INVENTORY BEHAVIOR AT REMOTE SITES

Final Report

NASA/ASEE Summer Faculty Fellowship Program -- 1987

Johnson Space Center

Prepared by:	William C. Lewis, Jr., Ph.D
Academic Rank:	Research Assistant Professor
University and Department:	University of Washington Mechanical Engineering Dept. Industrial Engineering Prog.

NASA/JSC

Abstract

An operations research study is conducted concerning
inventory behavior on the Space Station. Historical data
from the Space Shuttle is used. The results demonstrate a
high logistics burden if Space Shuttle reliability
technology were to be applied without modification to
Space Station design (which it will not be). Effects of
rapid resupply and on board repair capabilities on
inventory behavior are investigated.

Introduction

"Adde parvum parvo magnus acervus erit."
"Add little to little, and there will be a great pile."
Ovid, ca. 50 A. D.

Intuition is linear and works best on small systems.
Large systems with additive effects are not easy to
understand. They produce surprises, like Ovid's pile.
Large systems with nonlinear components and additive
effects are worse: they produce shocks. These systems
cannot be intuitively understood. They must be
mathematically understood, or not understood at all.

Inventory systems are nonlinear in important ways, and
they are additive as well. This report describes an
operations research study of local inventory required to
support remote site operations. The study attempts to
quantify inventory requirements for generic remote site
operations in such a way as to anticipate the order of
inventory demands early enough to permit their
incorporation into the design process.

As a test case, historical data from Space Shuttle
operations is used. This data was very kindly provided by
Rockwell International, which deserves much of the credit
for the usefulness of these results. In some ways, Space
Shuttle data is not appropriate for remote site modeling,
in that the Shuttle was designed for ground maintenance
after each 1 week mission, whereas the Space Station will
never receive ground maintenance, and has a 30 year design
life. However, Space Shuttle module reliability was
probably increased to the limits of cost effectiveness for
1970 era technology in an attempt to reduce maintenance
costs and.increase vehicle availability. Also, it is
tempting to rely on proven Space Shuttle technology when
designing the Space Station. The analysis does not show
how Space Station inventory will behave. It does show how
Space Station inventory would behave if Space Station
technology had the same overall failure rate and repair
time distribution as Shuttle technology, which it most

certainly will not.

More significant than the actual numbers is the demonstration that inventory stockouts (that is, exhaustion of inventory for some classes of modules) will always be a factor. If there is a module failure, and no modules of that type are in inventory, there are three courses of action: 1) ignore the failure, 2) repair an already failed module, or 3) get a replacement module by ex tempore resupply. If the failure is in a redundant system, option (1) may be attractive. If not, (2) or (3) will be the options of choice. The results of this analysis strongly suggest that (2) and (3) will both be desirable during Space Station operations, as they are at any remote site operation.

It is discovered on analysis of Space Shuttle operations records that systems comparable to the Space Shuttle require unexpectedly large inventory for long missions. It is further discovered that, even with a fairly large spares inventory, shortages will occur at a significant rate. Makeup of these shortages, either through repair and improvisation, or through special equipment delivery, may be a significant problem in space station and planetary expedition logistics.

The mathematics, the data, the analyses, and the conclusions are described separately to permit replication by the interested reader. A reader not interested in replication or time consuming development of mathematical intuition in this area may skip to the results section.

Mathematical Approach

The techniques used are elementary. Elementary statistical methods are to be preferred over advanced methods on grounds of clarity if they correctly model all significant aspects of a situation. The approach described below is correct in the sense that it provides an estimate of stockout probability distribution parameters which is moderately reliable and can be shown on theoretical grounds to be insensitive to actual item failure distributions, which are not ordinarily known. I make no claims to the method's optimality if actual item failure distributions are known exactly.

The statistical method used to set inventory stock levels and predict number of stockouts is developed below:

1) Independent actors, each performing the same act with an idiosyncratic distribution of times between performance, will as a whole perform the act with a negative exponential distribution of times between

performance. For example, each person in a dialing area
makes telephone calls on an individual basis. The time
between calls for any individual is a random variable, and
the distribution with which this variable is associated
will, in general, vary from one person to the next. That
is, each person's time between calls will have an
associated and unique, or idiosyncratic, distribution. If
calls made by all people in the dialing area are observed
(by monitoring call arrivals at the telephone exchange),
it will be observed that the time between arrival of new
calls will behave like a random variable with a negative
exponential distribution. The same argument holds true
for arrival of customers at a bank, arrival of cars at a
gas station, or "arrival" of equipment failures. The
density function for a negative exponential distribution
is:

$$f(Lambda, t) := Lambda * exp(-Lambda*t)$$

"Lambda" is the average arrival rate, in events/unit time.
"t" is the time between events. It is obtained by adding
up individual arrival rates for all independent actors.
The average time between arrival is, of course, 1/Lambda;
the variance in time between arrivals is $1/(Lambda^2)$.
Thus, high rates of arrival give a smaller variance in
interarrival time than do large rates. That is, the
process varies more at low arrival rates than at high.

   2) Any process with density function f(Lambda, t) of
interarrival time will have a Poisson distributed number
of events per unit time. That is, if instead of observing
time between events, one counts events per unit time, one
will find that the number of events per unit time is
Poisson distributed if and only if the time between events
has a negative exponential distribution f(x).

   The Poisson distribution is as follows:

$$P(Lambda, N) := exp(-Lambda)*(Lambda^N)/N!$$

"Lambda" is defined as before (and is indeed identical to
the Lambda for f(Lambda, t)). Lambda is also the mean and
variance of P(Lambda, N). Note that variance of
P(Lambda, N) increases with increased Lambda, whereas that
of f(Lambda, t) decreased under the same condition. The
higher the arrival rate, the more variance, hence the more
spread in observed number of arrivals over several
observations. "N" is the number of events. Note that N
is a cardinal number, that is a member of the set
{0, 1, 2, ... }. N cannot be 3.5, or -2, as these cannot
count events. This will cause certain minor problems
later, and has the immediate consequence that P(Lambda, N)
is not a density function, but rather a distribution. In
general, the probability of N failures per unit time will

be given by P(Lambda,N). The above can be re-written recursively as:

$$P(Lambda,0) := exp(-Lambda)$$

$$P(Lambda,N) := P(Lambda,N-1) * Lambda/N$$

The behavior of the function P(Lambda,N) as expressed recursively is clear. If Lambda>1, it P(Lambda,N) increases until N>Lambda, then decreases. If Lambda<1, P(Lambda,N) always decreases. Note that P(Lambda,N) is always positive and finite, as one would expect a probability to be.

3) In planning for inventory, one might be interested in having just enough inventory to avoid stockout. That is, one wants to have enough in inventory so that one never runs out, and at the same time wants nothing superfluous in inventory so that one does not waste money. This is impossible, as will be shown.

If failures are Poisson distributed, then the probability of having N failures or less is given by the cumulative Poisson distribution,

$$C(Lambda,N) := sum(0,N,P(Lambda,N))$$

Obviously, the probability of covering all failures when there are N items in inventory and Lambda is constant is just C(Lambda,N). By the same token, the probability of stockout must be 1-C(Lambda,N).

The behavior of 1-C(Lambda,N) is therefore of interest. Let us re-write:

$$1-C(Lambda,N) = sum(N+1,infinity,P(Lambda,N))$$

and we have seen above that P(Lambda,N) will be finite (although possibly very small) for all cardinal numbers N.

From this it is apparent that 1-C(Lambda,N) will always be finite, no matter how large N, and that (accordingly) there is always some probability of stockout, no matter how large N, the number of replacement units in inventory.

4) Fortunately, the above is not merely an impossibility proof. It also leads to an algorithm for sizing inventory. If one is satisfied with a fairly small desired probability of stockout per mission, Psd, then, if Lambda per mission is known, one need only specify an N large enough such that:

$$1-C(Lambda,N) < Psd \quad and \quad 1-C(Lambda,N-1) > Psd$$

N is the smallest number of items which can be kept in
store if Psd is to be met. N is uniquely specified by
this algorithm, which may be written as N(Psd). The
actual probability of stockout is, of course,

$$Psa := 1 - C(Lambda, N(Psd))$$

Lambda for the mission is Lambda per unit time multiplied
by number of units of time per mission. Note that, in
general, Psa < Psd. This is a result of N being a
cardinal number, and is the inconvenience mentioned above.
The probability of 1 stockout is P(Lambda, N+1), the
probability of 2 stockouts is P(Lambda, N+2) and so on.
The expected number of stockouts for the entire mission is
thus E(s) := sum(n=N, n=infinity, P(Lambda, n)). Ordinarily
only a few terms of this sum must be calculated before
P(Lambda, n) becomes negligible.

5) One must somehow specify Psd. For inventories
containing only one kind of item this is simple. One
defines a cost function Ma(N(Psd)) and minimizes it. For
inventories containing several noninterchangeable classes
of items, each with its own Psaj (j being an index
describing the kind), but supporting a single system (such
as the Space Station), this procedure is not satisfactory.
The cost function is properly a function of system
reliability, which is a function of subsystem reliability,
which is in turn a function of the number of working
modules within each subsystem, which is decreased by each
stockout. While the Psaj must obviously meet certain
criteria, because these criteria are dependent on the
criticality of the subsystem supported by units of kind j,
there emerges an additional, system wide, criteria: the
total number of stockouts during a mission (TSM). If j is
over about 10, one would expect system degradation and a
TSM nonlinearly related to the Psaj.

6) By the reasoning in (1), the random variable
representing time between stockouts for the inventory as a
whole should have a negative exponential distribution.
TotalLambda should have a value equal to the sum of the
hazard rates over all classes in inventory. The hazard
rate for each kind in inventory over the entire mission is
the expected value of stockouts the kind, E(sj). That is,

$$TotalLambda := sum(all\ j,\ E(sj))$$

7) Given the distribution types, and TotalLambda, one can
calculate probability of having N or fewer stockouts per
mission by the cumulative Poisson as above:

$$C(TotalLambda, N)$$

This can be used to calculate a range of stockouts most likely to be seen. We assume that things which happen less than 5% of the time are often not seen; thus, the smallest N seen would be Nsmall such that

$$C(TotalLambda, Nsmall-1) < 0.05$$
$$and$$
$$C(TotalLambda, Nsmall) > 0.05$$

Similarly, the largest N seen would be Nlarge such that

$$C(TotalLambda, Nlarge-1) < 0.95$$
$$and$$
$$C(TotalLambda, Nlarge) > 0.95$$

These are approximations, of course, and merely show that one has at least a 5% chance of seeing Nsmall or less and at most a 5% chance of seeing Nlarge or more.

That is a narrative description of the technique. Perhaps it should be illustrated graphically.

The inventory of one kind of unit could be sized by drawing its Cumulative Poisson [C(Lambdaj,N)], and using the curve to translate Psd into minimum N as in Graph A, which looks like a flight of stairs. It is obvious that increasing the number of units in stock from 3 to 4 will reduce Psa to below the 5% minimum specified.

By the same token, given TotalLambda for the completed inventory, stockouts could be determined by drawing its Cumulative Poisson [C(TotalLambda,N)], as in Graph B, and using the curve to translate probabilities of .05 and .95 into numbers of stockouts as shown. Again, it is obvious that the 5% limits of observation lie somewhere around 33 and 55.

Buried in the above analysis is the assumption that stockout hazard rates are constant. In fact, they increase as stock on hand is consumed. In the normal Space Station case, inventory will not be severely depleted during the mission. Perhaps half of it will be used up, and perhaps 1.5% of the inventory classes by module kind will have stockout. One could therefore confidently predict an increase in stockouts during the latter part of a mission. Unfortunately, there seems to be no easy way to solve analytically for stockout distribution as a function of time, and simulation would be necessary to determine this distribution. The distribution given is a valid average distribution for the mission; one would expect the lower limit to be more

CumPoisson
1.04

.99
.95
.71

π

.148
$\approx$

.000   C   1   2   3   4   5   6.00

Number of Units in Inventory

- - -
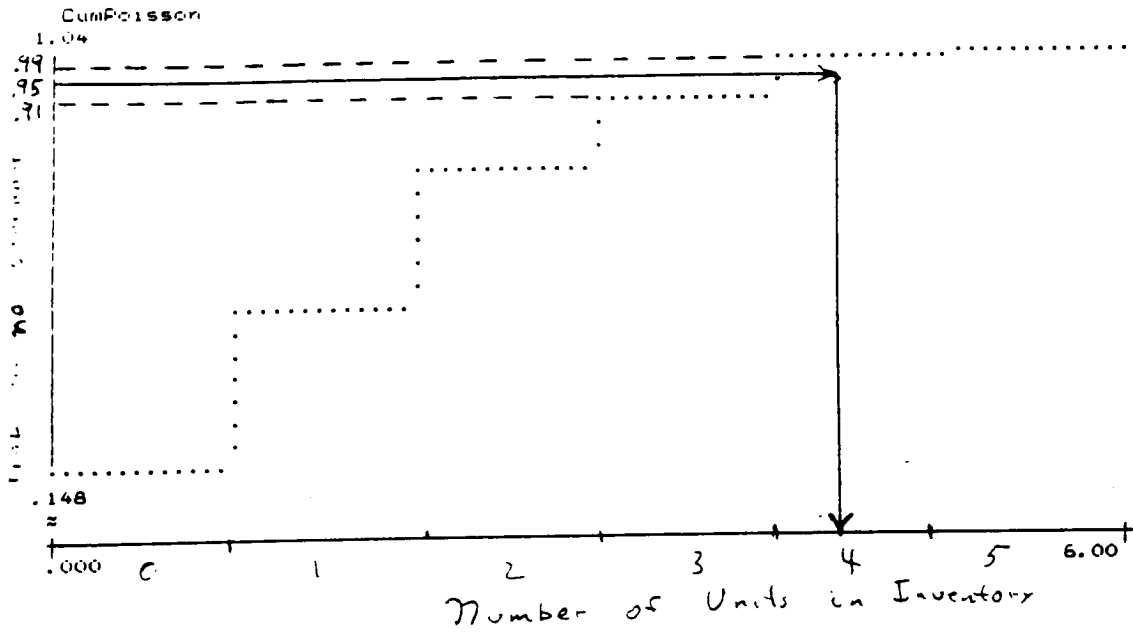
List of function values.
CumPoisson

x	CumPoisson(x)
.00000000	.18824707
1.0000000	.50261967
2.0000000	.76512079
3.0000000	.91124641
4.0000000	.97225386
5.0000000	.99263035
6.0000000	.99830180

$\lambda = 1.67$ failures /mission

$\left(\begin{array}{c} 6 \text{ modules} \\ \text{supported} \end{array} @ 7779 \text{ hrs } \dfrac{MTBF}{Module} \right)$

← 3 modules ⇒ 8.9% chance of stockout.
← 4 modules ⇒ 2.8% chance of stockout.
← 5 modules ⇒ .74% chance of stockout.

Graph A

$\lambda = 43.4$ stockouts/mission

CumPoisson
1.05

95%

P(N)

5%

-.0500   32 33   54 55   90.0

N   Number of Stockout Events
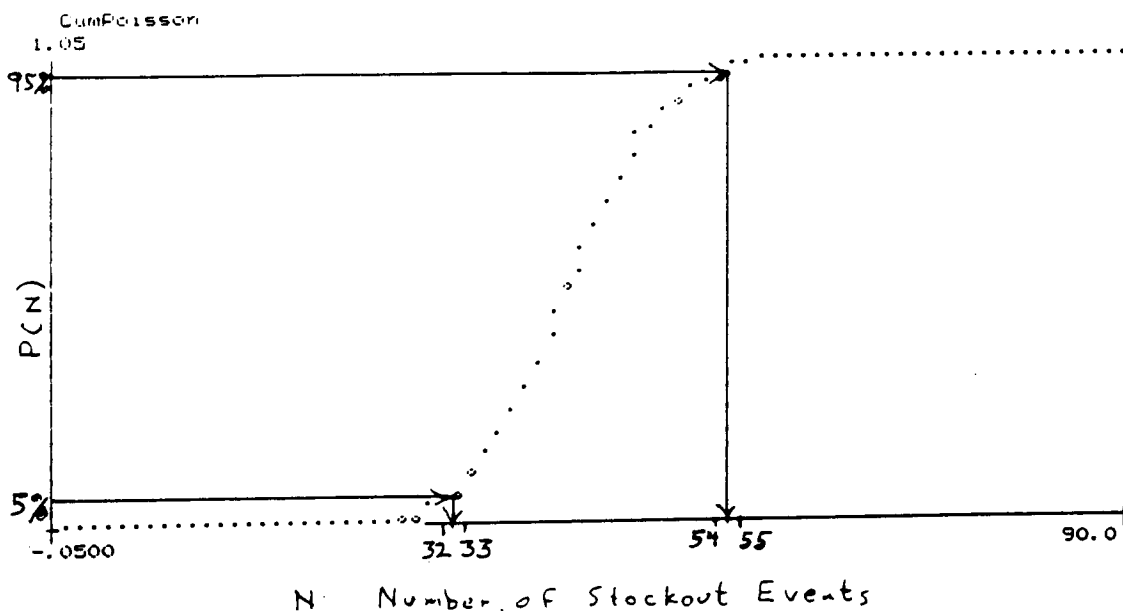
List of function values.
CumPoisson

x	CumPoisson(x)
30.000000	.020759638
31.000000	.030820003
32.000000	.044454650
33.000000	.062373498
34.000000	.085230079

←5%

List of function values.
CumPoisson

x	CumPoisson(x)
52.000000	.91406814
53.000000	.93423665
54.000000	.95043461
55.000000	.96320716
56.000000	.97309883

←5%

Graph B

18-9

typical of the mission's early part and the upper limit
more typical of the mission's later part. Again,
simulation is necessary to if the distribution is to be
numerically specified as a function of time.

Data

    Two kinds of data can be used. The first set is for
initial sizing studies, and is approximate. A uniform
MTBR (Mean Time Between Removals) of 8 000 hours,
MTTReplace (Mean Time To Replace) of 4 hours, MTTRepair
(Mean Time To Repair) of 80 hours, and average commonality
of 2 items per type are reasonable starting values.

    Historical data on Shuttle subsystem characteristics
was provided by Rockwell International. The data set used
was the Logistics Master Control File, Report No. 9100, as
of 07/13/87. Only QPEI / WUC, MDR, and MTBF data were
used. QPEI / WUC (Quantity Per End Item) is the number of
the item per individual OV. MDR (Maintenance Demand Rate)
is number of unexpected removals per each 1000 hours of OV
operation. Therefore, MTBR (Mean Time Between Removals) =
1000/MDR. MTTRepair (Mean Time To Repair) is the actual
time on the repair bench required to restore a removed
module to fully tested and operational status. RTAT is
the number of hours including handling required for module
repair. Generally, RTAT>>MTTRepair. Since handling in
the sense of transport to repair facilities will not be a
major factor at a remote site, RTAT was not used in the
analysis. MTTReplace (Mean Time To Replace) was
arbitrarily set to 4 hours.
    There are 2763 classes of modules in Report No. 9100.
Since STOCK2 accepts only 400 module types, some data
selection and trimming was necessary.
    The first decision was to use only data with MTTR
specified. Of 2763 classes, only 1702 (62%) specified
MTTR. A t test (see discussion and Table, below)
conducted on hazard rate and population supported per
class showed no reason to reject the hypothesis of equal
means, so the subset was accepted as representative. This
subset will be called "MTTR Subset".

    The second decision was to select a random subset of
MTTR specified data. This involved writing a program
(RAND) to generate 400 random numbers sorted in increasing
order over a range equal to the cardinality of MTTR
Subset. These numbers were treated as indices specifying
a second subset of MTTR Subset, called Subset A (which
accordingly has 400 classes). A t test (see discussion
and Table below) conducted on hazard rate, population
supported per class, and mean time to repair showed no
reason to reject the hypothesis of equal means, so the
subset was accepted as representative. Subset A was

actually used in analyses. It would be desirable to
conduct these analyses with a Subset B, C, etc., but time
did not permit this.

Statistical characteristics of the data are as
follows:

Data Set	Reciprocal Average Hazard Rate	Sample Standard Deviation
Complete	8291 hr	8286 hr
MTTR>0	8264 hr	8264 hr
Not given	8333 hr	8321 hr
Subset A	7231 hr	7231 hr

Data Set	Average MTT Repair	Sample Standard Deviation
Complete	51 hr	253 hr
MTTR>0	82 hr	318 hr
Not given	N/A	N/A
Subset A	93.3 hr	348 hr

Data Set	Average Population Per Class	Sample Standard Deviation
Complete	1.74	2.62
MTTR>0	1.85	2.91
Not given	1.55	2.07
Subset A	1.91	3.04

Note: Total population for any data set can be obtained by
multiplying Average Population per Class by Number of
Classes.

	Hazard Rate	MTT Repair	Population Per Class
Complete and MTTR Subset	t= 4.72E-4 Nu= 5.9E23	t= 3.41 Nu= 3.0E3	t= 1.27 Nu= 1.2E8
MTTR Subset and Subset A	t= 0.825 Nu= 9.5E21	t= 0.594 Nu= 570	t= 0.360 Nu= 1.5E6

For the Nu values given, the t random variable is distributed normally for all practical purposes. A $t < 1.282$ implies acceptance of the null hypotheses (that data set means are equal) on the 10% level.

The final table, listing t statistics, suggests that the MTTR Subset is very similar to the complete data set except, of course, in MTTRepair. The MTTR Subset and Subset A appear alike for all data types. Accordingly, the behavior of inventory for Subset A should statistically resemble that of the MTTR subset for all data types, and should statistically resemble that of the complete data set for all data types except MTTRepair.

Subset A is stored under file name STOCK2.STK on floppy disk, and will be included with STOCK2. Use STOCK2's export feature to change STOCK2.STK into STOCK2.ASC, which can be imported by a database program as a standard ASCII file with comma as delimiter.

Software and Computer

The software and hardware used are commonly available; those readers wishing to check or extend results may obtain a copy of the software and data on request to the author, if said request is accompanied by a 5 1\4 inch MS-DOS floppy diskette, formatted for a 360k drive. The program is written in Turbo Pascal, and it is advisable to use a database program, such as Borland's Reflex, for data analysis. It should run on any MS-DOS machine. Stock2R.com does not require a mathematics coprocessor, Stock2X.com does. The program is written in Turbo Pascal, by Borland International, and you will need a compiler for this language should you wish to modify and re-compile the program. Please use STK2R or STK2X to install the program if its messages are not properly displayed on your terminal screen. Instructions for the STK2 programs are in document INSTALL.TXT.

A program to generate random numbers uniformly distributed over a specified range is included as RAND and

RANDX. Again, the X suffix means a mathematics
co-processor is necessary. Installation is by RND and
RNDX. The source file is provided, but requires sort
procedures from Borland's Turbo Toolbox to compile.

A random subset of 400 module descriptions from the
Space Shuttle's operating history is on the diskette as
STOCK2.STK. It is suggested that this information (which
includes output data) be imported to STOCK2 by the "typed
file" initialization option. Detailed analysis of the
results should by choice be attempted with the help of a
good database program, as 400 items are too many for the
crude display facilities in STOCK2. After alteration by
the database program (Reflex by Borland may be a good
choice) the data can be re-imported to STOCK2. Use comma
as the delimiter.

Although hazard rates scale linearly, confidence
limits do not. Accordingly, a program, called LIMITS.com,
has been provided to scale the results of STOCK2 up or
down. To obtain confidence intervals for a target
installation, such as the Space Station, scale up the
global hazard rates from Stock2 by:

TargetLambda=
    Stock2Lambda * TargetModuleNumber/Stock2ModuleNumber

with Stock2Lambda and Stock2ModuleNumber being the values
from Stock2, and TargetLambda and TargetModuleNumber being
the values of the target installation. For example, if
Stock2 gives a repair Lambda of 3.16 for 400 classes of
modules with an average of 1.91 modules / class and the
Space Station is assumed to have 2578 modules, then the
repair hazard rate for the Space Station is:

$$3.16*2578/(400*1.91) = 10.7$$

To obtain Space Station confidence intervals, run LIMITS
and specify a hazard rate of 10.7. Use the same scale
factor for inventory.

Methods of Analysis

Item demand will change as the population of supported
modules ages. If MTBF only changes by a few percent
within an inventory period, this can be treated by setting
MTBF to the average for the period in question.
Calculations involving multiple periods use the
appropriate average MTBF for each period. For the
purposes of this report, a constant MTBF is the
appropriate choice.

A rough feel for this class of inventory problem can be acquired using Borland's Eureka, and interactive equation solver described under "Software and Computer." The algorithm above can be implemented quite directly, and run for simple cases. In particular, it is possible to graphically illustrate the relation between the Poisson and Cumulative Poisson distributions for various values of hazard rate, and to develop one's intuition for stocking levels and stockout probabilities.

In general, one finds that (a) stock levels increase nonlinearly as probability of stockout decreases and that (b) number of modules in inventory divided by number of modules in use tends to decline as population supported increases. The results of such investigations are not reported here, as they are not based on historical data.

When one wishes to use non-aggregated historical data, the sheer mass of computation involved overtaxes Eureka, and a specialized program must be used. This program permits direct entry of equipment characteristics, and executes the mathematical approach described above to determine replacement and stockout hazard rates for each kind of equipment and globally, for the entire inventory. This program, STOCK2, was used to conduct the following preliminary analyses:

1. Inventory behavior under varying uniform stockout probabilities.

    Specified probability of stockout for all parts types are set to a uniform value for each run. Inventory size, repair event hazard rate, and 90% confidence interval for number of repair events are calculated.

2. Effect of provision for module repair / cannibalization on inventory behavior and crew size.

    Probability of exhaustion is set independently for each part type so as to minimize stockout probability for part types with large repair times, and inventory level is kept constant by increasing stockout probability for part types with small repair times. The net result is to keep inventory size approximately constant while decreasing time required to repair or cannibalize failed modules. After establishing a set of probabilities at average specified probability of stockout of 0.1, probabilities were linearly transformed to obtain other average specified stockout probabilities, and the effect of this on repair time and inventory stock level were recorded.

3. Effect of mission length on inventory behavior.

    Longer and shorter mission times are considered with fixed average specified stockout probability. The

shortest time is the length of a Space Shuttle mission;
the longest is about that of a round trip Mars mission.
The behavior of inventory at an average specified stockout
probability of 0.01 is considered.

4. Effect of commonality on inventory behavior.
    Greater commonality was simulated by multiplying
population supported in Subset A by 5 and by 10, then
scaling all results up to the same Space Station number of
modules.

    There was not sufficient time to conduct these
analyses thoroughly. It is likely that only the stronger
effects were discovered, and that many weak effects remain
as yet unknown.

   Results

General:

    Space Station is assumed to have 2578 ORUs at 1.91
ORU/Class, for a total of 1350 classes.

    On at most 5% of the missions values will exceed the
Upper Limit; on at least 5% of the missions values will be
below the Lower Limit; the two limits thus form an
approximate 90% confidence range. An exact 90% confidence
range cannot be specified for events, as they cannot be
subdivided.

1. Inventory behavior under varying uniform stockout
    probabilities.

Specified Probability of Stockout	Inventory Size	Stockout Hazard Rate	Number of Stockouts Lower Limit	Number of Stockouts Upper Limit
1.00	0	853	806	901
0.316	967	155	135	176
0.100	1550	51	40	63
0.0316	2110	18	12	26
0.0100	2800	5.2	2	9

Graph C: No Repair
Log P(Stockout) vs. Inventory Size

Inventory Size (Modules) vs. Log P(Stockout per queue)
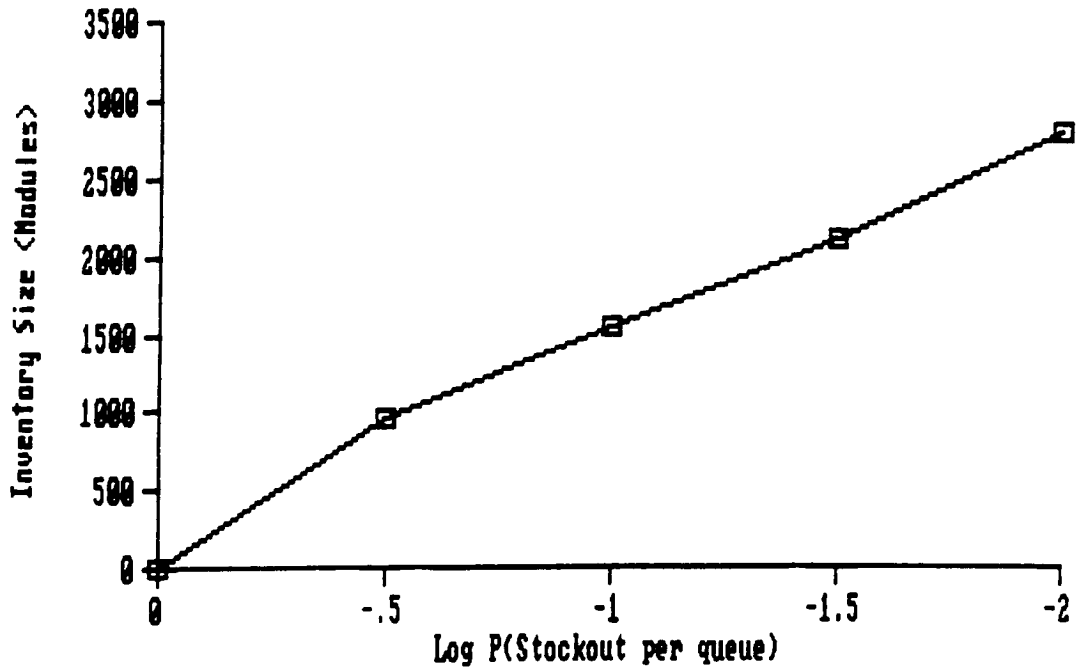
Graph D: No Repair
Log P(Stockout) vs. Stockout Hazard
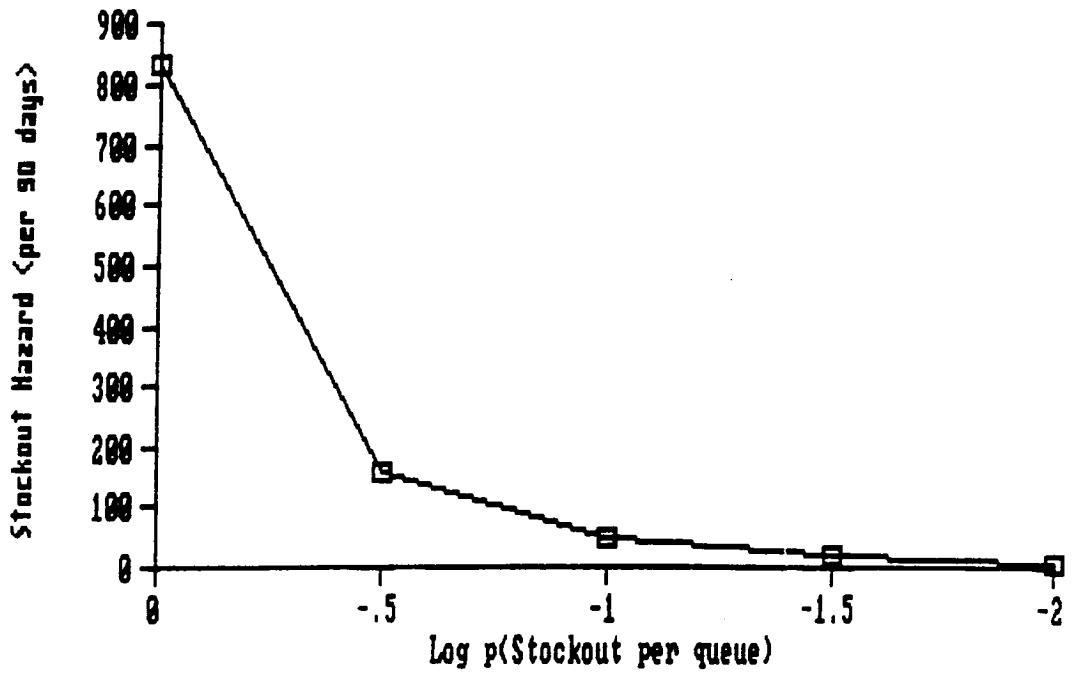
Stockout Hazard (per 90 days) vs. Log p(Stockout per queue)

2. Effect of provision for module repair /
   cannibalization on inventory behavior and crewsize.

Specified Probability of Stockout	Inventory Size	Stockout Hazard Rate	Number of Stockouts	
			Lower Limit	Upper Limit
1.00	0	853	806	901
0.316	1400	102	85	119
0.100	1910	46	35	57
0.0316	2390	18	11	25
0.0100	3150	4.0	1	8

Specified Probability of Stockout	Average Crew on Repair	Repair Hours Hazard Rate	Number of Repair Hours	
			Lower Limit	Upper Limit
1.00	123	88200	87700	88700
0.316	5.9	4220	4110	4320
0.100	1.6	1150	1100	1210
0.0316	0.52	421	388	455
0.0100	0.17	121	103	139

Graph E: Repair
Log P(Stockout) vs. Inventory Size
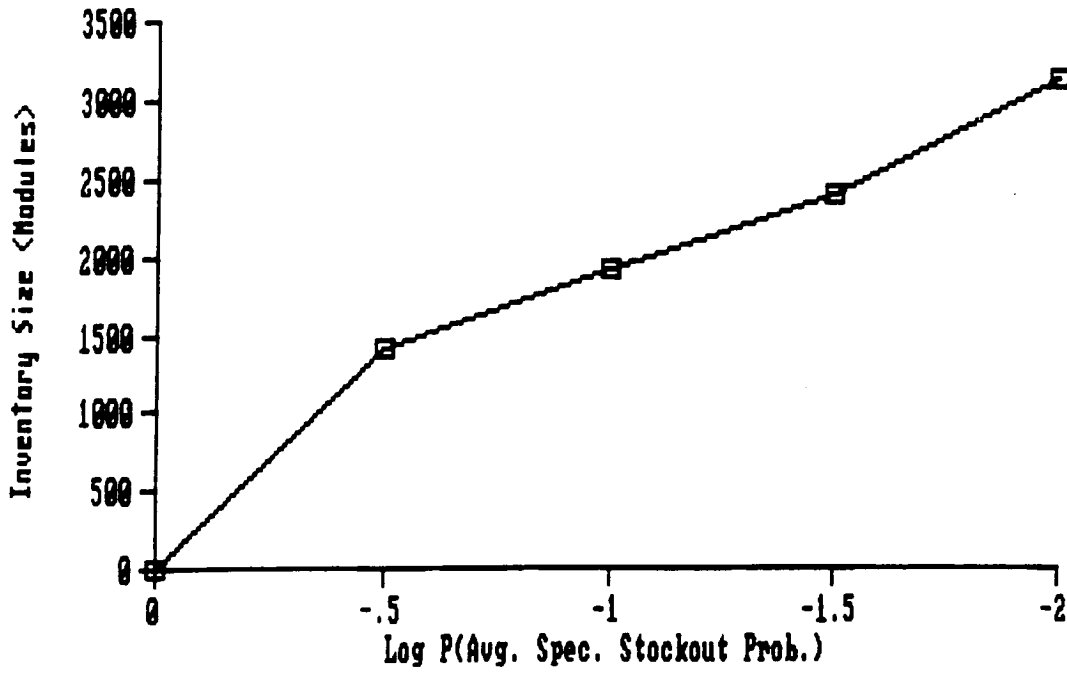
Graph F: Repair
Log P(Stockout) vs. Crew on Repair
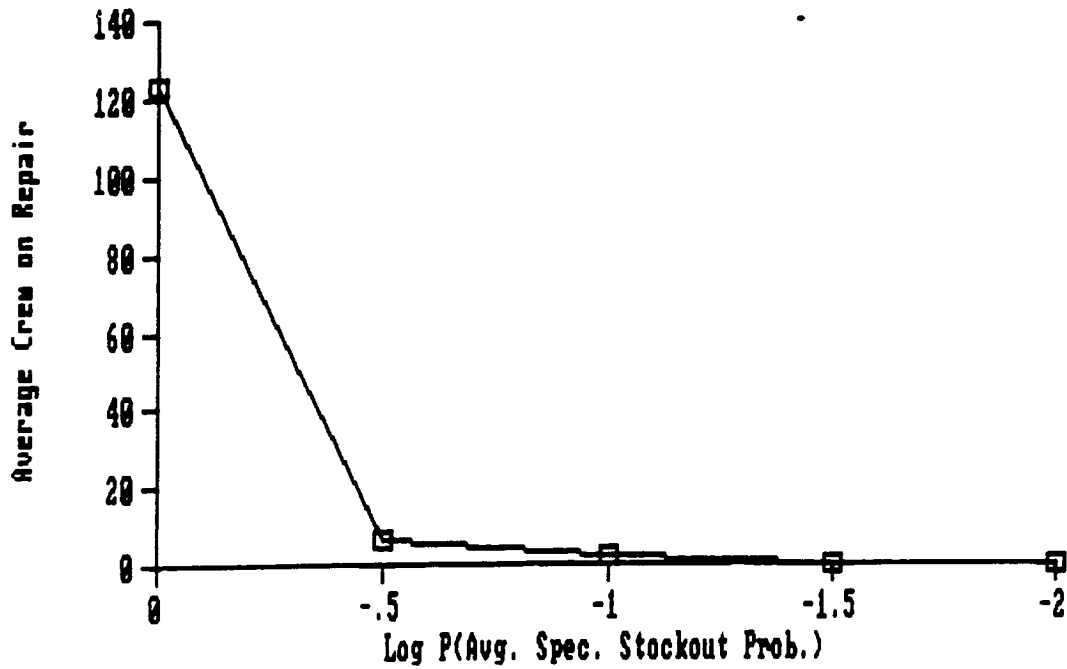
3. Effect of mission length on inventory behavior.

These calculations use data optimized for repair, as given in STOCK3.STK. Average specified probability of stockout is slightly less than 0.1.

Mission Length (Days)	Inventory Size	Stockout Hazard Rate	Number of Stockouts	
			Lower Limit	Upper Limit
7	347	10.6	6	16
14	542	15.7	9	22
90	1930	44.0	33	55
270	2560	72.0	58	86
900	12100	97.5	82	114

Mission Length (Days)	Average Crew on Repair	Repair Hours Hazard Rate	Number of Repair Hours	
			Lower Limit	Upper Limit
7	9.5	530	492	568
14	5.9	658	617	701
90	1.6	1110	1060	1170
270	0.9	1910	1830	1980
900	0.4	2660	2580	2750

4. Effect of commonality on inventory behavior.

Commonality	Inventory Size	Replacement Event Hazard Rate	Stockout Event Hazard Rate
1.91	1930	855	44
9.55	1350	852	15
19.1	1270	852	13

Discussion

As can be seen, there is significant intuitive content
to the mathematical treatment:
1) Hazard rates are additive.
2) Inventory levels are nonlinear with protection.
High protection from stockouts requires a big inventory.
Perfect protection requires an infinite inventory.
3) Moderate protection can be accomplished with a much
smaller inventory.
4) If there are a large number of separate
inventories, the hazards from "moderate protection" add up
to significant numbers of observed stockouts.

Mathematical technique, in this case, becomes most
interesting when applied to the interpretation of
historical data. The intuitive content requires scaling;
are we talking about an inventory of 100 or 1,000,000?
Will we run out of stock for 1 or 100 items? Historical
data for the Space Station is nonexistent; it has not been
built. Let us suppose, however, that the Space Station
was to be built using the same general order of
reliability as the Space Shuttle. This is unlikely, but
makes an interesting test of the mathematics, and gives us
some idea of the changes in Shuttle era technology
required to make a logistically tractable Space Station.

Inventory behavior under varying uniform stockout
probabilities is as expected. Stockout hazard rate
decreases nonlinearly as specified probability of
exhaustion increases. Inventory size increases
exponentially with specified probability of exhaustion.
In practice, an inventory size of 2800 supporting 2578
running might be chosen to ensure that there will usually
be fewer than 9 uncorrected failures per 90 day mission.
This is about one item in inventory for each item running.

The effect of provision for module repair and
cannibalization is to trade inventory volume for labor.
If one is willing to devote the labor of 1.6 people in
orbit, inventory can be reduced by 886 modules, from 2800
modules to 1914 modules, and number of uncorrected
failures reduced from a maximum of 9 to a maximum of 0.
This requires a distortion of inventory policy, however,
in that one must stock more of the high repair time
modules and fewer of the low repair time modules to obtain
the low average repair time which makes on orbit repairs
feasible. If only 0.52 people work on repairs, inventory
size is reduced only by 411 modules, and assigning 0.17
people actually increases inventory by 351 modules.

Increased mission length, or time between resupply,
tends to increase inventory. Accordingly, a doubling of
the Space Station's mission from 90 to 180 days would be

expected to increase its required inventory.
Unfortunately, not enough points could be taken in the
available time to quantify this effect reliably. The
algorithm would have to be rewritten so as to keep
stockouts/unit time constant (assuming repair) or to keep
total stockouts constant (assuming no repair or on demand
resupply). However, it is clear that increasing mission
length while keeping specified probability of stockout
constant will produce significant increases in inventory
levels. The 900 day mission requires 4.7 items in
inventory for every item running.

The mission length data provides a reality check of
the data and mathematics. The predicted removal hazard
rate for the Shuttle is 136 removals/mission (90%
confidence interval of 117 to 155), whereas Rockwell's
data indicate that there are actually 224.5
removals/mission during ground reconditioning. Since not
all of the actual removals are for unexpected faults, the
prediction seems reasonable.

Specified probability of stockout (SPS) was kept constant
for all mission lengths. The effect of SPS is to set a
global upper limit on number of stockouts. During the
short mission, many of the more reliable modules will not
fail, so that the actual probability of exhaustion for
that module time will be much less than the SPS for that
type, and the global limit will not be approached. As the
mission grows longer, the difference between SPS and
actual probability of exhaustion diminishes, and the
global upper limit is more closely approached. For
example, the results listed show a global stockout hazard
rate increasing from 10.6 to 97.5. The interesting side
effect is that the increase is not linear. That is, the
number of stockouts per day becomes less as the mission
gets longer, and fewer people are required to repair
failures. This suggests that failure repair may be a more
attractive option as missions grow longer. Alternatively,
by devoting more crew to repair activities, inventories
shown might be substantially reduced.

The effect of commonality on inventory behavior is to
decrease inventory. While only three points were taken,
they suggest that inventory size can be reduced by a
maximum of about a half by increasing commonality. It is
not clear to what extent commonality can be increased.

Operationally, the above shows conclusively that the
Space Station (and other remote site bases) cannot be
built under the same design philosophy as the Space
Shuttle. The Space Station has a 30 year design life, and
does not visit a refitting base. It is the refitting
base. Accordingly, it must be kept operational on a
continuing basis.

Achieving this will clearly require significant attention to module reliability, repairability, and on demand resupply. Modules must be more reliable to reduce inventory requirements; they must be repairable to minimize the effect of the inevitable stockouts; and rapid resupply of nonrepairable modules must be available on a routine basis. This is not a matter of opinion or engineering judgement. It is implicit in the mathematics.

Learning to do this is a subgoal of the Space Station project, which is (in large part) intended to provide experience in remote site base operation. Proper attention to support factors of the same general type as those considered in this report will result in a Space Station which can be supported, and which will require less ground support as it matures. The lessons learned in Space Station operation will later be applied to design of a Lunar Base, and lessons from the Lunar Base, in turn, applied to a Manned Mars Mission. In this way, logistics problems will be solved incrementally, in a timely and effective manner, and the manned exploration of the Solar System will continue at a steady pace, as foreseen in the Report of the President's Commission on Space and in the Ride report.

Further Work

First, analysis (3) should be carried out more rigorously, providing some convenient way of predicting behavior inventory levels for long missions when those for short missions are known.

The repair hour hazards developed above are really averages. Repair hour hazard is a function of time; it will be higher towards the end of the mission, when inventory is depleted, then it will be at mission start. The only feasible ways of determining the repair hour hazard distribution appears to be simulation. In extreme cases, the repair activity might absorb all available crew labor towards mission end. It would be useful to predict and avoid that.

The existence of high failure rate parts should be investigated further. Preliminary examination suggests that there are surprisingly few such parts in the STOCK2.STK data set, and that reliability improvement would require significant MTBF increase over many part types. For that matter, the effect of increasing MTBF is not properly considered above. It will, obviously, decrease inventory size, but will probably do so nonlinearly. It would be interesting to attempt to characterize the effect in the same way that commonality has been considered above, so as to determine the probable contribution of increased MTBF to reducing the logistics

problem.

STOCK2 could be extended to accept more than 400 data classes. It would be interesting to process the entire data set and compare the results with those of Subset A.

Application of STOCK2 to data sets from other equipment, such as planetary probes and modern commercial aircraft, would provide some badly needed perspective on inventory policy. Data from Antarctic bases, extended Oceanographic expeditions, and industrial plants in countries without a local industrial infrastructure might provide a look at successful inventory policy. Application of STOCK2 to the Hubble Space Telescope might prove particularly interesting.

Failure while in inventory has not been explicitly treated, as it is believed to be of negligible importance. If it is important, some fraction of the modules in inventory must be added to the population of modules supported. For example, if MTBF in inventory is twice that in service, than half of the modules in inventory must be counted as in service. This operation produces a convergent series, and one would expect the convergence to be rapid. Note also that if some of the modules are in active standby, possibly with automatic switchover, it may be desirable to count them as inventory rather than as running. The effect is to reduce the number of "running" modules to the minimum required for station operation. The risk, of course, is that one may find the inventory exhausted with only one "necessary" module in a redundant system running and half the mission to go. Obviously, some difficult and real decisions about what is "necessary" are called for before any such reduction can be considered.

Specified probability of stockout is the only independent variable in the current algorithm. It would be convenient for some kinds of analyses to modify it so as to accept number in stock as the independent variable, and in others to modify it so as to accept number of stockouts as the independent variable.

Optimization of the inventory under the following conditions would be of interest, and has not been addressed above:
*   Partial resupply (perhaps by small ELV). Total resupply is treated above as a shortened mission.
*   Constrained availability of Criticality 1 modules.
*   Repairability of only a subset of module types.

Characteristics of the inventory of replacement components used in repair have not been considered.

In addition to the above, there are related questions concerning ground based logistics, which must have the modules needed for restock available when needed, to include rapid resupply missions. There is also the question of proper techniques, tools, inventory, and so on, used for on orbit module repair. The scope of these questions is beyond what any one person could accomplish, and are suggested as tasks for an organization.

Acknowledgements

Bibliography

Data used is Rockwell International's Logistic Master Control File Baseline Report, also called Report 9100, dated 07/13/87. Permission to see the original data may be requested from the logistics branch of Rockwell's plant in Downey, CA.