containing the point. In the model we use to predict vapor concentration over and downwind from the surface pattern, the vapor concentration depends upon the mass median diameter (MMD) of the droplets on the surface. Each rectangle (ellipse) has two MMDs, one for the upwind half of the pattern and one for the downwind half (large droplets fall faster than small droplets so that the average MMD in the upwind half of the pattern is larger than that of the downwind half).

Delivery parameters for each weapon delivery pass include the attack heading, the desired mean point of impact of the pattern reference point (for a single weapon or for the middle of a stick of weapons), the aiming error, and the ballistic error. The location of the arriving weapons and the "actual" wind direction and velocity at the time of the attack are determined using Monte Carlo procedures. The downwind offsets and dimensions of the pattern rectangles (ellipses) are scaled by the ratio of the sampled wind velocity to the nominal wind velocity for the pattern.

Monaghan and McPherson Rough Surface Evaporation Model

There are several models for estimating surface evaporation of chemical agents [15, 16]. For use in TSAR/TSARINA, we have adopted a version of the rough surface evaporation model developed by J. Monaghan and W. R. McPherson [16], which provides estimates of both the residual surface deposition and the evaporation-created vapor threat over and downwind of the contaminated area.

Meteorological conditions (temperature, wind velocity and direction, and atmospheric stability), chemical/physical properties of the agent, type of surface, and the liquid loading are the primary factors affecting evaporation rates and the resulting downwind vapor. Higher temperatures and wind velocity increase the evaporation rate but higher wind velocity may actually decrease the downwind vapor concentrations. The other factors affect evaporation or absorption rates or the downwind dispersion of the vapor.

Although the Monaghan and McPherson model could be used for evaporation from other surfaces, it has been applied only to grasslands. The version of the model used for TSAR/TSARINA contains model parameters selected to fit test data for

evaporation from grasslands[16]. Liquid droplets sprayed on grassland impact at various heights above the ground surface and form liquid films on the grass and underlying debris. The thickness of the films decreases with increasing height and is an order of magnitude thinner near the top of the grass than at the bottom. Ventilation of the surface by the turbulent atmosphere increases from the ground to the top of the grass. The evaporation rate is proportional to the area covered by liquid film and is constant until the thinner films at the top of the grass are completely evaporated.

Surface Deposition Densities and Vapor Concentrations

TSARINA determines surface deposition densities at aircraft shelters, aircraft parking areas, taxiway segments, and designated buildings by adding up the surface deposition densities for each rectangle of each chemical weapon that contains the location coordinates of those facilities. These are output for TSAR on "40" cards for each attack.

In addition, steady-state surface deposition densities and vapor concentrations and parameters for determining residual deposition densities and vapor concentrations over time are output for up to 150 arbitrarily selected monitoring points on the airbase. For each attack, TSARINA creates an output dataset that summarizes the characteristics of the vapor produced by each half-rectangle of each weapon at each monitoring point. The characteristics include the steady-state values of surface deposition density and vapor concentration, the arrival time of the surface deposition (the downwind distance of the monitoring point divided by the wind velocity), and key time constants of the decay over time. TSAR then determines surface deposition densities and vapor concentrations at each monitoring point by adding up the contributions from each half-rectangle.

TSARINA also outputs on "40" cards the closest monitoring point to each aircraft shelter, aircraft parking ramp, taxiway segment, and designated building. At attack time TSAR uses the "40" card input of surface deposition density at the facility and the vapor concentration at the closest monitoring point, but at all other times TSAR uses the residual surface deposition density and vapor concentration at the closest monitoring point.

5. INDIVIDUAL CHEMICAL PROTECTION EQUIPMENT

When the effects of chemical protection and chemical attacks are to be assessed, the user must specify what chemical ensembles are worn at the different airbases, and what portions of these ensembles (MOPP) are worn before any attacks are sustained. The value for the preattack MOPP for each of up to three chemical ensembles is specified individually for each of the five generic task types and for off-duty personnel (entered using CT3/4 and CT3/7 and stored in the MOPPOL array). These choices, together with the atmospheric conditions that the user has specified for each base (on the CT17/1), the heat transfer and permeability characteristics of the ensembles (entered with CT43/1), and the temperature of the work place (a property of each facility's CW Type as defined with the CT44/1), determine how long each task may be continued and how long each crew must rest when they stop and, ultimately, how many sorties the organization is able to generate. Even in the absence of an air attack, the inefficiencies induced by wearing portions of the ensemble in anticipation of an attack can appreciably diminish an organization's sortie generation capability.

When an airbase is attacked, TSAR carries out a long sequence of calculations based on the data supplied by TSARINA, to reflect the physical damage from conventional weapons and the toxic effects of chemical weapons, as is discussed in Sec. IX.7.

The information transferred by TSARINA includes an estimate of the number of minutes it will take before the early, heavier portions of the surface chemical deposition will arrive at each on-base location of interest. These times are compared with the user's specification of the warning time that on-base personnel are to be assumed to receive and the time that it takes individuals to don various portions of their ensemble (specified with the CT43/4), to generate an estimate of which personnel will need to be hospitalized and which will be fatalities. For the first attack with chemicals the user may stipulate (CWRISK on CT3/5) that some fraction of the personnel masks do not fit well, and additional casualties will be assessed from this cause.

When the on-base distribution of the surface depositions of chemical agents and the resultant agent vapor have been determined, TSAR uses the threshold data (from the CT4/4 inputs) to determine the MOPP that personnel must wear at various on-base locations in the open, in aircraft shelters, or in any of the many facilities. If the user wishes to represent poor monitoring equipment he should set those thresholds so as to

imply that most of the ensemble must be worn even at quite low concentrations and to keep well suited when there is contamination anywhere on base. The control variable VARMOP (CT3/4) facilitates the latter constraint; when set to 1, personnel are assumed to wear the portion of the ensemble that is appropriate for their immediate work place (though still constrained by the intensity thresholds), but when VARMOP = 0, all on-base personnel must wear the MOPP that would be appropriate for the most highly contaminated on-base facility of the same CW Type in which they are to work. By raising and lowering the thresholds relative to the actual safe levels, and by use of VARMOP, the user can substantially vary the implied characteristics of an airbase's measuring, monitoring, and notification capabilities.

The on-base surface contamination and vapor concentration at each monitoring point are updated periodically every CWFREQ (CT3/4) hours after an attack, and the required MOPP at all work places is adjusted accordingly, on the assumption that the ambient conditions at the work place can be approximated by the conditions at the monitoring point closest to each work place. With these data, TSAR is able to represent the variation in working efficiency as the chemical and weather conditions vary throughout the day and the effect of these conditions on sortie generation.

These same considerations also affect the times for personnel rest and cool off. When a crew stops a task, a decision is made as to where and how long they are to cool off. They may be sent to a collective protection shelter, or they may stay at their work place as explained fully in the next section. The ambient conditions of the resting place and the ensemble that they must wear to be safe determine the time required for their temperature to fall to the required level.

6. REPRESENTATION OF COLLECTIVE PROTECTION SHELTERS

When personnel get overheated as a result of working in their IPE, it is necessary for them to stop work and cool off if they are to avoid collapsing from heat exhaustion, excessive perspiration, or dehydration. In general, personnel could relax at their work place, seek a facility where they can doff part or all of their IPE, or even go off base to obtain conditions conducive to achieving the necessary cooling and relaxation. The TSAR simulation has been designed to permit the user to represent most of the diverse behavior patterns that might be experienced at different airbases.

The user describes the conditions that are to be simulated with three distinct sets of controls: (1) the basic control variable USECP (CT3/4); (2) the locations of special CPS, along with data describing their nominal entry delays (CT43/5); and (3) individual facility data (when available) that defines the simultaneous entry capacity of the entry portal (CT37) for each building and the time to process an individual through that portal. By manipulating these controls, the user can specify a diverse set of situations, as explained below.

Values for the basic control variable, USECP, are defined as follows:

- USECP = 0 Personnel remain at their work place to cool off; no collective protection shelters are available
- USECP = 1 or 2 Personnel seek CP shelters with a frequency and extra delay time as specified with the CT43/6 options
- USECP = 3 or 4 Personnel seek CP shelters as defined by the CT43/6 cards but may have to queue at the entrances when too many seek admittance in too short a time; queuing characteristics are specified on the CT37 card for each shelter
- USECP = 1 and 3 CP shelters are not used if the base is not contaminated by chemicals
- USECP = 2 and 4 CP shelters are used without regard to the chemical contamination

The CT43/6 cards permit the user to define up to four distinct groups of collective protection shelters. A separate group may be specified for the personnel that carry out each of the four generic types of work: (1) flight line, including weapons loading and refueling, (2) backshop repairs of equipment and spare parts; (3) munitions assembly; and (4) runway, taxiway, and building repairs. Alternatively, two or more of these sets of personnel may be assigned to use the same group of CPSs, and some personnel may not need to be assigned special shelters. (Also, the flight line personnel will use "buildings" #31, #32, and #33 if special CPSs are not assigned for these personnel.) Each group of CPSs is a collection of buildings, each of which has distinct location, capacity, construction, and chemical protection characteristics.

As noted on the format for CT43/6, the user must specify (1) the building number of the first member of the CPS group that each set of personnel are to use, (2) the nominal entry time, and (3) the entry time distribution for that group of CPSs. Whenever a work crew must stop a job and cool off in a CPS, one building in the CPS group is selected for each five personnel; normally the location is selected at random, based on

the relative capacity of each of the members of the CPS group (as specified with CT37 for each building). The time required for personnel to cool down is calculated on the basis of the ambient environment within the selected building and on the assumption that the personnel will remove whatever components of their IPE are safe under those circumstances (the cooling time is the time it would take the rectal temperature to fall within DELTA degrees of the current equilibrium temperature within that building). When USECP = 1 or 2, a time is added to the cooling time that the user specifies as his estimate of the time that would be needed to get to the CPS, to get into the shelter, and to begin to cool off. The personnel are considered to be unavailable for a time that equals the calculated time to cool off, and the delay time, or "entry time" (unless the normal shift change intervenes, in which case the cooling off is assumed to be completed as part of their off-duty activities). A different procedure is followed when USECP > 2, as will be described in the next subsection.

The user is given several options for representing this entry time delay when USECP equals 1 or 2. In one of the simpler options, the user may specify that the added time is approximately the same for all buildings in a group of CPSs; the "Entry Time" entry on CT43/6 is used in that way when USECP < 3. Other options permit the user to capture some of the variability likely to be experienced by specifying that the actual delay time for each set of personnel that stop to cool off is to be selected from a distribution whose mean is the "Entry Time." The distribution to be used can be either a number less than 20, which refers to the correspondingly numbered distribution stored in subroutine TTIME (see App. I in Vol. III), or it can be a number equal to or greater than 20, which invokes a quite different set of options.

If the user wishes to simulate the situation in which personnel will always go into a CPS to cool off and will always stay only long enough to do that, a distribution whose number is less than 20 (including zero, of course) should be specified. If personnel will sometimes need to spend a longer time in the CPS for eating or performing other bodily functions, a distribution equal to or greater than 20 should be specified. When the latter is done, and the distribution number is formed by the two digits "X" and "Y" (where $20 \leq XY < 100$), the nominal "Entry Time" delay will be added to the cooling time in $(X - 1)DDX$ of the events, and it will be Y times as long in $1/X$ of the events. If it is desired that this long delay always occur in some particular location—for example, at some distant building with better accommodations—the user should set REMOTE to

unity. When that is done, all personnel who are to experience the longer delay will be assigned to the first member of the CPS group, independent of the relative capacity of the members, if that building has not been damaged; if it has been damaged, the selection is as though REMOTE = 0. (Also, if the capacity of the first member is zero, only those with the longer delay will be sent there.)

A related option permits the user to simulate the situation in which personnel will go into a CPS only on some occasions and will cool off by relaxing at their work place on others. This behavior is represented when $XY \geq 20$ and $Y = 0$; in this case the personnel select a CPS to cool off in for $1/X$ of the events, and cool off at their work place, without any special delay, for the others; the nominal "Entry Time" is added to the cool-off time when they go to the CPS.

Under some conditions the user may wish to place personnel in a CPS only occasionally but may also wish to assume that those who cool off at their work place will take some additional time to get there, in addition to the cooling off time. For this option the building number entered as the "1st CPF" on the CT43/6 should have a minus sign in front of the building number; that acts as a flag to identify this option. When this option is chosen (with $XY \geq 20$ and $Y \neq 0$), the personnel cool off at their work place in $(X - 1)/X$ of the events and are assumed to sustain the delay of "Entry Time"; for the other $1/X$ of the events the personnel select a CPS to cool off in and have an added delay of Y times the "Entry Time." (When $Y = 0$ this option defaults to that described earlier for $Y = 0$.)

Explicit Simulation of Queuing at Collective Protection Shelters

The several CPS options described above do not simulate the queuing that could develop at the CPS entries but assume that the user-specified "entry time" in conjunction with the several possible distributions approximates the effects of whatever queuing might occur. When USECP is greater than 2, queuing at the CP shelters is simulated explicitly, using the entry-portal capacity and processing time that may be entered for each building on the CT37 card. (If these data are not entered for a shelter when USECP > 2, the processing time is assumed to be the "Entry Time" from the CT43/6, and the portal capacity is assumed to be one.) These data specify that N persons may be processed into a CP simultaneously and that it takes an average of M minutes for each individual to enter; hence, when P persons must cool off, it is assumed that they take $T = M \times P/N$ minutes to be processed into the CPS. If additional work crews attempt to

enter the same CPS during that time, it is assumed that they will not begin to process into the CPS until the T minute period is complete. The total time that a team of personnel is unavailable for work is the time they take to get into the CPS, including any queuing delays before they may be processed, plus the time required within the CPS to cool off (we assume that no cooling occurs during the queuing and entry process).

If the user has specified a time distribution less than 20 for a CPS group (with the CT43/6), all work crews that must cool off select a CPS, and the processing time considered in the queuing computations is varied in accord with that distribution from one work crew to the next. If a distribution of XY is specified, where $XY \geq 20$, a CPS is entered for only $1/X$ of the events, with all other work crews cooling off at their work place. If $Y = 0$, only the queuing and processing times are added to the cooling time as just described; if $Y > 0$, an additional delay is added for $1/Y$ of the occasions that a CPS is used, and the total time that the work crew is assumed to be unavailable is taken to be the sum of (1) the queuing time, (2) the processing time, (3) the time to cool off, and (4) the additional time spent in the CPS for eating or other bodily functions, which are equal in magnitude to the "Entry Time" on the CT43/6. If REMOTE is 1 and $Y > 1$, all work crews with the additional delay will be assigned to the "1st CPS" if the building has not been damaged. When USECP > 2 , a negative sign on the building number "1st CPS" will have no effect on the simulation (but it does when USECP ≤ 2 ; see previous subsection).

7. ESTIMATION OF CASUALTIES FROM CHEMICAL EFFECTS

Typically, chemical weapons release their agent fill at fairly high altitudes. The resultant agent spray droplet rain falls to the surface where it begins to be absorbed by the surface material and to produce toxic vapor by evaporation. Both the residual surface deposition and the vapor concentration decrease over time as the liquid droplets are absorbed and evaporated and are gone after a few hours, the actual time depending primarily on the persistency of the agent, wind velocity, temperature, and the surface material. During this time toxic doses of chemical agent can be received either percutaneously (through absorption of the agent from the droplets or vapor by the skin or eyes) or by vapor inhalation.

In TSAR, casualties from chemical attacks are evaluated at the times of the attacks and at certain times between attacks. At the time of an attack, it is assumed that the

personnel receive warning of the attack, stop their current activities, don their attack MOPPs (see Sec. IX.5), and then are released from their jobs if they are working. The evaluation at attack time assesses casualties from the percutaneous effects of the agent rain from the attack and the percutaneous and inhalation effects of the agent vapor up to the time the attack MOPP has been donned. In addition, at the time of the first attack a crude estimate is made of the number of casualties caused by ill-fitting masks. Then, whenever a personnel state changes by being released from a job or by finishing a rest period, an evaluation is made of casualties to the personnel involved during the elapsed time in their previous state. However, the evaluation of casualties between attacks underestimates the actual number of casualties since TSAR does not currently evaluate all personnel states or accumulate vapor dosages over time for individual personnel (a preliminary design for a procedure to do so has been made but has not as yet been implemented) and uses only the dosage accumulated while in a particular state. In the applications to date, this has caused no serious underestimation of casualties because the attack MOPPs specified have provided sufficient protection that essentially no casualties would have been produced between attacks.

At the time of the first attack, personnel are assumed to be in their preattack MOPP. Then, upon being warned of the attack, the personnel don the protective clothing of their attack MOPP. The actual warning time depends upon the value of an input warning time parameter and the time taken for the chemical spray droplets to reach the ground. The warning time parameter (1WARN for the first chemical attack on a base and WARN for subsequent chemical attacks—CT3/5) defines the time (in minutes) before the attack at which all personnel are notified to don their attack MOPP (negative warning times imply that the personnel are notified that many minutes after the chemical weapons burst).

The personnel first put on their protective masks (if part of the attack MOPP) and then don the rest of the protective clothing. The total vapor dosage received is the sum of that received before and after donning the protective mask. The dosage received before donning the mask is the integrated vapor concentration at the facility where the personnel are located. This depends upon the outside vapor concentration and the filtering and air exchange characteristics of the facility. In TSAR, we assume a simple air exchange model between the inside and outside air leading to the equation

$$C_{in}(t) = (1 - F)C_{out} + [C_{in}(0) - (1 - F)C_{out}] \exp(-t/T) \quad (1)$$

where $C_{in}(t)$ is the vapor concentration inside the building at time t after the droplets reach the ground, C_{out} is the outside concentration (assumed constant for the first few minutes after the attack until the protective clothing is donned), F is the fraction of the vapor filtered from the outside air as it enters the building, and T is the time for one air exchange between the inside and outside air (see Fig. 11).

Defining time zero to be the time after the attack at which vapor first arrives and t_m to be the time when the mask is in place, the dosage, D_m , received before the mask is in place is then the integral of the concentration from zero to t_m , or

$$D_m = t_m(1 - F)C_{out} + [C_{in}(0) - (1 - F)C_{out}]T [1 - \exp(-t_m/T)] \quad (2)$$

The mask acts as a filter between the vapor in the local environment and the inside of the mask. Defining F_m to be the fraction of the vapor filtered by the mask and t_1 to be the time at which the protective clothing has been donned, the dosage, D_1 , received by the individual between the time the mask is in place and the rest of the protective clothing is donned is

$$D_1 = (1 - F_m) \{ (t_1 - t_m)(1 - F)C_{out} + [C_{in}(t_m) - (1 - F)C_{out}] T [\exp(-t_m/T) - \exp(-t_1/T)] \} \quad (3)$$

The total inhalation vapor dosage D received after the attack until the protective clothing has been donned is then $D = D_m + D_1$.

The percutaneous vapor dosage, D_p , received before donning the attack MOPP is obtained by replacing t_m in Eq. (1) by t_1 . The effect of this dosage is then evaluated assuming that the personnel are in their preattack MOPP.

The percutaneous effect of the initial surface contamination created by the falling chemical spray droplets depends upon the actual amount of droplet spray impinging upon the individual and the MOPP being worn. We assume that the MOPP is either the preattack or the attack MOPP, depending upon whether the warning time is sufficient that the attack MOPP is in place by the time the spray droplets reach the given location (i.e., we give no credit for part of the additional protective clothing to be on when the droplets arrive).

For personnel in the open, we use the relationship between the dose on the man and the agent density on the surface given in [17], as modified by the Aerospace Medical Research Laboratory,

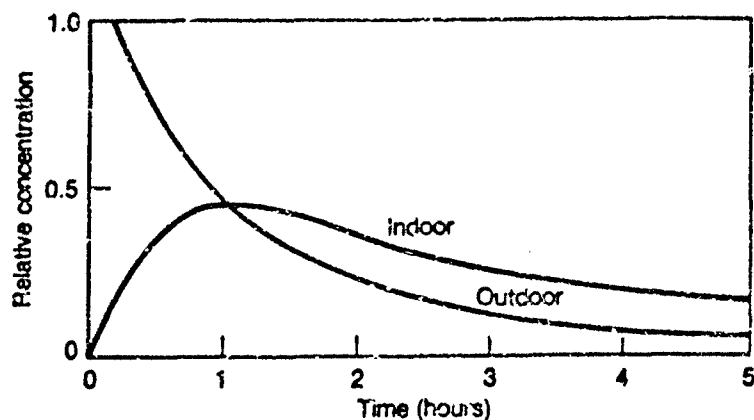


Fig. 11—Indoor and outdoor vapor concentration for a ventilation rate of one air exchange per hour

$$\begin{aligned}
 DM &= DG & DG &\leq 10 \text{ mg/m}^2 \\
 &= 1.62DG^{0.80} & 10 \text{ mg/m}^2 < DG \leq 1152 \text{ mg/m}^2 \\
 &= 17.8DG^{0.46} & DG > 1152 \text{ mg/m}^2
 \end{aligned} \tag{4}$$

where DM is the dose on the man (mg) and DG is the agent surface density (mg/m²).

For personnel not in the open (i.e., with some type of overhead cover), we use the same man dose/surface density relationship above except that the surface density is replaced by the density beneath the overhead cover. To obtain the surface deposition density beneath the overhead cover, the outside surface deposition density is multiplied by a factor representing the fraction of the overhead deposition actually getting under the cover (this fraction is generally zero for most buildings but may be a nonzero for buildings that are partially open—e.g., an aircraft shelter or hangar with an open door).

To represent one possibility for accidental exposure to toxic agents, a crude representation of the effect of ill-fitting masks is included in TSAR. The parameter CWRISK (CT3/5) specifies the fraction of the personnel that have ill-fitting masks at the time of the first chemical attack. This fraction of the personnel is then assumed to receive four hours of inhalation dosage, assuming the chemical environment of their work or rest place at the time of the attack (the four-hour time period was chosen arbitrarily as a typical time before the faulty mask is discovered and corrected).

Between attacks, the agent surface density and vapor concentrations at the monitoring points are updated every CWFREQ hours. This includes an update of the vapor concentration inside each building type (i.e., CWTYPE) associated with each monitoring point. A version of Eq. (1) is used for this update, by setting t_m equal to CWFREQ and assuming a constant value of the outside vapor concentration over the time period. Whenever a crew is released from a job or has finished a rest period, the vapor dosage accumulated during the time, t_w , in that state is taken to be the current vapor concentration in the corresponding building type for the closest monitoring point times t_w .

Combining Percutaneous and Inhalation Doses

When an individual is subject to both percutaneous and inhalation doses from a toxic agent, some method must be used to combine the toxic effects. Very little discussion of this problem was found in the literature on toxic agents, and evidently few data exist.

The method we have adopted to combine percutaneous and inhalation doses is from [18]. It has reportedly produced reasonably good results for the combined response for doses of two similar poisons administered together to insect populations and for the same poison applied to different locations on the insects.

Each individual in the population is assumed to have tolerance doses for the two different forms of the agent represented by a bivariate probability distribution of tolerance doses for the population as a whole. If Z is the tolerance dose for one form of the agent, an individual responds if the applied dose z is greater than Z , and fails to respond if z is less than or equal to Z . For two different forms of the applied dose, the fractions Q_1 and Q_2 of the population failing to respond to separate applications are

$$Q_1 = P(z_1 \leq Z_1) = P(z_1/Z_1 \leq 1)$$

and

$$Q_2 = P(z_2 \leq Z_2) = P(z_2/Z_2 \leq 1)$$

where $P(z_i \leq Z_i)$ is the probability that the tolerance dose Z_i for an individual is greater than or equal to the applied dose z_i .

It is next assumed that an individual fails to respond to a combined dose if the sum of the individual fractions applied dose/tolerance dose does not exceed unity. Then, the

fraction of the population Q failing to respond to the combination of doses is

$$Q = P(z_1/Z_1 + z_2/Z_2 \leq 1) \quad (5)$$

Two further assumptions lead to a unique solution for Q: (1) the standard assumption of lognormal distributions for each of the two tolerance doses, and (2) a bivariate lognormal distribution with complete correlation for the joint distribution of the tolerance doses. By the first assumption, $\log(Z_i)$ has a normal distribution (with mean $\log(u_i)$ and standard deviation S_i , say). Let

$$X_i = (1/S_i)\log(Z_i/u_i)$$

so that X_i has a normal distribution with mean zero and unity standard deviation. Equation (5) may then be written as

$$Q = P(10^{(x_1 - X_1)S_1} + 10^{(x_2 - X_2)S_2} \leq 1) \quad (6)$$

where $x_i = (1/S_i)\log(z_i/u_i)$.

Since Z_1 and Z_2 are perfectly correlated from the second assumption, X_1 and X_2 are also perfectly correlated so that $X_1 = X_2$ with probability one. Let $X = X_1 = X_2$. Then Eq. (6) may be written as

$$Q = P(10^{(x_1 - X)S_1} + 10^{(x_2 - X)S_2} \leq 1) \quad (7)$$

Define $f(x)$ as

$$f(x) = 10^{(x_1 - x)S_1} + 10^{(x_2 - x)S_2}$$

Since $f(x)$ is monotonically increasing, Eq. (7) may be written as

$$Q = P(X \leq f^{-1}(1)) = N(\Gamma^{-1}(1)) \quad (8)$$

where $\Gamma^{-1}(1)$ is the value x_0 satisfying $f(x_0) = 1$, and $N(\cdot)$ is the cumulative normal distribution with mean zero and unit standard deviation. The value x_0 is called the NED (normal equivalent deviation) for the combined doses z_1 and z_2 . From the definition of $f(x)$, x_0 satisfies

$$10^{(x_1 - x_0)S_1} + 10^{(x_2 - x_0)S_2} = 1 \quad (9)$$

and Eq. (8) reduces to

$$Q = N(x_0) \quad (10)$$

In the literature on toxic substances, $1/S_1$ and $1/S_2$ are called probit slopes for the responses of the population to applied doses of the two different forms of the agent. When the probit slopes are identical, let $S = S_1 = S_2$. Then Eq. (9) may be written as

$$10^{(x_1 - x_0)S} + 10^{(x_2 - x_0)S} = 1$$

so that

$$\begin{aligned} x_0 &= (1/S)\log(10^{x_1} + 10^{x_2}) \\ &= (1/S)\log(z_1/u_1 + z_2/u_2) \end{aligned} \quad (11)$$

In the literature, the toxic effects on humans are quantified in terms of the dosage required to cause a given effect in 50 percent of the population being considered (the median dose— u_i in the above development) and the probit slope (the reciprocal of S_i) of the distribution of response tolerances, usually assumed to be the lognormal distribution.

For toxic vapor, the median dosages for personnel incapacitation or lethality are denoted by ICT50 and LCT50, respectively, where CT stands for concentration \times time, and the usual measurement unit is mg-min/m³. For toxic liquid, the corresponding terms are ID50 and LD50 for incapacitation and lethality, respectively. Here, the D stands for dose; the usual measurement unit is mg.

The values of the median doses depend upon the toxicity of the agent and the MOPP—even very light clothing can increase the dose required for a given effect by an order of magnitude or so over that required when the dose is applied to the bare skin.

In TSAR, the development above, resulting in Eqs. (9) to (11), is used to evaluate the combined effects of percutaneous vapor and liquid and inhalation vapor. We assume that the probit slopes for percutaneous liquid and vapor effects are the same and combine the doses as

$$D_p = D_q/ED50 + D_v/ECT50$$

where D_p is the effective percutaneous dose (assumed to have a median value of unity and a probit slope equal to the common probit slope), D_q is the liquid dose, D_v is the vapor dose, and ED50 and ECT50 are median doses for the effect (incapacitation or lethality). This procedure is consistent with the development above (in Eq. (11)) for combining dosages. Then, the fraction of personnel suffering a given effect

(incapacitation or lethality) from the joint effect of the (combined) percutaneous dose and the inhalation dose is obtained from Eq. (10), after using the (combined) percutaneous and inhalation doses in Eq. (9) (along with the corresponding median doses and probit slopes) and solving for x_0 .

X. COMMUNICATIONS

TSAR allows for the representation of scheduled shipments of material from CONUS to the theater, special shipments from CONUS in response to theater requests, intratheater shipments of resources, and the transmittal of airbase status information. The schedules for each of these types of transfers are controlled by the user's specifications, as are the contents of scheduled CONUS shipments; the contents of the other transfers are generated endogenously.

1. SCHEDULED SHIPMENTS FROM CONUS

The user must initially specify resources scheduled to be delivered from outside the theater after the beginning of the simulation. These data are entered with CT31; the delivery times are arranged in a time-ordered queue in the CONUS array, and the cargo is stored in the CARGO array at the time of entry. The destination and time of delivery should be mentioned on the first of a set of cards when all the commodities on those cards are to arrive together.

The only resource classes that may not be shipped from CONUS are aircraft shelters and other facilities. No more than 99 units should be entered, except for munitions and TRAP, for which the limit is 6250. If more are required, enter the commodity twice for the same delivery. For POL, TSAR assumes that the unit of measure for shipments is hundreds of thousands of pounds, whereas fuel normally is stored and used in thousand-pound units. (Storage capacity for POL may be enhanced by specifying a shipment of Type #100 POL; units of measure are the same as for POL.)

When an arrival is noted in subroutine MANAGE, control is transferred to the RECSUP (receive supplies) entry point in the DOSHIP subroutine and the resources are added to the stock levels at the appropriate base. When new ground personnel, support equipment, or aircraft parts arrive, subroutine CHECK is called to check whether they may be used immediately; for ground personnel, the new personnel are added to the day and night shifts to maintain the ratio of the shift sizes in the same proportions as specified by the "target" levels for each personnel type in the initializing data for each base. When munitions arrive that have been shipped disassembled and the components required for

assembly have been specified, the shipment is broken down and stored in the form of the appropriate components.

When aircraft are ferried to the theater from CONUS, they are added to the inventory at the appropriate base and undergo a normal postflight inspection, except that attrition is not checked. The aircrew is attached to the base's flight staff and given 24 hours to rest before their first assignment. Aircrews that are ferried to the theater (arrive without aircraft) are treated in the same manner.


2. RESPONSIVE SHIPMENTS FROM CONUS

The user also may simulate the requisition and resupply of resources from CONUS for any class of resources except shelters and the other facilities. When activated (CT33), a requisition is submitted for resources that are lost in combat or during an airbase attack and, in the case of parts, for parts that may not be repaired in the theater. The resources requested are delivered after the delay specified by the user for each of the various resource classes. Arriving resources are treated identically to those described above.

3. INTRATHEATER SHIPMENTS

Resources (except aircraft, aircrews, shelters, and facilities) may be transferred from one airbase to another using an intratheater system.¹ The description of this system is controlled by the user's specifications of the schedules and the statistics governing their delays, cancellations, and losses on CT32/1 and CT32/2. These shipments do not involve specific resources (e.g., trucks or aircraft), nor are they capacity limited; they provide a representation of the times expended between the time that supplies are consigned for shipment and the time that a shipment reaches its destination and the cargo is added to base supplies. The algorithms governing the transfer of resources with the intratheater transportation system are outlined in Sec. XI.2. The factors that are considered in this representation and the card types that are used to input the relevant data are summarized in the following sketch.

¹Aircraft and aircrew transfer can be affected exogenously by specification of a different recovery base with a flight demand or endogenously by directing aircraft transfer, as discussed in Sec. XI.1.



The user may specify daily departure times on an individual basis for each origin-destination combination. The user may also control the mean departure delay, mean in-transit time, and the distribution of these values on an individual basis, using any of the 15 distributions that may be stored in the TTIME (true time) function. By manipulating the shape of these distributions, a fraction of the shipments may even be canceled;² the commodities that had been prepared for that shipment are then scheduled on the next shipment. The user may also specify a loss rate for the shipments between any two bases; the commodities on these shipments are not recovered.

The schedules for the various departures and arrivals are organized into time-ordered queues in the SHIP array with the subroutine SCSHIP (schedule shipments). These schedules are first organized at the time the program is initialized and subsequently at midnight at whatever interval the user has specified with the control variable SHPFQ. Any of the schedules may be changed at any time during the simulation in much the same manner as the demands for aircraft sorties (see the READFT subroutine for instructions).

The data stored for each shipment include pointers to the next departure from the same base to the same destination, to the next departure from any base, and to the next arrival at any location, as well as a pointer to the location of the first package to be included with the shipment. The resources themselves are stored in the SHIPQ array.

Resources are prepared for intratheater shipment with a call to the SHPRES (ship resource) subroutine, which checks on the availability of the quantity stipulated and decrements the shipper's stocks appropriately. For ground personnel, the work force is reorganized and reassigned as necessary with a call to subroutine REDPEO (reduce people). When the commodity is a faulty part rather than a serviceable part, that fact is denoted by a negative part number. When ground personnel, AGE, or serviceable parts are shipped, the numbers enroute to each base are tallied in the appropriate storage arrays for these resources; similarly, faulty parts enroute to a CIRF are tallied in that base's

²Delays greater than 18 hours are interpreted as cancellations.


portion of the PARTS array. The restrictions as to the size of the individual lots that are shipped are outlined in Sec. XVII. The quantities of these resources that are enroute are available for possible use in the theater resource management algorithms.

When an intratheater departure or arrival is noted in subroutine MANAGE, control is transferred to the subroutine DOSHIP. For departures, it is necessary only to update the appropriate pointers; for arrivals, processing follows the same procedures outlined for the receipt of CONUS shipments, except that provisions are also made for the receipt of faulty parts and their transfer to the appropriate on-base repair shop.

As TSAR is currently formulated, the only use made of the intratheater shipment system is for faulty parts and for the shipment of ground personnel, AGE and equipment, and serviceable parts when imbalances are noted among the airbases by the theater resource management system outlined in the next section.

4. INTRATHEATER RESOURCE STATUS REPORTS

Although an exact count is maintained on the status of all resources on all bases throughout the simulation and these data could provide the basis on which various theater resource management systems might be examined, it seemed inappropriate to presume that the information that would be available to such managers would be precise and up to date. Indeed, one of the greatest drawbacks associated with many centralized systems is their need for high-quality communication and transportation systems. Unless some of the inefficiencies of the systems that may actually be available to our forces could be represented, it would be reasonable to question the validity of the results of any examination of schemes for managing resources on a theater-wide basis. The following sketch indicates the factors that are considered:



All necessary data are entered using the CT36 cards.

In TSAR, each base may be designated to report the then-current status of the ground personnel, AGE, and aircraft spare parts at several different times during the day. These data are collected at those times for transmittal to "theater headquarters"—to the

theater resource manager. The user also specifies the time delay before that information would be formatted, transmitted, decoded, and available to that manager; the delays and the distribution of those delays are controlled base by base. Since only one location is provided for "in-transit" information, the data arrival time must be no later than the next data collection time; if it is, the transit time is shortened and a report of that action is printed. The completeness of the report may be controlled in two ways: The entire report may be lost in a specified percentage of cases, or some percentage of the individual data may not be reported.

When the user elects to activate the theater communication system, the theater manager's database is initialized with information that is accurate at zero time. However, if the user does not wish to turn control over to the theater manager until some later time the initiation of the reports may be delayed to NEWDTA (CT3/2) by initializing OLDDATA (CT3/2) to unity. This feature avoids the related data processing during the early stages of a scenario, during which the user recognizes resource transfers would be unnecessary or undesirable. When OLDDATA is first initialized to unity and subsequently set to zero by using the special code available in subroutine CONTRL, reporting will begin when the next report is due to be sent after NEWDTA.

The particular status reports that are transmitted when the communication system is activated are controlled by the user's choice as to which classes of resources will be managed; the user may select each or any combination of the ground personnel, AGE, and spare parts as explained in the next section.

Management of the theater reporting system is the function of the STATUS subroutine. This subroutine is used first to place the various transmittal and receiving times into the REPORT array heap and to initialize the manager's data base. Subsequently, this subroutine is called by MANAGE whenever a report is to be transmitted or received. The data in transit and the data on hand for theater management are stored together in the PEORPT, AGERPT, and PRTRPT arrays. The nature of this storage arrangement necessitates that reports in transit be received before the next is transmitted; users must assure that their schedules and delays make this possible.

When a report is to be received, a check is first made to see whether it was lost in transit. Subsequently, a random number is drawn for each element of the data storage arrays and checked against the specified data incompleteness rate. The last action in the STATUS subroutine is to store the time for the next day's transmittal or receipt of the corresponding daily report in the REPORT array heap.

XI. THEATER RESOURCE MANAGEMENT

TSAR's ability to represent operations at a set of airbases and to handle the transfer of various classes of resources among those airbases can be combined to provide a unique mechanism for pretesting policies that would exert a broad span of control over theater resources. Indeed, some may view TSAR's prime role as a testbed for examining the effectiveness of new policy proposals. Although TSAR's initial formulation is concerned only with the theater-wide management of aircraft, ground personnel, and equipment, in addition to faulty and serviceable aircraft spares, it could readily be extended to managing the other classes of resources.

The range of policy options that might be examined with TSAR is limited, obviously, to those sets of decision rules that are expressible in terms of resource status information—past, present, and projected—available within this simulation. In TSAR there are basically three sets of status information available: (1) accurate data regarding current status, (2) the delayed and imperfect data provided with the theater status reporting system, and (3) an approximate projection of each base's current capability to generate sorties. In addition, a limited amount of data exist regarding future sortie demands, as well as the completion times for all ongoing tasks. The range of decisionmaking policy options that could be evaluated with TSAR should be reasonably illustrated with the following discussions of those rules that are encoded in the current formulation.

TSAR offers several options for managing aircraft resources. Initially, aircraft may fall into three different categories: aircraft assigned to an operating base, aircraft in a theater reserve of "filler" aircraft, and replacement aircraft in CONUS. If the user designates a pool of "filler" aircraft, they may be used to offset degradations due to lost and damaged aircraft, as well as aircraft with excessive maintenance requirements and those that have been withdrawn to a rear base for maintenance. These fillers may be used in addition to or instead of a reserve of aircraft in CONUS for replacing losses. The user is provided with several options on how these aircraft are used and managed; these options are discussed fully in connection with CT3/2 in Sec. XIX. When specified, replacement aircraft are used exclusively for aircraft lost in combat or from the effects of air attack, or have been so badly damaged they must be salvaged. If the user specifies

that both fillers and replacement aircraft are available at the beginning of the simulation, the losses will be replaced with a filler and the replacement will join the filler pool when it arrives in the theater if it is not then needed at the operating base.

During the simulation, decisions governing the diversion and transfer of aircraft and the assignment of sortie demands when a base has not been specified are based on the estimates of the sortie generation capabilities at the airbases that are developed each evening in the BASCAP (base capabilities) subroutine. Aircraft that must be diverted from their assigned recovery base are sent to the base that has an open runway and the highest sortie generation capability per available aircraft. Such aircraft operate at the alternative location until the runway at their parent base has been reopened and the base's sortie generation capability per available aircraft is within a specified percentage (MULTI1 on CT4/2) of the capability at the alternative base.

On certain occasions an aircraft may have to recover on an airbase that does not normally support that aircraft type and would not have the types of resources needed to carry out all the normal maintenance. Since TSAR does not currently have provisions to simulate the transfer of people or equipment between bases to fix such aircraft, it is necessary to ignore such requirements to avoid stranding aircraft unrealistically. This is done by omitting the shops and tasks (including ABDR) from the appropriate CT29 cards so that unperformable tasks cannot be designated, and by specifying on CT17/12 the base/aircraft-type pairs for which deferred maintenance should not be initiated, but retained until the aircraft has recovered at an appropriate base. For those aircraft that would not be launched on a combat mission from such diversion bases, TSAR checks every two hours to return such aircraft to their parent bases when conditions permit.

At this time all other theater resource management rules, or algorithms, are collected for convenience in subroutines CONTRL (control), OBTAIN, REALLO, and REPRTY (repair priority). If it were desired to substantially expand these sections, additional subordinate routines could be readily appended. The algorithms now encoded deal with the following five resource allocation decisions:

1. Disposition of serviceable spare parts repaired at an operating base,
2. Periodic review and reallocation of ground personnel, equipment, and aircraft spares among the operating bases,

3. Acquisition of a spare part by an operating base,
4. Disposition of serviceable spares at a CIRF, and
5. Choice among repairs waiting at a CIRF.

In several instances the user may select from alternative sets of rules for these kinds of decisions. The choices vary in both complexity and required processing; unfortunately, there is a tendency for the more complex, often more efficient rules to require the greater amount of processing, hence to absorb greater computer resources.

The sets of decision rules that apply for any particular simulation are dictated by the initial value of the control variables SHPREP (CT2/2) and CMODE (CT1). When SHPREP is initialized, parts repaired at an operating base are not automatically replaced in stock or sent to the base where they were removed from an aircraft. Rather, a check is first made to see if the newly repaired part would reduce the number of NORS aircraft at the base where it was repaired—that is, whether (NORS Aircraft - Aircraft Missing Part) is less than SHPREP at that base; if so, it is retained on base. If not, all bases are checked and the part is sent to the base with the greatest need. Newly repaired parts are always considered for shipment when SHPREP is large.

The user's choice for CMODE defines the three internal control variables CTHEA, CCIRF, and SHOPRY, since $CMODE = 100 \times CTHEA + 10 \times CCIRF + SHOPRY$. CTHEA controls the classes of resources that are periodically reviewed and reallocated; CCIRF controls the treatment of an operating base's demand for a spare part and the disposition of newly repaired or newly acquired parts at a CIRF; SHOPRY controls the selection from among the parts waiting for repair at a CIRF.

The decisions and the bases for those decisions that are controlled by these three variables are outlined below. Although each set of algorithms acts independently in the manner to be outlined, there are instances in which one rule may be overridden or negated by another. An obvious example occurs when a CIRF is directed to ship all newly repaired or newly acquired spares to one of the operating bases and the operating bases are directed to order a part from central supply at the CIRF; such requests will always go unfilled if all parts have been shipped as soon as they become available. It is important to be aware of the effect of one's choices for the various control variables and of their possible interactions.

1. MANAGEMENT OF FILLERS, AIRCRAFT REALLOCATION, AND DIVERSION

TSAR options for managing the theater's aircraft resources are designed to simulate various decisions that theater managers would, in certain circumstances, make to reduce aircraft vulnerability to air attack or to enhance the sortie generation potential of their aircraft force. Included would be the replacement of lost aircraft, the insertion of reserve aircraft to offset aircraft immobilized by the need for extended maintenance, transfer of aircraft to dispersed operating bases, and various work-leveling decisions. Such situations could arise whenever air attacks are imminent or bases suffer disproportionate losses of support resources or aircraft, or when closed runways force aircraft to divert.

Management of Filler Aircraft

A pool of filler aircraft may be defined for the theater and used to offset the degradations due to lost or damaged aircraft, as well as aircraft with excessive maintenance requirements. This pool may be used in addition to, or instead of, a reserve of aircraft in CONUS. It is assumed that an aircrew is available for each aircraft in the pool. The control variables FILLAC, MAXMNT, and FLEVEL (CT3/2) provide options as to how these aircraft are used and managed.

The value for FILLAC defines the circumstances under which a filler aircraft is assigned to an operating base. The five conditions are:

FILLAC	Conditions for Filler Usage
1	An aircraft is lost in air operations, or airbase attack.
2	An aircraft is lost, or is transferred to a rear base for battle damage repair.
3	An aircraft is lost, or is transferred to a rear base for maintenance, including damage.
4	As in 2, or when the expected repair time for an on-base battle-damaged aircraft exceeds MAXMNT, and the FLEVEL conditions are met (see below).
5	As in 3, or when the expected maintenance time for an on-base aircraft exceeds MAXMNT and the FLEVEL conditions are met.

Whenever a filler aircraft is assigned to a combat unit to replace a combat loss, a replacement is ordered from CONUS, if stipulated by the replacement policies prescribed with the CT33.

The value of FLEVEL affects the decision to augment on-base aircraft and controls the disposition of both aircraft repaired at a rear base and aircraft transferred from CONUS to the filler pool. To requisition an augmented aircraft, or to return aircraft from the rear, it is necessary that the current number of on-base aircraft be less than the quantity designated by the value of FLEVEL. Those quantities are:

FLEVEL	
0	Number of aircraft less than the initial number of assigned aircraft.
1	Number of non-battle-damaged aircraft less than the initial number of assigned aircraft.
2	Number of aircraft less than the base's shelter capacity.
3	Number of non-battle-damaged aircraft less than the base's shelter capacity.

When these conditions are not met, aircraft newly repaired at a rear base and aircraft that have arrived from CONUS are consigned to the pool of filler aircraft.

Computing Sortie Generation Forecasts to Support Aircraft Management Decisions

The basic evidence needed to assign sortie demands to airbases, or to reallocate aircraft among airbases, are forecasts of each base's capability to generate sorties of different kinds with the different aircraft types. Naturally, one cannot expect to obtain such forecasts with anything like the accuracy achieved in the simulation proper, but that simulation can indicate only the sorties that have been flown during a previous period for a particular set of aircraft and flight demands. To obtain more general estimates, a procedure has been incorporated into TSAR that provides approximate assessments of the airbase capabilities that are used to support such decisions. The forecasts developed with this procedure are updated daily and are derived so as to capture the effects of resource shortages that result from either consumption or base damage.

The substantial processing required to develop these forecasts is conducted only when the user has initialized the control variable STATE (CT4/2) greater than zero. There are two steps in the procedure: The first is conducted at program initialization and generates the expected resource requirements per sortie for each type of aircraft on each mission type. These requirements include the expected manhours for each type of personnel, the expected utilization of each kind of support equipment, part, and munitions, and the likelihood that any of the shop facilities will be required. These computations are carried out in subroutines RREQTS and REQTS1.

The second step of the procedure is to contrast, for each type of resource, the on-base assets with the per-sortie requirements. Taking the quotient as the constraint imposed on sorties by each type of resource, the basic procedure is to determine the lower bound of all such constraints for each type of aircraft and each type of mission. The calculations are carried out daily at 1930 just before the sortie demand data are read in and scheduled, and demands without a base specified must be allocated; the logic is in subroutine BASCAP and is called from MANAGE.

The actual computations in subroutine BASCAP are somewhat more complicated than just outlined for several reasons. For the mission-dependent munitions, the calculation takes into account whatever lower-priority combat loads could be loaded as well as the preferred SCL; and for parts, the forecast is modified to account for the serviceable items that would be expected to be generated by parts repair. Furthermore, three different forecasts are derived: The first forecast is made without regard to the number of aircraft on base; the second forecast introduces the additional constraint that no more than $N \times S$ sorties may be flown, where N is the number of aircraft of the type considered that do not have mission-critical "holes," and S is the maximum number of sorties that an aircraft could be flown between 0500 and ENDAY.

The third forecast provides an approximate accounting for other aircraft types that may be on base and that have common resource demands. The base's capability to generate sorties with each type of aircraft is determined by dividing the level of available assets by the aggregate demand of all aircraft types for each kind of resource (where the demands are weighted by the number of aircraft of each type).

These three forecasts are stored in the CANFLY array for each base, each aircraft type, and each type of mission. In addition, the value of the second forecast, for the mission with the highest estimated sortie potential, is stored in the SORCAP array for

each base and aircraft type. The data in these arrays provide the basis for the various aircraft management decisions during the ensuing 24-hour period.

2. PERIODIC REVIEW AND REALLOCATION OF RESOURCES

The available numbers of ground personnel, support equipment, and serviceable aircraft parts may be reviewed periodically in subroutine REALLO, and actions will be taken to redress serious imbalances that are noted. The nature and timing of these reviews are controlled by CTHEA and the user's choice for C4TM and C4INT (both on CT4/1). The first review occurs at the hour C4TM of the simulation, and subsequent reviews are at intervals of C4INT hours. The delayed and imperfect status data reported to the theater manager by the theater communications system are used in these reviews. The particular classes of resources reviewed at those times are dictated by the value of CTHEA shown in Table 6.

Ground Personnel

For each type of personnel, we first establish which base's staff has the largest and the smallest proportion of their nominal complement (adjusted for the actual aircraft on base and the numbers of sorties flown in the last two days); and then send 20 percent of that type of personnel from the best-staffed base to the worst, when certain conditions are met.

1. The gaining base has less than 75 percent of its nominal requirement.
2. The losing base has more than half its nominal requirement, and
3. The losing base has over twice as many staff members per aircraft as the gaining base.

The adjustment of the "nominal personnel complement" uses a hybrid proxy for current demand that takes into account the actual number of aircraft on hand and the number of sorties actually flown in the last two days.

AGE and Equipment

The logic applied to each type of AGE and equipment at each base is identical to that used for reallocating ground personnel.

Table 6

PERIODIC THEATER-WIDE RESOURCE CHECKS

CTHEA	Personnel	AGE	Parts
0	—	—	—
1	—	—	X
2	—	X	X
3	X	—	X
4	X	X	X
5	X	X	—
6	—	X	—
7	X	—	—

Aircraft Parts

When parts are reviewed, a check is made on whether there are more parts of each type in the central supply (i.e., at the CIRF) than were specified to be held in reserve (by the user's initialization of the nominal or "target" level at the CIRF with CT23). If there are, a check is made as to which base has the greatest need; and the parts are shipped, one at a time, until the surplus is exhausted.

To determine which base is to receive a part, the operating bases are each checked and the total number of assets of that type of part on each base is determined by summing the serviceable items, those enroute, and the reparable when the base's repair shop is functioning. That number is then reduced by the number of aircraft needing that part on that base. At this point two alternate logics are used, depending upon whether the control variable STATE (CT4/2) is zero or not. If STATE is zero, the asset count, when positive, is divided by the base's nominal allotment of that part (adjusted by a hybrid proxy for demand that involves the current number of aircraft on base and the numbers of sorties flown in the last two days); if negative, the result is multiplied by nominal part requirement. If STATE is not zero, the asset count is further reduced by expected on-base demands for that part during the time that a part would need to be in transit; that estimate is based on the requirement-per-sortie data and the base's current sortie generation potential. For either value of STATE the final result is interpreted as the relative availability of that part type on that base, and a part is shipped to that base with the numerically lowest value of relative availability. This process is repeated until there are no parts of that type at the CIRF in excess of the specified reserve; the whole process is then repeated for the next type of part.

3. ACQUISITION OF SPARE PARTS

Whenever an aircraft "hole" is reported, that aircraft's operating base may under certain conditions request and, if other conditions are fulfilled, obtain a spare part from another operating base or from the theater's central supply. The procedures used are controlled by the value of CCIRF, which also controls the rules governing the disposition of newly repaired and newly acquired parts at the CIRF. The procedures adopted are as shown in Table 7.

The procedure and conditions that govern the four different responses to a base request follow:

- a. When $CCIRF = 1$, a simple mode of lateral resupply is simulated. Whenever a "hole" is reported, the bases that the user has specified (a maximum of 14 using CT23/74) are checked one by one in numerical order, and the first base that fills the specified conditions ships a part to the requesting base. Those conditions are, first, that the number of reparable minus the number of "holes" at the requesting base is less than the value of ORDER2 (CT3/1), and second that either the donating base has at least two serviceable parts, or the donating base's adjusted stock requirement—i.e., $(\text{Nominal Stock Level}) \times (\text{Current Number of Aircraft plus a third of yesterday and today's sorties})$ divided by $(\text{Nominal Number of Aircraft})$ —is less than one-quarter of a part. As the value of ORDER2 (CT3/1) varies from a positive integer to zero to a negative integer, the policy for requesting lateral resupply can be varied from very liberal to very strict.
- b. When $CCIRF = 2$, the procedures parallel those for $CCIRF = 1$, except that all bases are checked and the base with the largest number of serviceable parts is chosen; if the donating base has only one serviceable part, the current value of its nominal stock level must again be less than one-quarter of a part.
- c. For values of $CCIRF$ greater than 2, the first action taken by the ordering base is to check whether the theater's central supply has a part that may be shipped. If so, it is shipped if the requesting base fulfills the following condition: The sum of the ordering base's number of reparable, plus the number of serviceables already enroute from the central supply, minus the number of "holes" in aircraft at that base must be less than the value of

Table 7

ACQUISITION OF SPARE PARTS

CCIRF	Base Requests for Parts	CIRF Disposal Policy
0	No response	Return to sender
1	Filled by first base fulfilling conditions	Return to sender
2	Filled by base best fulfilling conditions	Return to sender
3	Filled by CIRF when conditions permit; otherwise same as 1	Retained in stock
4	Filled by CIRF when conditions permit; otherwise same as 2	Retained in stock
5	Filled by CIRF when conditions permit; otherwise same as 1	Send to base with the most NMCS aircraft
6	Filled by CIRF when conditions permit; otherwise same as 2	Same as 5
7	Same as 3	Send to base with most NMCS aircraft if in excess of required reserve
8	Same as 4	Same as 7
9	Filled by CIRF when conditions permit; otherwise CIRF directs lateral resupply	Same as 7

ORDER1. Again, a negative value of ORDER1 defines a strict lateral resupply policy, under which parts can be requested only when the number of outstanding "holes" exceeds the tangible assets by the specified (negative) level.

If a part is not shipped by the CIRF, the requesting base then attempts to obtain a part from an operating base by a lateral resupply action. For CCIRF = 3, 5, and 7, the same procedure is used as when CCIRF = 1. For CCIRF = 4, 6, and 8, the procedure is that used when CCIRF = 2.

- d. When a part cannot be shipped by the CIRF and CCIRF = 9, the central manager checks the other operating bases to determine which can best afford to ship a part to the requesting base. This check of the other bases is based on the status information as reported through the theater reporting system.

To select the donor base the following ratio is computed for all other bases: (available parts plus enroute parts) divided by (the current level of the nominal base requirement). The base with the largest value for this ratio is directed to ship a part to the requesting base, if that value is greater than one-quarter. If it is not, but there are at least two serviceable parts at that base, one is shipped.

4. DISPOSITION OF NEWLY REPAIRED OR NEWLY ACQUIRED PARTS AT A CENTRAL SUPPLY POINT

For CCIRF = 0, 1, and 2, newly repaired parts are returned to the base where the reparable was generated; newly acquired parts are placed into the local stock. For CCIRF = 3 and 4, all such serviceables are placed into stock at the CIPF and are only shipped in response to a base demand.

For CCIRF 5 and 6, the bases are checked to see which bases, if any, need the part to satisfy a demand; a newly repaired part is sent to that base with the lowest value of [serviceables + reparables + enroute - holes] multiplied by the bases' current demand proxy, if that value is negative. For CCIRF = 7, 8, and 9, any newly acquired part that is in excess of the central supply's stipulated reserve is shipped to the most needy base. That determination is made in the same manner outlined in conjunction with periodic resource reallocations; that is, it is sent to that base with lowest ratio of (serviceables + reparables + enroute - "holes") divided by the base's current demand proxy (or, for a negative sum multiplied by that requirement). These calculations are based upon the status information reported by the theater reporting system when STATE has been initialized.

5. REPAIR PRIORITY DETERMINATION AT A CIRF

When broken parts must wait to be repaired at an operating base, their position in the appropriate shop's wait queue is based upon the local supply and demand, when the control variable ORDWT (CT3/1) has been initialized as unity, as outlined in Sec. VI.3. At a centralized repair facility a similar procedure is used, adjusted to account for the lack of local demand as such.

Whenever a reparable concludes the administrative delay at a CIRF and must wait to be repaired, it is ordered FIFO (first-in, first-out) if ORDWT is zero, or by ascending

values of the quantity RANK (i.e., items with low values receive priority over those with high values) when $ORDWT = 1$:

The quantity RANK is defined as:

$$RANK = -(100 + IMPORT) \times VALUE \text{ when } VALUE > 0$$

and

$$RANK = -VALUE \times MTBF \text{ when } VALUE < 0$$

where

$$VALUE = 2 \times (HOLES - SERVICables) - ENROUT$$

summed across all bases for the part in question,

MTBF is the mean time between failures (expressed in sorties),

and

IMPORT is a measure of a part's relative importance.

The values used to compute VALUE are the then-current values at each base if theater communications are not being simulated; if they are, the data available to the theater manager (Sec. X.4) are used for the computation.

The factor IMPORT is a proxy for the importance of a particular type of part to the missions that can be flown by whichever types of aircraft use the part.¹ When $SHOPRY = 1$, IMPORT is simply the number of mission types for which the part is essential, divided by the total number of mission types that can be assigned to the aircraft types that use the part. When $SHOPRY = 2$, the proxy IMPORT is an estimate of the number of critical "holes" that would be generated in the average CIRF-BASE transportation time, at the current sortie demand rate.

The effect of this prioritization scheme is to work on the parts that have created the largest number of "holes," with the part with the larger IMPORT getting priority if the numbers of HOLES are the same for two part types. For parts that have not yet generated any HOLES, the part most likely to cause a HOLE (i.e., least value of $-VALUE \times MTBF$) gets the best priority.

¹The PRTCRT array is generated during initialization: For each part, each entry in this array contains, in packed form, a record of which aircraft types use that type of part, and which of that aircraft's missions require that part. The relative importance of a particular part, as used above, is defined as the sum of (number of mission types for which the part is required) divided by (number of mission types that can be flown) for the several types of aircraft using the part.

Since the prioritization for the parts will eventually become out-of-date due to the unpredictability of failures and repairs, all parts and equipments in the wait queue at each shop at the CIRF (as well as at the operating bases) are reprioritized twice a day, at 5:30 AM and 17:30 PM, starting on day 3. Furthermore, for those parts for which it is known that the required number of people, or the needed equipment, are not on base, RANK is set to 32750. And if the repair requires an AIS tray that is not functioning, RANK is set to 32600. Then, as repairs are checked, the search through the queue is stopped if the RANK is as large as 32600, or after 100 parts have been checked if the RANK is as great as 1000.

6. USE OF THE THEATER PARTS REPAIR FACILITY AS A BACKUP

The concept of a CIRF as described previously is an integral part of the spare parts logistics structure for the theater. The algorithms (Sec. VI) that are used to generate base stocks of spare parts distribute the appropriate numbers of LRUs to the operating bases. However, the SRUs are distributed both to the CIRF and operating bases, taking into account the proportions of the LRUs that will be repaired onbase and at the CIRF (as implied by their respective NRTS rates). In addition, these algorithms place the numbers of LRUs and SRUs in the CONUS-theater and intratheater pipelines, as would be expected on the basis of the stipulated peacetime flying program and the user-specified logistics structure.

Another role for an off-base parts repair facility in the theater could be as a backup, to handle repairs when a base has lost its planned capabilities, or did not receive key equipments scheduled for shipment from CONUS. In addition to the men and specialized equipments required for such an operation, some quantities of SRUs would also be required to permit the LRUs that are sent back from the operating bases to be repaired. But in these circumstances, TSAR's automatic parts generation (and distribution) algorithms would not provide SRUs for the backup facility. Stocks, in addition to these "procured" for the operating bases could be provided, of course, using CT23. Another option is made available by initializing the control variable TSAR to 3. When this is done, all operations are the same as for TSAR = 2, except that provisions are made to permit the repair facility to acquire SRUs from operating bases. For situations in which an operating base ships an LRU to the facility for repair, and knows which SRU is faulty, the appropriate SRU is added to the shipment. In other

circumstances in which an SRU is required at the facility, and none are in stock, enroute, or being repaired at the facility, the theater manager has one shipped from the base best able to supply it.

XII. DATA INPUT

The first step of the input process is to define the dimensions of the storage arrays and to zero out their storage locations; this is the primary function of subroutines INIT, INITO, and INIT1. Subroutine INIT also explains how TSAR may be redimensioned to tailor it to the programmer's special requirements simply by changing the appropriate values in the several PARAMETER statements. Subroutine INIT also includes copies of the 33 primary sets of labeled COMMON, a list of the arrays that are found in COMMON, and data clarifying which array dimensions may be modified.

The second step in the input process is to read the input data provided by CT1 through CT49, using subroutine INPUT, INPUTA, INPUTB, INPUTC, and INPUTD. Base-specific data stored in individual datasets are entered using subroutine BEDOWN. The definitions, formats, and procedures for entering these data are outlined at length in Sec. XIX in Vol. II. The user has considerable latitude as to what is to be included; many portions of TSAR may be inactivated simply by omitting a card or by providing a null entry for certain data.

This input process has many built-in checks; but because not all possible user errors have been anticipated, the user should adhere precisely to the instructions. Most input cards are screened by subroutine TESTER, which will catch a variety of common errors; the meaning of the errors caught by TESTER must be inferred by reference to the source code for that subroutine.

Subroutine INPUT calls on subroutine INPUTD to read airbase attack data and airbase damage data and to organize the attack times in a heap. The INPUTD subroutine is designed so that these data may either be input directly with the TSAR data deck or read from disk, where they have been stored by the companion model TSARINA, which computes the required damage and chemical contamination data from a description of the attacks and the location of resources among the various airbase facilities.

In addition to simply storing data, subroutine INPUT, assisted by subroutines REVIEW, AUDIT, and WRAPUP, carries out many data organization and initialization tasks, and performs additional tests on the completeness and internal consistency of the data. The initialization process also arranges resource shipments from CONUS in a time-ordered queue and computes the entries for the SHPTSK (shop tasks) array. Any

errors detected that are considered "fatal" for execution are explained by an error message that begins with an exclamation mark "!" and execution is terminated after initializing; this permits such messages to be readily located with a system editor, since this usage is unique. Anomalies in the data that would rarely be introduced intentionally, but are not considered fatal, are denoted with the word "warning" in the error messages.

The user has two options for obtaining a record of the input data: To simply echo input data as it is entered, the user places a "1" in column N of the first input card, if Card Type #N is to be listed; if the Type #N Card has subtypes, all are listed. The other option lists the data after they are stored and after the various special initialization actions have been carried out; this option is requested with the special card that precedes the sortie demand data; again a "1" is placed in column N if the data read with Card Type #N is to be reproduced (this option is not functional for all Card Types). The subroutine INLIST and the support routines LIST1, LIST2, LIST3, LIST4, and LIST5 respond to these demands. The user should note that the data are printed directly from storage and that they frequently have been modified or "packed" differently than when they were input. The definitions provided for each array in App. C of Vol. III will permit the user to interpret these listings.

The last steps in the input procedure are managed with subroutines INITIZ and TRIALS. The pointers identifying the available space in the several dynamic storage arrays are initialized in INITIZ. The last step in subroutine INITIZ is to list the status of personnel substitutability at each airbase.

Initialization is completed by subroutine TRIALS. Subroutine ICHECK is called first to initialize the PARTRQ array, provision battle damage spares, estimate the MRBF for each part, and generate a record of the shops that borrow personnel and equipment from other shops. Then subroutine RREQTS is called to compute the expected requirements for personnel, equipment, munitions, and shop facilities for each type of mission and each type of aircraft when the control variable STATE has been initialized to a value greater than zero. These estimates are used subsequently in subroutine BASCAP to provide daily projections of each base's sortie generation capabilities.

When the user has elected to let TSAR initialize the parts data and the parts pipeline to the depot, as outlined in Sec. VI.1, subroutine COMPRT (compute parts) is also called by TRIALS. When this option is chosen, the user must first have stipulated certain base characteristics and the NRTS policies for each part and each type of base

(using special versions of CT23). Subroutine COMPRT manages subroutines IPARTS and IPART2, which compute the total numbers of each type of part for each type of base; POS plus BLSS are derived for in-place units and a WRSK kit is created for units that are deployed to the theater. Listings of the results of these computations and of the pipeline contents are controlled by PPRINT.

Subroutine AVGTME (average times) is called to compute the average time that each aircraft maintenance shop can be expected to spend on on-equipment tasks and off-equipment repair jobs for each type of aircraft, when base resources are unlimited. These estimates take into account the likelihood that the different tasks will arise, parts will be broken, and the parts will not be condemned or shipped to another base for repair. When the control variable STATE has been initialized to a value greater than zero, subroutine BASCAP is called from TRIALS to generate the initial forecast of each base's sortie generation capabilities. These approximate sortie projections are derived by comparing the average resource demands for each type of mission and each type of aircraft with the available quantities of those resources at each base, as outlined at the beginning of this section. These forecasts are subsequently updated each evening at 1930 and are used with a variety of algorithms concerned with managing aircraft assignments and transferring aircraft from base to base.

If theater resource reports are to be transmitted during the simulation, TRIALS next calls subroutine STATUS to initialize the theater manager's data base with up-to-date and complete information regarding the resources that will be managed. The intratheater shipping schedule queue is organized next.

The next step in TRIALS is to input the initial set of sortie demand data. This is done by calling the entry point DAYONE in subroutine READFT (read flight data). As explained at greater length in Sec. VII.1, these data can be replaced or modified each day at 1945 or, if the flights are periodic, they may be used to control the demand for sorties for several days or throughout the simulation. Finally, the heap in the array PERIOD (periodic and scheduled tasks) is initialized in subroutine MANAG.

The input procedure up to this point has been primarily concerned with acquiring, checking, and manipulating the data that describe the various tasks and the initial resource levels and schedules, and with initializing various queues and heaps. The initial status of the aircraft and the maintenance shops is established with CT41 and CT42; and when parts are initialized automatically, NMCS aircraft may be generated to provide

sufficient parts to stock the pipelines. If only CT41 is used, it is presumed that for the situation being simulated there has been an opportunity to work off all unscheduled maintenance tasks except for NMCS aircraft, and to upload the aircraft for the designated types of missions at the beginning of the simulation. Similarly, the parts stockage generation option presumes that all on-base parts are serviceable. Thus, the various shops are inactive and no jobs have been interrupted or are waiting.

To reflect a situation in which aircraft maintenance tasks remain to be completed and various parts are being repaired or are waiting to be repaired, subroutine ZSHOPS is called by subroutine TRIALS. This modification of the initial conditions is controlled by the several CT42 cards, where the aircraft maintenance that is outstanding may be expressed by a three-part distribution for each type of aircraft at each base. Thus, one might specify that 20 percent of aircraft Type #2 at Base #3 has two tasks outstanding, 30 percent have three tasks, and 50 percent have five tasks. Subroutine ZSHOPS selects the required tasks at random, consistent with their nominal probability of occurrence, and computes the time remaining as a random fraction of the normal task time. If parts repair jobs are specified on CT42/1, the appropriate numbers are selected and placed in the administrative delay queue or in an in-process status by an equivalent random process. Other CT42 cards may be used to specify particular numbers of aircraft that require particular maintenance tasks, or particular types of parts and equipments that are undergoing repair.

The default air crew status (established in subroutine INPUTA) is that all air crews will be available for assignment at any time after 0015 on the first day.

When all phases of the initialization process are completed the simulation begins, unless program execution is terminated because VERIFY was set to 2 by the user or by the program as a result of fatal input errors that TSAR detected.

XIII. SIMULATION CONTROL

The MAIN routine initiates and concludes the simulation but delegates the control of the three main phases—input, simulation, and output—to subordinate routines. Input has been discussed, and printout of the final results is controlled by subroutine OUTPUT. Control for the simulation proper is passed first to subroutine TRIALS, which is responsible for the last portions of the initialization process and for running the simulation the designated number of trials. TRIALS manages the storage of the initial conditions for the first trial, for regenerating zero-time shop activities, and, when spares stocks are computed internally, for recomputing the initial spares for each trial. Control is passed by TRIALS to subroutine MANAGE, which exercises primary control throughout each trial of the simulation.

The basic task performed by subroutine MANAGE is to examine the earliest event that will occur in each of ten separate groups of events and to determine which of these ten is to occur next. Simulated time is then advanced to that time, and control is transferred to the appropriate subroutine for processing that event.

If the next event in each of two or more of these ten groups of events is to occur at the same time, the first event examined is processed first. The order in which the groups are examined is:

1. Completion of civil engineering reconstruction jobs
2. Completion of on-equipment aircraft maintenance tasks
3. Completion of parts repair and equipment repair jobs
4. Periodic and user-scheduled events
5. Completion of aircraft delays (sorties and preflight delay)
6. Initiation of aircraft maintenance at end of postflight delays
7. Aircraft sortie launch events
8. Completion of munitions assembly jobs
9. Completion of personnel rest periods
10. Arrival of resupply shipments

This order has been established primarily with a view to minimizing unnecessary processing. Thus, shop reconstruction is checked before maintenance personnel are released, so that parts repairs that are awaiting initiation would not need to be checked twice. And on-equipment tasks and aircraft delays are completed before flights are checked so that aircraft that are becoming available for launch are so designated at launch time.

When USECW = 0 (i.e., the chemical warfare features are not in use), control is transferred to BSEREP, RUNAC, RUNSHIP, and DOBILD for the first, second, third, and eighth groups, respectively; when the CW features are being used, control is initially transferred to STOPIT, where the effects of the special ensembles and chemical contamination on the personnel are checked before control is transferred to the same four subroutines to complete the action. Control is transferred to FLIGHT and DOSHIP for the seventh and tenth groups and to STARTM when the postflight delay is completed; for the fourth and fifth groups the nature of the event or delay determines which subroutine takes control; subroutine MANAGE transfers control to the appropriate location. Control is transferred to LETGO to release personnel that have had to rest and to check where they may be needed.

Many of the user-defined management control variables may be changed during the simulation. The timing for such changes is specified in the input data using CT49 cards or may be selected endogenously, thus providing a form of dynamic adaptive control. Subroutines ADAPT and MODIFY provide for the management of the user-supplied logic that controls such adaptive behavior.

When processing has been completed by the subordinate routine(s), control is returned to subroutine MANAGE and the next earliest event is selected. The entire simulation proceeds in this manner until the user-stipulated simulation length is exceeded, at which time MANAGE returns control to TRIALS to print the trial results and to initiate the next trial, or to print the overall results of the several trials.

The MODIFY subroutine is used to change the value of various TSAR control variables in response to endogenous and exogenous requests for change; the capability for generating endogenous changes in various factors provides a primitive form of adaptive behavior. Although no current use is made of this potential in the TSAR code, the structure is fully available in subroutines ADAPT and MODIFY. To date, the

primary use of subroutine MODIFY has been for exogenous changes at specified times; the CT49 cards currently provide the mechanism for manipulating many of the variables in this way, and many more could easily be added as desired. The change mechanisms that currently are available with CT49 are listed in Table 8.

The standard CT49 format is composed of several groups of three five-column fields; the day and hour for the change that is to be made is entered in the first field. The entry in the second field combines the Type Number for the change to be made, and part of the description of that change; the third field provides the remainder of the change description. The data in the second field can be thought of as being composed of TYPE \times 100 plus NUM, and the data in the third field (labeled VALUE on the card format) is frequently interpreted as $V1 \times 1000$ plus V2. (In those cases where only a single value is called for, it should be right-justified in the VALUE field.) In the explanations given below of the various types of changes that are currently available in subroutine MODIFY, we will refer to the several variables TYPE, NUM, VALUE, V1, and V2.

All values entered with CT49 should be entered in exactly the same units as are specified at the normal location for entering the value for these factors; i.e., in hours, in minutes, in hours and minutes, in TTU, or in whatever units are normally appropriate.

Space remains to define change in Types #10, #13 through #18, and #32 through #40; furthermore, additional space could easily be added to accommodate more user-designed types of changes.

Table 8

CHANGES AVAILABLE FOR TSAR CONTROL VARIABLES

Type	Num	Value		Description of Change
		(V1 Task Type	V2)	
1	Base		V2	V2 is used as a multiplier of the original value of the HURRY factor for this base and generic task type (for the five generic tasks)
2	0	V1	V2	V1 is the TEST1 value of TEST used for the debug time windows, and V2 is the number of the trial
	1		TESTAC	Defines the number of the aircraft to trace with special activity tests
	20 + x		SPAREX	Changes the value of SPAREX
3	1,7		START	Defines the beginning of the seven debug time windows (TTU)
	8,14		STOP	Defines the end for the seven debug time windows (TTU)
4	1		PRINT	Defines new values for key output control variables
	2		PPRINT	
	3		CPRINT	
	4		RPRINT	
	5		APRINT	
	6		DPRINT	
	7		DOU TL	
	8		Spare	
5	1		SLEEP	Defines new values for key time-related control variables. Enter the changes in the same units that were called for originally
	2		REST	
	3		ENDAY	
	4		LOADTM	
	5		LSTTOD	
6	1		STATFQ	Defines new values for specified control variables
	2		CANMOD	
	3		EXPED	
	4		FILLAC	
7	1		JOBCON	Defines new values for specified control variables
	2		MNTLMT	
	3		DOPHAS	
	4		ASSIST	
8			CMODE	Defines new values for CTHEA, CCIRF, and CIRFLG (see Sec. XI)
9	1		AUTHPC	Redefines the "authorized" collapse probability
	2		AUHT	Redefines the limiting rectal temperature
	3		DELTA	Redefines DELTA in degrees Centigrade ($\times 100$)

Table 8—continued

Type	Num	(V1	Value V2)	Description of Change
11			SCROLL	Since SCROLL defines the <i>last</i> day to display the special "scroll" output, this change can be used to start the "scroll" when the change is effective and to then stop it after SCROLL days
12	NUM		SCROL1	Defines the number of the first aircraft to be listed in the "scroll" (SCROL1), and the number of aircraft to list (NUM)
19	BASE		WX#	Changes the meteorological conditions (CWATTK(5,BASE)) to WX#
20	CRIT		Task Number	Defines a new value of task criticality for an on-equipment task
21	NUM		Task Number	The task time is changed by NUM TTU (the sign of the change is the sign of the task number)
22	NUM		Task Number	The probability of the task is changed by NUM percent (the sign of the change is the sign of the task)
23	NUM		Task Number	The first team size is changed to NUM personnel
24	NUM		Task Number	The second team size is changed to NUM personnel
25	BASE	1	V2	Changes the number of flight surfaces to be considered for a MOS to V2
	BASE	2	V2	Changes the required length of the MOS to 100 × V2 feet
	BASE	3	V2	Changes the required width of the MOS to V2 feet
26	BASE	1	V2	Changes the minimum postattack landing/takeoff delay to V2
	BASE	2	V2	Changes the special maintenance delay to V2
	BASE	3	V2	Changes the facility reconstruction delay to V2
	BASE	4	V2	Changes the delay prior to RRR to V2
	NA	5	V2	Changes the value of CEDELY to V2
	NA	6	V2	Changes the value of SHPDLY to V2

Table 8—continued

Type	Num	Value		Description of Change
		(V1	V2)	
27	BASE	0	V2	All values of EXTRAK are changed to V2 at BASE
	BASE	1	V2	Changes the value of EXTRAK for flight-line personnel to V2 at BASE
	BASE	2	V2	Changes the value of EXTRAK for preflight personnel to V2
	BASE	3	V2	Changes the value of EXTRAK for backshop personnel to V2
	BASE	4	V2	Changes the value of EXTRAK for munitions assembly personnel to V2
	BASE	5	V2	Changes the value of EXTRAK for civil engineering to V2
	BASE	6	V2	Changes the value of EXTRAK for off-duty to V2
28	BASE		V2	Changes the value of the special aircraft decontamination switch [CWATTK (13,-)] to V2
29	1	ACTYPE	V2	Changes the postflight delay time to V2
	2	ACTYPE	V2	Changes the preflight delay time to V2
	3	ACTYPE	V2	Changes the number of missions that an aircraft may be assigned to V2
	4	ACTYPE	V2	Changes the administrative delay for transferred aircraft to V2
30	TYPE	1	V2	Changes the time for civil engineering procedure #Type to V2 minutes
	TYPE	2	V2	Changes the time function for civil engineering procedure 4Type to V2
	TYPE	3	V2	Changes the alternate procedure for civil engineering procedure #Type to V2
31	ACTYPE x 10 + Number		B1 x 100 + B2	Permits change in an aircraft transfer directive between base B1 and B2; useful for zeroing demand, or changing to 9 or less

XIV. SUPPORT SERVICES

Twenty-four subroutines and 11 minor subroutines and functions support the main simulation. Each performs one or more specific functions and many are called upon from a variety of different locations. The functioning of each of these support routines is described at least briefly in this section. These discussions are ordered alphabetically; the subroutines discussed include:

ASSAY	Evaluates the fatal and nonfatal casualties
BREAK	Computes breakrate modifiers for on-equipment tasks when VBREAK is unity
CALCLO	Establishes the percentage of fatal and nonfatal casualties
CHECK	Checks on outstanding demands for newly available resources that have been released from aircraft maintenance tasks and parts repair jobs
CKTEMP	Evaluates temperature rise and fall for work crews in a CW environment
CWMOPP	Determines the required MOPP
CWSHIFT	Assists SHIFT in checking for heat or toxic casualties
CWTIME	Estimates permissible task time in CW ensemble
ENDAC	Eliminates all records associated with an on-base aircraft
FTIME	Generates time requirements for civil engineering jobs
GOREST	Manages disposition of personnel in a CW environment when they are released
HEAP	Enters and removes items from a heap
INTRUP	Enters and removes time-ordered items from the interrupted task array
KILLAC	Eliminates all records for an aircraft that is lost in combat
LOSSES	Determines the specific number of items lost
NORRPT	Enters and removes records of aircraft "holes" from the NORQ array
REDPEO	Reduces or increases the number of ground personnel and reorganizes the shift structure
RESET	Resets simulation time for a continuous simulation of $NTRIAL \times SIMLTH$ days duration when $EXTEND = 1$
SHIFT	Adjusts the size and activity of the work force when shifts are changed

STRTSK Stores and retrieves required and deferred on-equipment tasks

TRIAGE Determines overall casualty rate and prorates between conventional and chemical effects

TTIME Generates time requirements for aircraft maintenance and theater communication delays

UPDATE Updates the chemical contamination of each monitoring point and determines the required MOPP and equilibrium temperature

WAIT Enters and removes time-ordered items from the waiting- task array

The other 11 service items are summarized briefly at the end of this section.

SUBROUTINE ASSAY

This subroutine is an entry point in subroutine TRIAGE. For each group of personnel subject to casualties from the same conventional and/or chemical weapons effects, this subroutine determines the number of fatalities, the number of nonfatal casualties attributed to conventional weapons, and the number of nonfatal casualties attributed to chemical weapons effects. This is done by using a Monte Carlo procedure applied to the output of TRIAGE for the total percentage of fatalities, and the percentages of nonfatal casualties attributed to conventional and to chemical weapons effects. The classification of each nonfatal casualty as due either to conventional or to chemical effects (but not to both) is done so that the individual hospitalization times can be determined by sampling from the appropriate hospitalization time distribution, which is entered with CT43/6.

SUBROUTINE BREAK

This subroutine is used when VBREAK is initialized to unity to modify the probabilities with which on-equipment tasks are required as a function of achieved sortie rate. Called at the end of each day by subroutine OUTPUT, this subroutine computes the sortie rate achieved during the preceding day for each type of aircraft; the estimate is given by the total sorties flown by each aircraft type and the total number of such aircraft surviving in the theater. The appropriate breakrate modifier is then computed separately

for each shop and each aircraft type on the assumption that the nominal breakrate applies for a sortie rate of one per day, and that the actual breakrate falls linearly for each additional sortie per day by the percentage specified with CT44. The resultant value is stored in the second element of the TSKPR array.

SUBROUTINE CALCLO

This subroutine is called by subroutines GOREST (when a job is stopped, both at attack times and between attacks), LETGO (when personnel have completed a rest period), COOLOS (at the time of an attack, for personnel who are resting), and SHIFT (at shift change time). It sets up the required arguments and then determines the percentage fatalities (from combined conventional and chemical weapon effects), the percentage hospitalized from conventional weapon effects, and the percentage hospitalized from chemical weapon effects, for groups of personnel subject to the same weapon effects. At the time of an attack, it determines the percentage of conventional losses (from TSARINA CT40 input), establishes the percentages of fatalities and nonfatal casualties by a call to subroutine CWLOSS, and then calls subroutine TRIAGE to establish the total percentage fatalities (from combined conventional and chemical weapon effects), the percentage of nonfatal casualties attributed to conventional weapons, and the percentage of nonfatal casualties attributed to chemical weapons. Between attacks, calls to this subroutine return the percentage fatalities and the percentage nonfatal casualties to chemical weapons effects during the period of time specified as an argument—e.g., the time spent working on a task before resting.

SUBROUTINE CHECK

This subroutine is used to check on resource demands that may be outstanding and is called whenever resources are released from a previous event or are delivered to a base. The five sources for such demands are interrupted, waiting, and deferred on-equipment aircraft maintenance tasks and interrupted and waiting parts repair jobs. To reduce processing somewhat, the call to this subroutine may specify a shop number, an aircraft, a part, a type of personnel, a type of equipment, or any combination. In the subsequent search among outstanding demands, no attempt is made to initiate tasks that do not require the resource specified.

The search is ordered to examine on-equipment tasks before parts repair jobs; in each case, interrupted items are examined before those that are waiting. At night (after ENDDAY), deferred tasks are checked after repairs have been checked. All five queues are searched when a shop or a type of ground personnel is specified. If an equipment type is specified, only on-equipment tasks that are waiting (and, at night, deferred) are checked. When an aircraft is specified, only on-equipment tasks are examined. When an LRU is specified, the aircraft waiting for that part are examined to select the one with the fewest holes, and if two or more have the same number of holes, the aircraft with the earliest projected ready-to-fly time is selected. When the part specified is an LRU, only jobs waiting for repair are examined.

For the shops that may lend personnel or equipment to other shops, subroutine CHECK checks the shops that are listed in a TSAR-generated list of borrowing shops to see whether the newly released personnel or equipment are needed for either on-equipment or off-equipment jobs. If so, the resource is lent to the shop that requires it. These checks are made for all on-equipment shop tasks before the off-equipment jobs are examined.

SUBROUTINE CKTEMP

This subroutine estimates how much the rectal temperature of an average (70 kg) man will rise or fall as a function of his environment; the estimates are based on an adaptation of the Goldman heat stress model^[11]. When called to work, a work crew's initial temperature is assumed to equal (approximately)¹ the equilibrium temperature at the work place, and to increase in a manner dependent on their chemical ensemble, the strenuousness of the task, and the ambient conditions. Based on a user-stipulated allowable temperature (or an equivalent allowable collapse probability), this subroutine determines the length of time that the crew may work. Subroutine CKTEMP is also called whenever a task is stopped to estimate how long the crew must rest for their temperature to fall to the prescribed level; it is also called every two hours to compute and store the equilibrium temperature for personnel, appropriately protected, in all facilities.

¹See Sec. IX for details.

FUNCTION CWMOPP

This function subroutine determines the MOPP required at a facility (building, shelter, taxiway or runway arc, or ramp) at any time subsequent to a chemical attack (prior to a chemical attack, the required MOPP is determined from the MOPPOL array; see subroutine UPDATE). The required MOPP depends upon the facility chemical protection category (CWTYPE), the residual chemical threat at the MP of the facility (or, when VARMOP = 0 on CT3/4, at the MP with the highest chemical threat for the CWTYPE of the given facility), and the MOPPs specified in the required MOPP (RQDMOP, CT44/4) array. The RQDMOP array contains the intensity thresholds for donning different MOPPs, and the required MOPP is the highest numbered MOPP needed for the given surface depositions and vapor concentrations (for up to three chemical agents), taking into account the attenuation afforded by the CWTYPE for the facility (CWPROT, CT44/1).

SUBROUTINE CWSHIFT

This subroutine is called by subroutine SHIFT when USECW > 0. All tasks assigned to shops that are then changing shift are checked by calling subroutines STOPIT and GOREST to see if any of the crew are casualties because of toxic effects or from heat prostration; if not, the tasks will be continued during the following shift without interruption, if the size of the work force permits. If any of the crew are casualties the job is interrupted and the residual personnel are added to those available (because they are going off duty and will cool off as required while off duty).

Subroutine CWSHIFT then calls subroutine REDPEO to readjust the assignments of personnel to the various squadron and wing organizations, whenever any personnel have been lost or gained during the preceding shift. This is necessary because these readjustments are not made at the time of the change, as is done when USECW = 0, but are done only at shift time; this difference was introduced to avoid the numerous readjustments that are to be expected in a chemical environment.

SUBROUTINE CWTIME

This subroutine is called by subroutine TTIME to estimate how much longer a task will take when personnel are wearing encumbering chemical ensembles, and to check how long the crew will be able to work before the expected rectal temperature

will exceed the user-specified limit. After checking what part of the ensemble must be worn at the location where the work is to be carried out, the total task time can be determined with the user-supplied task degradation data (see CT43/3). Using subroutine CKTEMP, an estimate then is made as to how long the crew can work taking into account the difficulty of the task and the heating and cooling effects of the ambient temperature, humidity, and wind velocity; the limitations that could be imposed by excessive wetting or dehydration are also checked using subroutine DEHYDR, and the most restrictive criteria determine the time that the work crew is permitted to work.

SUBROUTINE ENDAC

This subroutine is used only when an aircraft has been damaged or destroyed by an enemy airbase attack. It is called from subroutine ATTKAC after that subroutine has appropriately decremented the personnel, equipment, and parts associated with the aircraft at the time of the attack.

For damaged aircraft ENDAC is used to eliminate any flight assignments; if an aircraft has been destroyed, ENDAC then removes it from the aircraft delay queue (when appropriate) and removes all references to required, interrupted, or waiting tasks. All ongoing tasks are then stopped and the surviving personnel and AGE are released for other jobs. No times are recorded for these tasks. The last step is to call subroutine KILLAC in order to erase any deferred task records, to eliminate the aircraft from the base inventory, and to order a replacement aircraft, as appropriate.

FUNCTION FTIME

This special function provides the user substantial flexibility for specifying how the required time for civil engineering jobs vary for different types of jobs and different levels of damage. In basic terms, the formulation consists of a delay, or startup time, plus a damage-dependent reconstruction time. For each type of civil engineering facility repair job, the user specifies the time (t) required to repair a "unit of damage" and indicates how the total time (T) will vary with the total number of units of damage (D) by entering a coded number (C) for the functional relationship. This subroutine uses those data to estimate total time as follows:

$$T = \text{Delay} + t \times D^b$$

where

$$\text{Delay} = f(B)$$

$$b = g(P)$$

and, since $C = 12 \times P + (B - 1)$,

$$P = C/12 \quad \text{the largest integer multiple of 12 in } C$$

$$B = 1 + C - 12 \times P$$

The data tabled in FTIME for f and g provide 12 values for the delay (0,1,2,3,4,6,8,12,18,24,36, and 48 hours) and seven values for b (.5,.75,.9,1.0,1.1,1.25, and 1.5). To specify a time proportional to the total damage without any initial delay, C would be 48—i.e., $P = 4$, $B = 1$, so that $b = 1.0$ and $\text{Delay} = 0$.


SUBROUTINE GOREST

This subroutine is called from subroutine STOPIT when a job is stopped (and other subroutines when a munitions load team is released) to determine the number of casualties from airbase attack weapon effects (by a call to subroutine CALCLO), the number of personnel collapsing from excessive heat buildup, and the number of personnel requiring rest. The personnel who are nonfatal casualties to weapon effects or collapse from heat prostration are placed in the "clinic" (by calls to subroutine CLINIC) for a period of time determined by sampling from the appropriate hospitalization time distributions entered with CT43/6. The personnel requiring rest either rest at the work location or go to collective-protection shelters according to the value entered for USECP on CT3/4 (see Sec. IX.6). For personnel going to protective shelters to rest, this subroutine picks an appropriate shelter and handles the entry times and any queue formed at the shelter.

SUBROUTINE HEAP

When it is not necessary that timed events be ordered, but only that the earliest event be readily located, a data collection that has been called a "heap" permits more efficient processing. On the average only two positions need be interchanged when a new event is entered into a heap.

The heap illustrated in the following sketch has 11 entries; the smallest "time" in the heap is 817 and the largest is 860. The smallest value is at the "top of the heap." Except that no entry is larger than any entry lower down on a branch that stems from it, there is no particular order maintained in a heap. The various entries are all stored in an array and are interconnected with a system of pointers. In the example, the next entry would be placed in position 12; if its value were less than 842, positions 6 and 12 would be interchanged, and then positions 3 and 6 would be interchanged. If the new entry were less than 817, positions 1 and 3 also would be interchanged, and the new entry would be at the "top of the heap."



This subroutine has four entry points, one to enter an item (INHEAP), one to remove the item with the lowest valued time (OUTHEP), a third to extract an item (EXHEAP) from within the heap, and a fourth to modify the time for an item in the heap (MODHEP). To extract an item from the heap or to modify an item, it is necessary to know which column it occupies in the parent array, if it is to be found readily; but when these actions are required before an event has become the one with the earliest time, that information logically is known.

The size of the calling array is a variable and the number of entries in that array may be fixed or variable. This subroutine operates on three rows of whatever storage array is specified in the calling statement. The time of the event is in the first row, a pointer to the event's position in the heap is in the second, and a pointer back to the event from its position in the heap is in the third row.

One peculiar property of this data structure should be noted: If several events are entered that have the identical time associated with each of them, they will not be removed in the same order in which they were entered.

SUBROUTINE INTRUP

Aircraft on-equipment tasks and parts repair jobs that have been interrupted are queued in the INTTSK array. Each shop on each base stores a pointer to the first and the last of its interrupted tasks and its interrupted repairs in the array SHOPS. Whenever resources are available to start an interrupted activity, the first item in these queues is the first to be examined. If the user wishes priority to be given to that item that has been in the queue the longest, the control variable ORDIT is initialized to zero and the queue is managed locally in the main routines.

If the user wishes to have the events ordered such that the one with the lowest value of a variable called TIME is first, ORDIT is initialized to unity, and subroutine INTRUP is called to manage the queue. The value of the variable TIME need not be a time per se, and, as discussed elsewhere, differing events are queued in accordance with differing definitions for TIME.

This subroutine has separate entry points for entering an item (ININT) and for removing an item (OUTINT). The code is written so that any item may be removed, not only the one that is first in the queue. Three rows of the INTTSK array are involved in queue management; the value of the variable TIME is stored in the seventh row, and pointers to later and to earlier items are in the fifth and sixth rows, respectively.

SUBROUTINE KILLAC

This subroutine is used whenever an aircraft is lost in combat, and it also completes the work begun in ENDAC when an aircraft is destroyed on base. The two basic functions performed by this subroutine are to erase any reference to tasks that may have been deferred on the aircraft and to eliminate the aircraft from the base inventory.

SUBROUTINE LOSSES

This subroutine generates the specific number of items that are lost when N items suffer a loss probability of P . If the control variable NONUNI (CT2/1) is zero, the value returned for one item is determined by comparing a random number with P ; for more than one item the value returned is that integer closest to the expected losses—i.e., $N \times P$. If the control variable NONUNI is unity, the value returned is a sample drawn from the binomial distribution (determined by comparing N random numbers with the value P).

SUBROUTINE NORRPT

Whenever a part has been removed from an aircraft and has not been replaced immediately, or whenever a part on an aircraft has been found to be defective but has not yet been removed, a record is made of the particular aircraft and part. The data on these reports, or "holes," are stored in array NORQ using subroutine NORRPT.

Whenever an entry is made in NORQ (using entry RPTNOR), this subroutine first adjusts the number of aircraft on the base that are missing a part and then adjusts the count of the "holes" in that aircraft. If the rules for intratheater resource transfer permit, an order may be issued (by a call to subroutine CONTRL) to ship a part of the required type to the base that reported the "hole." The discussion of subroutine CONTRL outlines the rules that are followed in different circumstances (see Sec. XI).

The last step is to place the aircraft number in the NORQ queue that contains the numbers of those aircraft assigned to that base and missing that part. The aircraft number is ordered in the queue by the amount of time remaining until the aircraft would have been ready to fly if the part had been available; the aircraft with the least time is first. For subroutine NORRPT to manage these queues, the array NORQ stores the aircraft number, the time remaining, and a pointer to the next report in the queue. The fourth element of the PARTS array—PARTS(PART, 4, BASE)—contains the pointer to the position in the NORQ array of the first aircraft that requires that type of part.

Whenever the aircraft "hole" has been filled, this subroutine is called through entry REDNOR to take the record out of the queue in NORQ. This is done after the tallies noted earlier are updated.

SUBROUTINE REDPEO

This subroutine reduces the number of ground personnel on a base when some are shipped to another base and reorganizes the number that remain after an airbase attack. Subroutine SHPRES calls in the first instance and subroutine BOMB in the second. Calls to this subroutine prescribe the type (PEOP) and the number (NUM) of personnel to be withdrawn; NUM = 0 when the survivors of an air attack are to be reallocated to the day and night shifts. Distinct procedures are used for aircraft maintenance personnel and for civil engineering personnel.

The first step in this subroutine is to identify whether personnel of the designated type are assigned to two or more on-base organizations (the ALTPEO array provides the necessary data on the equivalent types of personnel). If they are, the personnel are redistributed among the several organizations in the proportions implied by the "target" force levels. The next step is to establish what numbers will be on the day and night shifts after reorganization. The new shifts are sized in the same proportions as the "target" force levels, except that no shift is allowed to be smaller than the "minimum shift size" entered with CT17.

If some personnel at work during the present shift must be released, parts repairs are interrupted first; if more personnel are required, aircraft maintenance tasks are interrupted. If more people have been directed to be transferred than can be found, the number to be transferred is reduced accordingly; this situation can arise if personnel are being used in other than their "parent" shop (where they cannot be located readily).

The procedures for the civil engineering personnel are comparable except that they are all in a single organization and the choice of tasks to be interrupted is based on the facility priority list (CT39); personnel are released from the lowest priority task first. When a civil engineering task is interrupted, the work remaining to be done (i.e., the current damage level) is estimated on the assumption that the remaining work is the same fraction of the total job as the remaining time is of the total time. The quantities of unused building materials are estimated in the same manner and they are returned to stock.

SUBROUTINE RESET

When the control variable EXTEND (CT1) is initialized to unity, the simulation may be extended to an indefinite length and is not restricted to 65 days. This is done by resetting the various time values in the simulation data base at the end of each trial, but without reinitializing any of the resource status values; thus, the second trial is just an extension of the first and so on. This subroutine performs all the necessary time adjustments when called by subroutine TRIALS.

SUBROUTINE SHIFT

This subroutine is called at two-hour intervals by subroutine MANAGE and changes the size of the on-duty work force for the personnel assigned to whichever work centers (shops) have a shift change at that time. Both the day and the night shifts are assumed to be 12 hours in length. Shifts that begin between midnight and 10:00 AM, inclusive, are designated the "day" shift. Using CT18/1, the shift schedule is prescribed independently for each shop. The work schedules are the same on all bases for shops of like number; however, the number of personnel on the different shifts is controlled independently for the different bases using CT21.

The basic function of this subroutine is to check whether more people are currently engaged than will be available on the next shift, and if so, to interrupt a sufficient number of activities that the required number of personnel may be released, or, if more people will be on duty, to attempt to assign the extra personnel to interrupted or waiting activities. The complications arise from personnel that may (1) be allowed to work a specified amount of overtime if they can complete their task within that time, or if they are engaged on an aircraft that has been scheduled for a late takeoff; and (2) have been lent to another shop and will not be found when their parent shop is checked.

The first step taken when a shop changes shifts is to reset a flag and zero a counter in the sixth and seventh positions of the PEOPLE array. Then, if USECW > 0 (i.e., the CW features are being used), the on-duty personnel that are cooling off and are assigned to shops that are changing shifts are released from the COOLER array; they first are checked to see that they have not been affected by toxic effects while they were cooling off, and if not, they are added to the available personnel. It is assumed that they will complete any additional cooling that is required during their off-duty shift. Then, again when USECW > 0, subroutine CWSHIFT is called to check the condition of all

personnel currently at work and to establish the initial personnel conditions for the tasks that will be able to be continued into the next shift. Following that, subroutine SHIFT checks each parts repair job and each aircraft maintenance task for each shop that is changing shifts. At the first encounter with an as yet unchecked type of personnel, the net change in shift size is established (and the flag is set to one). If the new shift is sufficient to handle all ongoing repairs and aircraft tasks, the flag is set to two. If the follow-on shift cannot handle the current work load, parts repairs and on-equipment tasks are interrupted (in that order) until a sufficient number are released, at which point their flag is set to two. The most recently initiated parts repairs or on-equipment tasks are interrupted first. The counter in the seventh position in PEOPLE is used to maintain a record of the number of personnel that remain to be released.

If the personnel on a particular activity can finish their task within the allowed overtime period (for this decision it is presumed that the exact completion time is known), or if they are working on an aircraft that is scheduled for a late takeoff, they are allowed to continue; they are credited to the required reduction, and subtracted from the "available" personnel for the subsequent shift. Thus, at the beginning of a shift the number of personnel available can be a negative number equal to the number of personnel working overtime; as each group is released, the "available" personnel remains at zero or less until fewer than the designated number on the next shift remain assigned. (Note that overtime is not allowed when USECW > 0.)

When personnel have been lent to another shop that may have its shift change at a different time, the flag and the counter are still operative; when the various activities are checked and the "borrowed" personnel are noted, they will be released if their flag value is zero or one. Otherwise, the activity continues. In effect, members of the new shift take over for those on the previous shift.

To avoid overlooking personnel assigned to shops that have no activity underway at the time the shift is changed, ground personnel are next checked type by type, and the PEOPLE data are modified as appropriate, when their parent shop has a scheduled shift change and the personnel flag is still zero.

For shops that have had a net increase in work force (measured by the counter REM), subroutine CHECK is called to start any outstanding jobs.

The only exceptions to the preceding description occurs for the "flight line" shop—Shop #25—and the shops associated with the preflight tasks—reconfiguration,

weapons loading, and refueling (Shops #27, #28 and #29, respectively), and for civil engineering. Personnel attached to Shops #25 through #29 who must be released are required to complete their current task, without regard to allowable overtime, because such tasks tend to be fairly short and it seemed likely that such critical tasks would be completed in wartime. Civil engineers do not work overtime.

SUBROUTINE STRTSK

This subroutine manages the storage of unscheduled on-equipment maintenance tasks in the RQDTSK and DEFTSK arrays. At the time an aircraft lands and the unscheduled tasks are identified in subroutine PSTFLT, a tentative mission is selected for the next flight and the tasks are separated into those that are required and those that may be deferred. Separate entry points are provided to store (STTASK) and to remove (REMTSK) a task, and a flag in the calling statement identifies the array to which the task belongs.

Each array is used to maintain an ordered set of tasks for each aircraft; two pointers in the ACN array determine the positions in the RQDTSK and DEFTSK arrays where the first tasks are stored for each aircraft. The tasks are ordered as they are identified, and for the required tasks the sequential shop structure that is defined with CT29 is preserved by entering the minus value of the first task identified for each group of shops whose work may be pursued simultaneously. The end of each set of tasks for an aircraft is identified by a zero entry in the task number position.

SUBROUTINE TRIAGE

For each group of personnel subject to casualties, this subroutine determines the percentage of personnel lost to the combined effects of conventional and chemical weapons and prorates the percentage of nonfatal casualties (hospitalizations) between conventional and chemical weapons effects. It first determines the total percentage of personnel lost to conventional weapons effects by combining the losses to direct effects (as determined from TSARINA CT40 input) and the losses to indirect effects (for accidents, accidental exposure, etc., from CT39/99) to produce the total loss percentage to the conventional attack (the two sources of losses are treated as independent effects). Then the percentage of fatalities from conventional weapons effects is determined by multiplying this loss percentage by the fraction of conventional casualties that are

fatalities (from CT4/2); the remainder of the total loss percentage from conventional weapons is the nonfatal percentage loss from conventional weapons. Next, the total percentage fatalities from the attack is determined by combining the percentage fatalities from conventional weapons and the percentage fatalities to chemical weapons. Next, the total loss percentage is determined by combining the total loss percentages from conventional weapons and from chemical weapons, and the total nonfatal casualty percentage is determined by subtracting the combined fatal loss percentage from the total loss percentage. The combined nonfatal casualty percentage is then prorated as the nonfatal loss percentage to conventional weapons and the nonfatal loss percentage to chemical weapons in proportion to the individual nonfatal loss percentages for conventional and chemical weapons.

FUNCTION TTIME

This function selects the "true" time for a job on the basis of a mean task time and a time distribution that are specified in the calling statement. For both on-equipment and off-equipment aircraft maintenance tasks, the user is restricted to the use of nine distinct distributions; for intratheater transportation and communication delays, up to 15 distributions may be specified (i.e., six additional distributions).

Twenty-five data points are stored in the local DIST arrays to represent each distribution. Several log-normal and uniform distributions with different variance to mean ratios are available currently in TTIME in TSAR and these could be changed easily to satisfy special user requirements. These data are interpreted as 1000 times the ratio of the true value to the mean value. The entry selected is determined by the draw of a random number between 1 and 25. The true time value is returned in TSAR time units (multiples of three minutes).

The user may add delays or speedup factors to the true time calculation. The nominal task times generated in TTIME are modified by use of several control variables to represent such efforts to shorten and otherwise expedite jobs. If the mean time and the random variate are designated as TM and F, the actual task time is generated as:

$$\text{HURRY} \times F(\text{TM} - \text{REDUCE}) - \text{SAVE}$$

where the variables HURRY, REDUCE, and SAVE may be specified separately at each base for on-equipment tasks, preflight tasks, parts repair jobs, munitions assembly jobs, and civil engineering tasks (see CT17/2).

SUBROUTINE UPDATE

This subroutine is called at the beginning of the simulation, immediately after an attack, and periodically (every CWFREQ hours, subsequent to a chemical attack) to determine the currently required MOPP at each facility (in each building, in each shelter, on each taxiway, and on each ramp). It also establishes the personnel equilibrium temperature in the open, in shelters, and in each building. Prior to a chemical attack, the required MOPP is determined from the MOPPOL array (CT3/4) that contains the preattack MOPP for personnel in each of the five generic task type categories (on-equipment, preflight, backshop, munitions assembly, and civil engineer) and for off-duty personnel. At any time subsequent to a chemical attack, the currently required MOPP is determined by use of the function subroutine CWMOPP.

SUBROUTINE WAIT

This subroutine manages the on-equipment and off-equipment jobs that must wait and are stored in the WAITSK array, in the same manner that subroutine INTRUP manages interrupted activities. Each shop on each base has a pointer to the first and the last on-equipment task and to the first and the last parts repair job that is waiting for action by that shop.

This subroutine is used only when the user wishes to have activities ordered in their queues by the value of the parameter TIME; to be so ordered, the control variable ORDWT is initialized to unity. The items are ordered such that the one with the lowest value of TIME is first. And as with the INTRUP subroutine, the value for the TIME variable need not be a time per se; the specific definitions in use are explained in connection with the calling routines.

The mechanics of this subroutine are identical to those in INTRUP, except that the queues are maintained in the 7th, 8th, and 9th positions of the WAITSK array.

ADDITIONAL SERVICES

There are 11 additional services; five are used in conjunction with the INLIST subroutine for formatting the listing of input data (i.e., LIST1 through LIST5). Four of the services interpret TSAR time to provide the time of day in TSAR time units, or hours and minutes, and the day and the hour (TOD, HRMIN, DAY, and DATE). The last two functions control the time horizon for projecting aircraft supply and demand (function

THF) and the length of the time intervals used in that process (function TU). The user may specify these factors with entries on CT4/2, or use the encoded default values. As currently coded, the default values for the the time horizon are 8 hours between 4 AM and 4 PM, 12 hours from mid- night to 4 AM; 16 hours from 8 PM till midnight, and 20 hours from 4 PM to 8 PM; the 16 time intervals are defined by the function TU.

XV. OUTPUT

In a simulation that involves multiple trials and as wide a variety of activities as TSAR, a great abundance of data might be reported. The output options that are provided with TSAR permit the user to examine a substantial portion of what we judged to be the more relevant results, but all possible outputs certainly are not available. For the additional, more specialized kinds of results that some users may find necessary for their particular problems, custom additions should be appended at the time they are required. Special provisions have been included to assist the user in providing such additions. If all such output were included, the costs in time and dollars for storing the data, and the space for displaying them, would have to be borne by all users. Most of the output options currently available are illustrated in connection with the sample problem in Sec. XXI, Vol. II.

The current output options are controlled by the variables PRINT, STATFQ, CUMSTA, DOUTIL, CPRINT, PPRINT, RPRINT, APRINT, DPRINT, TPRINT, DOPOST, and SCROLL. There are also provisions for user-customized output—see Sec. XV.3. Printed input information that precedes the simulation results in output listings that are discussed in Sec. XII. (The various debugging statements and outputs controlled by the variable TEST will not be discussed in this section.) For the individual trials, PRINT (CT2/1) controls the data printed that relate to the number of flights and sorties flown and the numbers of maintenance tasks performed. It also controls the daily reports of work-rest times experienced in the generic tasks and the cumulative reports regarding personnel casualties, fatalities, and hospitalizations. STATFQ (CT2/1) controls the collection and display frequency of shop performance statistics, including statistical data on the resource constraints that cause on-aircraft maintenance delays and backshop repair delays; these statistical data may be obtained separately for each trial, or the results may be aggregated over all trials, depending upon whether CUMSTA (CT2/1) is 0 or 1. DOUTIL (CT2/5) controls the collection and display frequency of activity data for all types of on-base personnel (except air crews) and a detailed manhour report for each type of personnel; the average percentage of available personnel of each type are listed for every odd-numbered hour each DOUTIL days. Furthermore, when DOUTIL is initialized, the manhours expended are accumulated for each type of personnel and are

listed at the end of each trial (when PRINT \geq 8) and at the end of the NTRIALS. PPRINT (CT3/3) controls the display of the numbers of serviceable, repairable, and enroute spare parts at each base, both at initialization and whenever shop statistics are printed, while CPRINT (CT3/4) controls the display of chemical contamination information, and other special chemical data. RPRINT (CT2/5) controls the listing of several special runway and taxiway repair activities, APRINT controls other attack-related results, DPRINT (CT2/5) controls the extent of output provided with the special summary of current aircraft status and deferred aircraft tasks that is activated with the CT2/4, and TPRINT controls the listing of commodity arrival reports. DOPOST initiates disk storage of selected results for postprocessing—see Sec. XV.4.

SCROLL (CT2/1) provides the user an opportunity to observe the behavior of several aircraft in some detail. When SCROLL is used, a record of the daily activities for each of up to NSCROL (CT2/1) consecutively numbered aircraft is listed at the end of each day for the number of days specified by the value of SCROLL. The number of the first aircraft is #1, unless otherwise specified, and the results for NSCROL aircraft will be listed, unless a smaller number is specified. The four numbers listed immediately following each aircraft number are: (1) the number of sorties initiated that day, (2) the number of the base the aircraft is assigned to, (3) a coded number summarizing the aircraft's maintenance status, and (4) the number of "holes" in the aircraft; these data are current as of midnight. Following these data, the times for the beginning and end of each flight, for each on-equipment task completed during the day, and for other special activities are listed, along with a description of the completed activities.

In addition to these various data that may be obtained for each trial, the final results also include a day-by-day record of the average number of sorties flown, and the standard deviation thereof, for each mission and for each base, when more than one trial is run. Other results printed for multiple trials include the average numbers of maintenance personnel that were available, the numbers of aircrews and ground personnel that were hospitalized or were fatalities, the numbers of manhours lost because of the requirement to cool off and because of hospitalization, and the aggregate numbers of equipment, spares, and munitions lost during air attacks. Multiple trial results also include the averages, by day, of the numbers of aircraft that are possessed (overall and by base), the aircraft that have been lost overall, the aircraft that have been damaged (overall), the aircraft still damaged (by base), the aircraft that are NMCS (overall and by base), and the cumulative totals of the NMCS aircraft-hours (overall and by base).

1. OUTPUT CONTROLLED BY THE VARIABLE PRINT

The data provided for each trial for a particular value of the variable PRINT include all items down to and including those listed for that value in Table 9.

2. OUTPUT CONTROLLED BY THE VARIABLE STATFQ

When STATFQ (CT2/1) is initialized to a value greater than zero, data on the duration of aircraft maintenance tasks, parts repair jobs, equipment repair jobs, and delays in initiating aircraft maintenance and parts repairs are stored using the subroutine TIMES. These data are printed at the end of each STATFQ days, at the end of each trial, and at the end of the simulation by calling subroutine DELAYS from subroutine OUTPUT. In each case, the results presented are based on the cumulative data to that point in the simulation if CUMSTA is 1; if CUMSTA is zero the results are cumulated independently for each trial. The results at the end of each trial also include the delay data for those activities that are still waiting at that time, on the assumption that all delays end at that time.

The first set of results presents the number of activities and the average length and standard deviation of the time that they required for on-equipment tasks, off-equipment jobs, and equipment repair jobs at each shop on each base.

The standard time, or resource unconstrained time, as calculated during the input process in subroutine AVGTME, is also listed for the on-equipment and off-equipment activities; the values computed in AVGTME for the various aircraft types are weighted in the output by the numbers of sorties flown by the various aircraft types at each base.

The second set of data provides a count of the ready aircraft that were canceled by a crew shortage and a count of the additional numbers of crews that would have been needed to satisfy the minimum flight requirements.

The last set of data provide a statistical summary of the causes and the duration of aircraft maintenance delays. For each base, for each of the other nine classes of resources, and for each individual resource type that caused an on-equipment task to be delayed, the results include the number of such delays and the average value and standard deviation of their duration. If any of the aircraft have "holes" at the time of the report, the number of holes is listed with the parts data for each base.

Data of the several types controlled by STATFQ are listed only when there are results to be reported; null data are suppressed.

Table 9

OUTPUT DATA CONTROLLED BY THE VARIABLE PRINT

PRINT	Output Data
-1	EOT: Storage array status if any overflows occur in one or more of the 18 dynamic storage arrays.
0	EOT: Cumulative flights and sorties flown, demanded, and the percentage of sorties flown of those demanded; totals for each base and each mission. ^a EOT: Fatalities, hospitalizations, and buddy care statistics. EOT: Cumulative on-equipment tasks, parts and equipment repairs by base and by shop. EOT: Cumulative hours NMCS at each base.
1	EOT: Sorties flown, demanded, and the percent of those demanded by base and mission, ordered by priority. EOT: Personnel work and rest times. EOT: Cumulative sorties canceled due to ATC constraints. EOT: Readiness indices ^b at each base. EOD: Aircraft possessed, lost, damaged, fillers, reserves, and transferred. EOD: Sorties and damaged aircraft by base. EOD: Daily reports listed for PRINT = 2. Report of the damage for each airbase attack. Number of tasks waiting by base and shop every six hours at operating bases.
2	EOD: Sorties flown, demanded, and the percent of those demanded that were flown by base and mission. EOD: On-equipment and off-equipment tasks completed during the day by base and by shop. EOD: Current supply of munitions and TRAP by type and base. EOD: Numbers of tasks and repairs being processed, and the numbers of repairs waiting by base and by shop (also listed at noon if PRINT = 3). EOD: Status of AIS and dynamic storage, and spares disposition. EOT: Remaining supplies of munitions and spares. EOD: Fatalities, hospitalizations, and buddy care statistics. EOD: Personnel work and rest times. Number of NMCS aircraft by base every six hours. Numbers of aircraft possessed, damaged, and with one or more "holes" by base, at three-hour intervals. Notice of the initiation of runway or taxiway repair. Resources surviving after each airbase attack.

Table 9—continued

PRINT	Output Data
3	EOD: Flights flown and demanded by mission and base. EOD: Numbers of sorties launched each hour at each airbase. EOD: Numbers of repairs waiting, tasks and repairs interrupted. EOD: Cumulative manhours on aircraft tasks, parts, and equipment repairs by shop and by base. EOD: Readiness indices ^b at each base. EOD: Cumulative sorties canceled due to ATC constraints. Current supply of spare parts at each base every six hours. Notice of initiation and completion of facility repairs. Cumulative distribution of aircraft maintenance times.
4	The numbers of interrupted tasks and repairs at noon. Available munitions by type every six hours.
5	Hourly listing of the number of aircraft waiting at each shop on each base.
8	EOT: Record of all on- and off-equipment tasks that remain interrupted or waiting.

EOT = End of trial

EOD = End of day

^aSortie data are available by base, aircraft type, mission, and priority.

^bThe readiness indices provide a cumulative measure of how quickly aircraft were prepared for flight. The index is the average percentage of each base's aircraft that were ready to fly within 2, 4, 6, and 8 hours after the previous sortie, or, when MLIST = 1, the percentage for each half-hour up to 24 hours.

3. PROVISIONS FOR CUSTOMIZED USER OUTPUT

On several occasions users have asked why one or another additional output measure had not been provided in the TSAR output. Unfortunately, it is impossible to anticipate all user interests for output measures without so greatly adding to the data storage and output listings that most users would deem many of them unnecessary and even a nuisance. Hopefully, TSAR output satisfies the widest practical range of user requirements within the limits of the data that are stored and listed.

The new data collection arrays—USERS1 and USERS2—respond to user comments by providing a structure within which the users will be able to collect the desired data by only adding a few assignment statements. The necessary Common statements, arrangements for zeroing storage arrays at time zero, as well as controls for listing such data, are all provided.

USERS1 and USERS2 permit the user to store a variety of custom results, and to list them each day, at the end of each trial, and at the end of all trials. All that the user must do is to enter the necessary assignment statements (and to be sure that the OUT Common is included in the relevant subroutines).

The USERS1 array stores three sets of up to 20 data elements for each base. The first set of data are accumulated for 24 hours and printed at the end of each day; the second set cumulates the same activities throughout each trial for printing at the end of each trial; and the third set of data collects 20 other activities throughout each trial for printing at the end of each trial. USERS2 collects the second and third data sets at the end of each trial so that the multitrial averages can be printed with the other multitrial results at the end of the run. Output for this feature is managed with the Custom Output control variables entered on CT2/5. These entries control the number of data elements that will be listed (1) daily from the first data set, and (2) and (3) at the end of each trial, and each run, from the second and third data sets.

To use this feature the user must enter the appropriate assignment statement(s) in the source code. Such statements will be of the form:

$$\text{USERS1}(N, L, \text{BASE}) = \text{USERS1}(N, L, \text{BASE}) + 1$$

where

N is the user selected element number for collecting the data of interest; and

L is 1 for actions to be listed daily, or

3 when end-of-trial results only are desired.

Normally, the values for N should be selected consecutively from 1 to 20, and the highest active value should be entered on CT2/5. The outputs will be identified as "CUSTOMIZED USER OUTPUTS" and listed in a simple tabular form. Naturally, if the user wishes a more descriptive format for any special requirements, the output routines may always be modified; the locations of these listings are quite obvious in subroutines OUTPUT and SUMUP.

4. PROVISIONS FOR SUBSEQUENT POSTPROCESSOR ANALYSIS AND GRAPHICS

In response to frequent suggestions, facilities were introduced into TSAR in 1987 to output to disk a variety of results in a form appropriate for postprocessing. Such data will provide the basis for generating a wide variety of special reports, and for creating graphic representations of selected results. Long-term storage of these machine-readable records will also make it possible to later review the results of TSAR runs without the need to find the printed output. This section describes these provisions in TSAR; development of the postprocessor, using whatever software packages are desired (e.g., SAS), will be left to the user.

There were three design objectives for this facility:

- Limit the increases in TSAR processing time.
- Limit the increases in TSAR core storage requirements.
- Limit space requirements for long-term storage.

The approach selected to provide data for postprocessing fulfills these objectives.

Two broad categories of data are needed to satisfy postprocessing requirements. The first type consists of a variety of performance characteristics that are reported at the end of each day, each trial, and at the conclusion of the several trials in the run; typically, these data are collected and listed in TSAR for several aircraft types and missions, for several shops, or for various resource categories and fill relatively long records. The other type of data involve relatively rare events such as parts shortages and parts cannibalizations, and can be described with a short record. The design chosen to handle these two types of data and to meet the three design objectives noted above is characterized by:

- Long records (130 characters) are all written to device 8 and short records (30 characters) are written to device 9.
- Each record is identified by a number that defines the type of data, and by the trial, day, base, etc. Segregation into groups of like records and other types of sorting operations are deferred until the first stages of the postprocessor.

- The user initializes DOPOST to unity on CT2/5 and designates which data are to be written onto the disk for subsequent postprocessing, using a special supplementary card that must follow CT2/5. Eighty distinct types of records may be specified; 50 are currently assigned.
- No special provisions have been introduced to provide daily summaries for various activities when cumulative records are readily available; the postprocessor can easily generate such incremental data.
- Relatively rare events, such as cannibalizations, are written on device 9 as they occur.
- Since delay statistics are only generated when the user has initialized STATFQ, thereby indicating how often such results are to be listed, a procedure has been introduced so that results are made available for the postprocessor without necessarily requiring that the results be listed in the basic TSAR output.
- All postprocessor records are of the form:

PP xy TRIAL DAY BASE Option Data

where xy is the number of the record type.

The formats are 'PP', I3, I3, I2, I3, I2, 23I5 for Device 8,
and 'PP', I3, I3, I2, I3, I2, 3I5 for Device 9.

Data Available for Postprocessing

The present TSAR source code provides for up to 80 different types of records that can be generated for postprocessing. Each type of record is identified by a number in columns 3-5, and by trial, day, base, etc. The number of such records could easily be increased beyond 80. The first two records in the output file (identified as records 998 and 999) provide a variety of dimensional data, etc., for controlling the postprocessor. The other record types currently in use and their identifying numbers are listed below; the actual format statements used with each of these records are listed in App. L. of Vol. III.

- 1 Daily sorties flown by mission and aircraft type
- 2 Daily sorties demanded by mission and aircraft type
- 3 Cumulative sorties flown by mission and aircraft type, by day

- 4 Cumulative sorties demanded by mission and aircraft type, by day
- 5 Daily sorties by base, day, and hour
- 6 Daily aircraft tasks by shop
- 7 Daily parts repairs by shop
- 8 Daily equipment repairs by shop
- 9 Cumulative on-equipment tasks by shop—number and average time
- 10 Cumulative parts repairs by shop—number and average time
- 11 Cumulative equipment repairs by shop—number and average time
- 12 Cumulative AIS repairs by station number
- 13 Periodic report on aircraft status and on deferred tasks (CT2/4)
(Output as specified by DPRINT except individual aircraft data)
- 14 Periodic report of personnel availability (see DOUTIL on CT2/5)
- 15 Completed runway and taxiway removals and repairs
- 16 Fatalities, hospitalizations, etc.
- 17 Work-rest times, etc.
- 18 Theater summary of serviceable and reparable spare parts
- 19 Cumulative NMCS hours
- 20 Parts Stocks: Serviceables and reparable by type
- 21 Aircraft lost/damaged in combat, by air attack; air and ground
aborts, cannibalizations, NMCS hours, aircraft transfers, and
sorties by mission—daily, cumulative, and multirial
- 22 Aircraft on base, and sheltered aircraft, damaged and destroyed
aircraft and shelters for each air attack

Causes for Aircraft Task Delays: Cumulative Number and Average Time

- 41 Personnel by type
- 42 Equipment by type
- 43 Parts by type
- 44 Munitions by type
- 45 TRAP by type
- 46 Material by type
- 47 POL
- 48 Facilities by number

Causes for Back-shop Delays: Cumulative Number and Average Time

- 49 Personnel by type
- 50 Equipment by type

User's Customized Outputs

- 51 USERS1(-1-) — End of day
- 52 USERS1(-2-) — End of trial
- 53 USERS2(-1-) — End of run
- 54 USERS1(-3-) — End of trial
- 55 USERS2(-2-) — End of run

Isolated Events

- 62 Runway closure times; MOS and extended MOS opening times
- 63 Times when the shelter access percentage changes
- 65 Cannibalizations by part type
- 66 Cross-cannibalizations by SRU and LRU numbers
- 67 Holes by part type (including LRU "holes" due to SRUs)
- 68 Losses sustained from a UXO explosion

End-of-Run Records

- 74 UXO, mines, and craters removed on runways and taxiways
- 75 Sorties by hour, day, and base
- 76 Daily sorties by base and mission, and theater (sd)*
- 77 Total sorties by base and mission, and theater (sd)*
- 78 Listings in SUMUP for the theater as a whole
- 79 Listings in SUMUP for individual bases
- 80 Listings generated by subroutine SUMMRY from CWOUT

*These sortie data are reported as $10 \times$ Sorties to permit use of integers and preserve tenths of sorties.

REFERENCES

1. R. R. Fisher et al., *The Logistics Composite Model: An Overall View*, The RAND Corporation, RM-5544-PR, May 1968.
2. T. C. Smith et al., *A User's Manual for SAMSOM II: The Support Availability Multi-System Operations Model*, The RAND Corporation, RM-4923-PR, November 1967.
3. D. E. Emerson and L. H. Wegner, *TSARINA--A Computer Model for Assessing Conventional and Chemical Attacks on Airbases*, The RAND Corporation, N-3010-AF, forthcoming.
4. L. H. Wegner, *The Taxiway Repair Schedule Problem: A Heuristic Rule and a Branch-and-Bound Solution*, The RAND Corporation, N-1883-AF, October 1982.
5. D. E. Emerson, *An Introduction to the TSAR Simulation Program: Model Features and Logic*, The RAND Corporation, R-2584-AF, February 1982.
6. D. E. Emerson and L. H. Wegner, *TSAR User's Manual--A Program for Assessing the Effects of Conventional and Chemical Attacks on Sortie Generation: Vol. I, Program Features, Logic, and Interactions*, The RAND Corporation, N-2241-AF, August 1985.
7. D. E. Emerson and L. H. Wegner, *TSAR User's Manual--A Program for Assessing the Effects of Conventional and Chemical Attacks on Sortie Generation: Vol. II, Data Input, Program Operations and Redimensioning, and Sample Problems*, The RAND Corporation, N-2242-AF, August 1985.
8. D. E. Emerson and L. H. Wegner, *TSAR User's Manual--A Program for Assessing the Effects of Conventional and Chemical Attacks on Sortie Generation: Vol. III, Variables and Array Definitions, and Other Program Aids*, The RAND Corporation, N-2243-AF, August 1985.
9. AFM 67-1, *Air Force Supply Manual*, Vol. II, Part 2, Chap. 11.
10. AFLCM 171-46, War Readiness Spares Kit/Base Level Self Sufficiency Spares (WRSK/BLSS) Requirements Computation System, DO-29-MP, March 4, 1983.
11. H. M. Berlin, L. Stroschein, and R. F. Goldman, *A Computer Program to Predict Energy Cost, Rectal Temperature, and Heart Rate Response to Work, Clothing, and Environment*, ED-SP-75011, Edgewood Arsenal, Aberdeen Proving Ground, Maryland, November 1975.
12. *Draft Standard, Thermal Environments, Analytical Determination of Thermal Stress*, ISO/TC159/SC5/WG1, Working Group "Thermal Environments," of subcommittee 5 "Ergonomics of the Physical Environment," International Standards Organization, n.d.
13. R. Saucier, *A Mathematical Model for the Atmospheric Transport and Diffusion of Chemical Contaminant (NUSSE2)*, Chemical Systems Laboratory, Technical Report ARCSL-TR-81071, Aberdeen Proving Ground, Maryland, November 1981.
14. R. K. Dumbauld and C. R. Bowman, Second Supplemental Technical Report for U.S. Marine Corps Chemical Logistics Evaluation Comparison of the Aero 14/B, TMU 28/B, and BLU 80/B Spray Delivery Systems, U.S. Army Dugway Proving Ground, Utah, March 1981.
15. Kenneth S. K. Chinn, *A Simple Method for Predicting Chemical Agent Evaporation*, Report Number DPG-TR-401, U.S. Army Dugway Proving Ground, Utah, September 1981.

16. J. Monaghan and W. R. McPherson, *A Mathematical Model for Predicting Vapour Dosages on and Downwind of Contaminated Areas of Grassland*, Suffield Technical Paper No. 386, Defense Research Establishment, Suffield, September 1971.
17. Private communication from the Air Force's Aerospace Medical Research Laboratory.
18. P. S. Hewlett and R. L. Plackett, "A Unified Theory for Quantal Responses of Drugs: Non-interactive Action," *Biometrics*, Vol. 15, p. 591, December 1969.