

AD A117089

Palm's Theorem for Nonstationary Processes

G. B. Crawford

DTIC FILE COPY

DTIC
ELECTE
JUL 20 1982

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

Rand

82 07 19 168

This publication was supported by The Rand Corporation as part of its program of public service.

Library of Congress Cataloging in Publication Data

Crawford, Gordon, 1936-
Palm's theorem for nonstationary processes.

"R-2750-RC"

Bibliography: p.

1. Inventory control--Mathematical models.
2. United States. Air Force--Procurement--Mathematical models.
3. Poisson processes. I. Rand Corporation. II. Title.

TS160.C7 358.4'18'068 81-15840
ISBN 0-8330-0359-3 AACR2

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

20.1 R-2750-RC

Palm's Theorem for Nonstationary Processes

G. B. Crawford

October 1981

DTIC
ELECTE
JUL 20 1982
S D H

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

Rand
SANTA MONICA, CA. 90406

PREFACE

Many of the current stock requirements and capability assessment models used by the Air Force have been influenced by logistics research carried out at The Rand Corporation under the auspices of Project AIR FORCE. This report summarizes a part of this research that substantially broadened the applicability of many of the classical requirements models. Although the research, which dates back to 1976, could significantly affect the accuracy and applicability of all similar models, it has not been widely used outside of Rand and the Air Force. The generalizations provide a basis for precise calculations of requirements and capability during the abrupt transition from peacetime to wartime. The calculations avoid the misleading assumption that the transition may be approximated by a post-hostilities steady-state solution.

The dynamic modeling techniques used have been applied by D. B. Rice in the *Defense Resource Management Study* for the Office of the Secretary of Defense, February 1979, and by J. A. Muckstadt, *Comparative Adequacy of Steady-State Versus Dynamic Models for Calculating Stockage Requirements*, The Rand Corporation, R-2636-AF, November 1980. A proof of the principal result was given by R. J. Hillestadt and M. J. Carillo, "Models and Techniques for Recoverable Item Stockage When Demand and the Repair Process Are Nonstationary--Part I: Performance Measurements," The Rand Corporation, N-1482-AF, May 1980.

This report gives a comprehensive treatment of these dynamic modeling techniques. It is intended to provide the layman with the necessary background to judge when the procedures are appropriate. In addition, the report presents a rigorous proof of a general form of the underlying theorem as well as statements and proofs of several useful results that have not appeared in the literature.

Preparation of this report was sponsored by The Rand Corporation from corporate funds.

SUMMARY

Like most models for calculating stock requirements, the models used by the Air Force to calculate requirements and allocations have traditionally assumed that the failure process generates arrivals approximating a steady-state Poisson arrival process. Although many real-world arrival processes are approximately Poisson, few exhibit steady-state behavior in the long run.

For example, a tactical NATO war can be expected to give rise to a transition from peacetime to wartime flying levels that cannot be captured with steady state models. Two properties characterize such a scenario: The escalation of hostilities will be abrupt, and the success of either side in the first few days of the war may be irreversible.

To understand our capabilities and requirements in real-world scenarios, we must model the transition from peacetime to wartime failure rates and know the distribution of the number of parts in the repair pipeline. Only then can we hope to assess the effects of our limited supplies of replacement parts on our ability to fly the demanding schedules expected in the early days of a war.

Classically these models have used a steady-state result known as Palm's Theorem to model the number of spare parts in the repair pipeline. The research reported here generalizes this theorem, which allows the precise calculation of the distribution of the number of parts in the repair pipeline at any time during a time-varying or dynamic scenario, which may include abrupt transitions in the level of activity. The report is intended to be readable with only a lay knowledge of probability and statistics. The appendix contains technical proof of the general dynamic form of Palm's Theorem.



Session For	☒	☐	☐
GRA&I			
TAB			
Source			
Classification			
Distribution			
Publication			
Availability			
Special			

Per file

ACKNOWLEDGMENTS

The author gratefully acknowledges his indebtedness to several people who made significant contributions. Cindy Williams and Allan Abrahamse contributed to keeping the mathematics honest as well as understandable. Helen Turin undertook the Herculean task of making it all readable. Their combined enthusiasm kept the process of publication as interesting as the research.

CONTENTS

PREFACE	iii
SUMMARY	v
ACKNOWLEDGMENTS	vii
Section	
I. INTRODUCTION	1
II. PALM'S THEOREM	4
Palm's Theorem	5
The Dynamic Form of Palm's Theorem	8
III. APPLICATIONS	10
Example 1. Fire Department Requirements	11
Example 2. People Waiting at a Bus Stop	12
Example 3. Air Force Stock Requirements	13
IV. STATIONARY SURVIVAL TIME DISTRIBUTIONS	16
The Caveats--Probable Perturbations in War Plans	25
Lack of Sensitivity to the Choice of Repair Time Distribution	27
Appendix: NONHOMOGENEOUS FORM OF PALM'S THEOREM	31
BIBLIOGRAPHY	37

I. INTRODUCTION

Many traditional approaches to inventory problems use a classical theorem for steady-state Poisson processes that appeared in the literature in 1943 in an article by C. Palm. The theorem has received wide application. It first described the number of telephone exchanges in use (thereby providing a mathematical approach to deciding how many exchanges are needed). Section III gives several examples of its applicability.

The theorem owes its widespread use to the popularity of a Poisson model to describe arrival processes. This popularity is deserved: There are good empirical and theoretical reasons (Section III) why a Poisson process may closely approximate real-world arrival processes. (Additionally, Poisson models have a lot of very nice mathematical properties.)

Although close approximations to a Poisson process are common, real-world arrival processes that exhibit steady-state behavior in the long run are rare. Approximations to fit dynamic, or nonsteady-state, arrival processes into this steady-state model have run the gamut from very bad to very clever. In some applications the approximation to steady state may be close enough to be justifiable. In the case of an Air Force faced with the need to make the transition from peacetime to wartime flying it is not.¹

The increased fire power available to most of the world's forces suggests that if hostilities break out between major powers the escalation of flying activity will be abrupt and demanding. Additionally, the complexity of the avionics in our tactical fighter forces has been increasing at a surprising rate. This complexity exacerbates the supply problem: Complex avionics are difficult to repair and expensive to replace. Adding to the expense, inflation in the aerospace industry has been almost twice the national average. In short, the need for an accurate requirements models is real. The cost of using more accurate models is the cost of using slightly more data and

¹See Muckstadt, 1980; and Sec III.

carrying out more complex computations. At a time when computers are becoming increasingly ubiquitous and computational costs are going down, there is little to recommend using simplistic, and often misleading, models.

The research described here is due to several people. In 1976 at Rand, R. J. Hillestadt, T. O. Lippiatt, and D. Gavor (Gavor is a consultant from the Naval Postgraduate School in Monterey) realized that the classical form of Palm's Theorem could be generalized to include dynamic or nonsteady-state arrival processes. At the same time the Operations Analysis Office at Hq Pacific Air Forces (PACAF) was using similar ideas to assess the day-by-day adequacy of stock to support the transition from peacetime to wartime flying in a Korean scenario (Crawford, 1977).

The mathematical ideas were first used extensively at Rand in an unpublished study of the requirements for spare engines by M. B. Berman, T. F. Lippiatt, and R. Sims. The ideas were also given as examples in the "Defense Resource Management Study" (Rice, 1979).

In May 1980, Hillestadt and Carrillo first published a proof of the generalized Palm's Theorem in a Rand publication with a description of its application to stock requirements and performance measures. A similar but less general proof appeared with a good example of the fallacy of using steady-state models in Muckstadt, 1980.

Nonsteady-state arrival processes are at the heart of the DYNAMETRICS/Consolidated Support Capability Measurement System computer model built at Rand by Hillestadt and Carrillo. That model has become the standard readiness assessment model for the Air Force.

Despite the level of interest at Rand and in the Air Force, the mathematical ideas discussed in this report have received limited use elsewhere, and the published research has been somewhat disjointed.

This report provides both an introduction for the layman and a statement and technical proof of the dynamic form of Palm's Theorem. It is intended to help the layman develop a feeling for what is, and what is not, an appropriate application. A general statement of the theorem is given along with less general forms that are easy to use and exact in many applications.

2
Section II explains the importance of the theorem to stock calculations and presents an intuitive proof of the classical form of Palm's Theorem for steady-state arrivals and discrete time. It also states the dynamic form of Palm's Theorem. (The appendix contains a technical proof, which builds on the ideas of the steady-state proof). Section III³ gives several different steady-state examples to provide an understanding of what is, and what is not, a Poisson arrival process and when the dynamic form of Palm's Theorem applies. Several different statements of Palm's Theorem are developed in Section IV⁴ to facilitate its application. These forms are used to prove what has been known as the worst-case approximation theorem. Section IV⁴ ends with a discussion of the application of the dynamic form to the calculation of war readiness spares requirements for aircraft.

The examples and the discussion illustrate that there is little justification for trying to approximate a dynamic arrival process by a steady-state arrival process. The computational difficulties encountered in the dynamic form are minimal, whereas the cost incurred by overestimating or underestimating inventory requirements is often very high.

II. PALM'S THEOREM

Inventory systems for expensive reparable items and retail inventory systems for expensive merchandise are frequently modeled according to an $(S - 1, S)$ inventory model. For reparable parts this model works as follows:

An initial inventory level of S serviceable units is provided. When a failure occurs, the failed unit is put into the "repair cycle." If a replacement unit is available in stock, it is put into service. If not, the consumer must wait until the next serviceable unit comes into stock from the repair cycle. In this process the S units become randomly split between serviceable units in stock and units in the repair cycle. From the consumer's point of view, the system will meet his needs unless he happens to have a failure when all S units are unserviceable.

In this description, "repair cycle" is really a misnomer; some units that enter this cycle may not be reparable. In that case, they are discarded and "repair time" includes the time elapsed from the detection of the failure until the receipt of a serviceable unit from outside sources.

In the case of a retail inventory system for expensive merchandise the scheme has a parallel structure. An initial inventory of S units is assumed. When a unit is sold, a replacement is ordered from the supplier, and the random order and shipping time replaces the repair cycle time.

It is assumed that the user of such a scheme is concerned with balancing two complementary costs:

(1) The indeterminate, and often intangible, cost of a failure or loss of potential sale when shelf stock has dropped to zero. (The cost in this case may be quite high--the cost of an idle airliner awaiting a serviceable replacement for a vital part, or of a missed sale in the retail system.)

(2) The cost of providing a sufficiently large number S of shelf units to provide adequate protection against "stock-outs."

For a given choice of S , the complete probability distribution of the number of orders that occur when the stock level is zero depends on the probability distribution of the total number of items in the "repair cycle" (Lu, Brooks, and Gillen, 1969). Thus the problem of determining an appropriate stock level S hinges on finding the distribution of the number of items in the repair cycle at a given point in time.

This report addresses the problem of solving for the distribution of the number of items in the repair cycle. To facilitate finding the distribution of this random quantity, it is commonly assumed that the failure (sales) and repair (order) processes satisfy the conditions of a theorem due to C. Palm (1943). This theorem (cf. Feller, 1957, and Takacs, 1962, for different statements and proofs) originally described the number of telephone trunk lines in use; in keeping with a more general application, this report will consider a failure or sale as an "arrival" (a failure generates the arrival of a unit at the repair cycle) and the repair time or order time as the "survival time" for that arrival. Therefore the discussion will be of the number of survivors instead of the number of units in the repair cycle at time T .

In these terms, Palm's Theorem states that if the arrival-survival process satisfies the following conditions, the number of survivors at any given time is a Poisson random variable (r.v.) whose mean is given by the formula below.

PALM'S THEOREM

If:

- (1) The number of arrivals at any time interval $\{t: r < t \leq s\}$ is a Poisson r.v. with mean $(s - r)\eta$, and
- (2) The probability that an arrival at time t survives until¹ time T is given by $\bar{F}(T - t)$, and
- (3) The survival times are independent of one another and independent of the arrival process, then the number of survivors at any time T is a Poisson r.v. with mean λ ,

¹Here, and throughout, until T means *at least* until time T .

$$\lambda = n\mu,$$

where μ is the mean survival time,

$$\mu = \int_0^{\infty} \bar{F}(s) ds.$$

It is not necessary to give a complete proof of this form of Palm's Theorem because it follows immediately from the dynamic or nonstationary form given below and proved in the appendix.

Feller proves the stationary form of this theorem by assuming that the survival time is an exponential r.v. and solving certain differential equations that can be inferred from the assumptions. Takacs and most authors use a moment generating function proof that has elegance but may leave the reader with little intuitive feel for what is happening.

To help build an understanding for Palm's Theorem, both the classical stationary version stated above the nonstationary version given below, a discrete time version of the stationary form is presented.

Assume that the number of arrivals on day 1 is a Poisson r.v. with mean λ and that the probability that an arrival on day 1 survives until day k is $\bar{F}(k-1)$. (Additionally, assume all of the independence inherent in conditions (1), (2), and (3) above.) In this case, the number of survivors on day k is a Poisson r.v. with mean $\lambda\mu$, where

$$\mu = \sum_{i=0}^{\infty} \bar{F}(i)$$

is the mean of the survival time.

Proof: Consider the distribution of the number N_1 of arrivals on day 1 that survive until day k :

The probability of m arrivals on day 1 is given by the Poisson density function

$$\frac{e^{-\lambda} \lambda^m}{m!},$$

and the probability that any one of these arrivals survives until day k is given by

$$p = \bar{F}(k - 1).$$

Thus,

$$\begin{aligned} \Pr(N_1 = n) &= \sum_{m=n}^{\infty} \frac{e^{-\lambda} \lambda^m}{m!} \binom{m}{n} p^n (1-p)^{m-n} \\ &= \frac{e^{-\lambda} p^n \lambda^n}{n!} \sum_{m=n}^{\infty} \frac{\lambda^{m-n} (1-p)^{m-n}}{(m-n)!} \\ &= \frac{e^{-p\lambda} (p\lambda)^n}{n!}. \end{aligned}$$

This calculation is a restatement of the simple and well-known result (Feller, 1957) that if a Poisson process is randomly censored then the censored process is Poisson also. This simple property is key to the steady-state and dynamic form of Palm's Theorem.

It has been shown that the number of arrivals on day 1 that survive until day k is Poisson with mean $\lambda \bar{F}(k - 1)$. Adding over day k and all previous days it follows (because the sum of independent Poisson random variables is Poisson) that the number of survivors is a Poisson r.v. with mean

$$\sum_{i=-\infty}^k \lambda \bar{F}(k - i) = \lambda \sum_{i=0}^{\infty} \bar{F}(i) = \lambda \mu,$$

where μ is the mean survival time.

This simple proof, although given here for the discrete time case, will work whenever one can break up the time before day k into periods where $\bar{F}(k - t)$ is constant. The hypothesized stationarity of the arrival and survival processes are only incidentally useful in that they permit a simple notation. (The dynamic form of Palm's Theorem and its proof given in the appendix builds on these observations.)

Note also that the result does not depend on the form of $\bar{F}$, but only on μ , the mean survival time. This useful property does not carry over to the nonstationary form. In cases where μ is known but $\bar{F}$ is not known, the worst case approximation theorem given in Sec. IV may be used. Section V gives empirical evidence that if the mean and variance of the survival process are known, the actual choice of an $\bar{F}$ that has these first two moments is unimportant.

THE DYNAMIC FORM OF PALM'S THEOREM

Throughout the remainder of this report the following assumptions are made regarding the arrival and survival process.

(1) The arrival process is Poisson with mean function $M(\cdot)$; that is, the number of arrivals $X(k) = t$, $r < t \leq s$, is a Poisson r.v. with mean $M(s) - M(r)$.²

(2) If $X(k) = t$, the probability that $X(k)$ survives until time T is given by $\bar{F}(t, T)$, and the survival process is independent of $\{X(n), n \neq k\}$. If $T < t$, $\bar{F}(t, T) = 0$.

(3) The survival time for $X(k)$ is independent of the survival times for $\{X(n), n \neq k\}$.

The sequel abbreviates these conditions by saying that the arrival process is Poisson with mean function $M(\cdot)$, and the survival process is independent of the arrival process.

If the above conditions (1) through (3) hold, and if for fixed T the function $\bar{F}(t, T)$ is a measurable³ function of t , then the number of survivors at time T is a Poisson random variable with mean $\lambda(T)$,

$$\lambda(T) = \int \bar{F}(t, T) dM(t). \quad (2.1)$$

The added complexity of the Stieltjes integral (Halmos, 1950) above has merit: whereas the steady-state version of Palm's Theorem

²This definition of a Poisson process differs somewhat from the usual. See the appendix for a discussion.

³This is about the weakest assumption regarding F that can be made and have the integral (2.1) make sense. It is satisfied, for instance, if $F(\cdot, T)$ has no more than a countable number of discontinuities.

was stated for continuous time and then proved for discrete time, the above formulation and the proof given in the appendix hold for both cases as well as combinations. In the discrete time case where $M(\cdot)$ has all its mass at lattice points (the number of arrivals on day i is a Poisson random variable with mean $m(i) = M(i) - M(i - 1)$), and the probability that an arrival on day i survives until day k is given by $\bar{F}(i,k)$, then (2.1) is equivalent to

$$\lambda(k) = \sum_{i \leq k} \bar{F}(i,k)m(i). \quad (2.2)$$

In the continuous time case, if $M(\cdot)$ can be written as an integral:

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

then (2.1) reduces to

$$M(t) = \int_{-\infty}^T \bar{F}(t,T)m(t) dt. \quad (2.3)$$

The proof of this form of Palm's Theorem follows the lines of the proof given above; one divides the time before T up into periods wherein $\bar{F}(\cdot, T)$ is "almost" constant and sums over these time periods. Then one uses the simple result that a censored Poisson process is Poisson and the sum of Poisson random variables is a Poisson r.v.

III. APPLICATIONS

Because the assumption of a Poisson arrival process is fundamental to any application of Palm's Theorem, it is worthwhile to develop some notions regarding a Poisson process and when it is and is not a good approximation of an empirical arrival process.

Researchers are often concerned with some group of entities (people, aircraft, etc.), each of which may give rise to some event of interest (make a telephone call, have a radio failure, etc.) in each time interval. If the entities are numbered 1, 2, 3, ..., n and associated with the i th entity is an indicator random variable $x(i)$ equal to one or zero according to whether that entity gave rise to the event in some fixed time interval, then the number of events or number of arrivals is given by $y = \sum x(i)$. Suppose that $\Pr\{x(i) = 1\} = p(i)$. If the entities act independently and all the $p(i)$ are equal to some one value p , y has a binomial distribution.

If n is fairly large and p is small, the Poisson distribution with mean np provides a very good approximation to the distribution of y (Feller, 1957). It is also well known (although infrequently stated) that this result is true even if the $p(i)$ are not the same, provided that they are uniformly small. In that case the mean of the Poisson distribution is $\sum p(i)$.¹ Thus the Poisson distribution is a good approximation to an arrival process generated by a collection of entities acting independently of one another, each with a small probability of generating an event in a given short time interval.

Because complete independence is the exception and not the rule in nature, these results might seem more academic than useful. However, from Paul Levy's work on infinitely divisible distributions (Feller, 1957), any arrival process where the number of arrivals in disjoint time intervals is independent is a process of the Poisson type. Therefore, the Poisson arrival process will be a good approxi-

¹For a proof, compute the log of the characteristic function of y and take limits.

unless there is a dependence among entities that causes the number of arrivals in disjoint time intervals to be dependent. For an example of this dependence, consider a squadron of aircraft that happens to be low on spare radios. If a rash of broken radios occurs, the subsequent time interval will see a drop in radio failures (there are fewer radios left to fail). The number of entities must therefore be large enough with respect to the $p(i)$ that an upswing in arrivals does not produce a downswing in the following time intervals.

Within the framework just discussed let us look at several examples of Poisson arrival processes and the application of Palm's Theorem for those processes.

EXAMPLE 1. FIRE DEPARTMENT REQUIREMENTS

The stochastic arrival of fire alarms in a metropolitan area is a good example of many entities independently giving rise to low probability events. An examination of New York City alarm rates (Carter and Rolph, 1975) has shown that this process is distinctly non-stationary: June has four times as many alarms as February, and the hour from 8 to 9 p.m. has ten times as many alarms as the hour from 2 to 3 a.m. Day-of-the-week differences are also significant. Carter and Rolph develop estimates of alarm rates as a function of time using methods of time series analysis and linear regression, and they use this time-varying intensity function and a Poisson arrival process to predict the number of alarms.

Suppose we wish to determine the size of the Fire Department needed at some future date. The work of Carter and Rolph provides the procedures for estimating the number of alarms. If that number is coupled with an estimate of the distribution of the random time required to answer an alarm, Palm's Theorem may be used to estimate the time-varying mean number of fire engines and crews needed. In this model the number of engines and crews actually needed will be a Poisson r.v. For any given level of funding, therefore, one can calculate the probability of not being able to cover all alarms during peak periods and use this criterion to calculate requirements. In this

application the assumption of Poisson arrivals and of independence between the arrival process and the time-to-answer process is justified.

EXAMPLE 2. PEOPLE WAITING AT A BUS STOP

Excluding early morning and late afternoon times when the number of people catching a given bus may be highly dependent from one day to the next, arrivals at a bus stop fit the description of a number of low probability events generated by a large number of independent entities. The process is nonstationary with peaks before the scheduled arrival of buses. "Survival time" is taken to be the amount of time spent waiting for the bus. If bus schedules are accurate, then given the time of the n th arrival, its survival time is essentially a degenerate random variable. However, this survival time is independent of all other arrivals, and that is all that is required by the dynamic form of Palm's Theorem. Therefore, with Palm's Theorem, the number of people waiting at a bus stop is a Poisson r.v. whose mean is given by:

$$\lambda(T) = \int \bar{F}(s, T) M(s) ds.$$

If buses adhere to their schedule without deviation and these schedules are universally known, but walking time to the bus stop is imperfectly known, this example results in an intensity function $m(t)$ with smooth peaks. Then Palm's Theorem is still applicable and the usual result is that the number of passengers waiting at any given time T is a Poisson r.v. with mean $\lambda(T)$ although both the arrival process and the survival process are highly dependent on a nonrandom external event. In that case, $\lambda(T)$ will have zeroes at the scheduled arrival times. These zeros follow peaks immediately prior to scheduled arrival times.

This example also illustrates a result that is sometimes misunderstood: Although the number of passengers waiting for the bus at time T is a Poisson r.v. with mean $\lambda(T)$, this stochastic process (number of passengers waiting at T) is *not* a Poisson (arrival) process.

Depending on the model under consideration, the size of the population at different times T_1 and T_2 may be independent or may be highly correlated. Accordingly, Palm's Theorem provides a tool to calculate the distribution of the size of the surviving population at a given time T_1 but does not, per se, say anything about the joint distribution of the populations, at two different times T_1 and T_2 .

EXAMPLE 3. AIR FORCE STOCK REQUIREMENTS

Muckstadt (1980) gives an example of a dynamic schedule where flying initially gives rise to a steady-state Poisson demand for aircraft engines of 0.8 units per day. Upon the onset of hostilities, this rate jumps to 3.16 and then tapers off exponentially, resulting in a 30 day average of 1.0 unit per day. Throughout it is assumed that repair takes five days. For this process the standard Air Force calculation correctly computes a stock requirement for peacetime operations assuming a steady state arrival rate of 0.8 per day. With the steady state form of Palm's Theorem this implies that the number of nonserviceable parts in the repair process will be a Poisson r.v. with mean

$$\lambda = (.8 \text{ units per day}) \times (5 \text{ days}) = 4.$$

To compute the stock level for spare engines, the Air Force goal is to provide sufficient stock to cover the number of units in repair with probability of at least 0.8--that is, to provide s units of shelf stock where s is the smallest integer such that

$$\sum_{k=0}^s \frac{e^{-\lambda} \lambda^k}{k!} > .8$$

For $\lambda = 4$ this calculation yields a requirement for 7 spare units. The standard Air Force calculation for the first 30 days of hostilities follows the same policy; it assumes steady-state demands at the 30 day average rate and provides enough to cover the number of units

in repair with probability 0.8. This policy requires eight units total to cover the period of hostilities.

The arrival rate after the onset of hostilities will clearly be nonstationary. Table 1 uses Muckstadt's results and compares the standard calculation with one that uses the same policy (guards against stock outages with probability of at least 0.8) but uses the postulated dynamic arrival rate and assumes several different forms for the repair time distribution. In cols. (3) and (6) times are assumed to be always five days. In columns 4 and 7 we have assumed that repair times are exponential random variables with a mean of five days. Columns 5 and 8 make the same assumption and further assume that there was a standdown before hostilities resulting in the completed repair of all units before the escalation in flying.

Several things in this comparison are interesting. As expected, the steady-state result significantly underestimates the stock needed. From the steady-state form of Palm's Theorem, the calculation doesn't depend on the shape of the distribution or repair times; however, the dynamic calculations clearly do. The assumption of constant repair times results in a higher number in repair during the surge, but that number drops sharply following the surge in flying. This phenomenon is examined in Sec. IV.

A standdown before the surge does not appreciably affect the maximum number of units required, but it does cause the requirement to rise more slowly in the first few days. The standdown may thereby provide a small hedge against the time required to ship additional stock or to reconstitute battle damaged maintenance facilities. Although the standdown can mitigate the stock requirements slightly, a steady-state approximation may grossly underestimate the number of parts required to support a surge.

Table 1

COMPARISON OF STEADY STATE AND DYNAMIC STOCK CALCULATIONS

<u>Steady State Calculation</u>							
Day	Expected Demands	Expected Number in Repair		Required Stock Levels			
Before Day 1	.8	4.		7			
Days 1 to 30	1.0	5.		8			

<u>Dynamic Calculation</u>							
		<u>Expected Number in Repair</u>			<u>Required Stock Levels</u>		
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Day	Expected Demand	Degenerate Repair Times	Exponential Repair Times	Exponential Repair with Standdown	Degenerate Repair Times	Exponential Repair Times	Exponential Repair with Standdown
Before Day 1	.8 or 0	4.0	4.0	0.	7	7	0
1	3.0	6.2	6.0	2.7	9	9	5
2	2.7	8.1	7.4	4.7	11	11	7
3	2.5	9.8	8.3	6.1	12	12	9
4	2.2	11.2	8.8	7.0	14	12	10
5	2.0	12.4	9.0	7.5	15	12	11
6	1.8	11.2	9.0	7.8	14	12	11
7	1.7	10.2	8.9	7.9	13	12	11
8	1.5	9.2	8.6	7.8	12	12	11
9	1.4	8.3	8.3	7.6	11	12	11
10	1.2	7.5	7.9	7.3	10	11	11
15	.7	4.6	5.7	5.5	6	9	8
20	.5	2.8	3.8	3.7	4	6	6
25	.3	1.7	2.4	2.4	3	5	5
30	.2	1.0	1.5	1.5	2	3	3

IV. STATIONARY SURVIVAL TIME DISTRIBUTIONS

Tools and heuristics are developed here to facilitate the application of Palm's Theorem when the user has incomplete information about the mean arrival function $M(t)$ and the survival distribution $\bar{F}(t, T)$. A "worst case" approximation theorem is proved, showing that in many inventory problems of practical interest $\bar{F}$ may be approximated by a degenerate repair time distribution (that is, by assuming the repair time is constant) with conservative results. Following a peak in the demand rate, the degenerate distribution gives an upper bound on the expected number of units in repair.

To simplify notation, we impose the following restrictions on M and $\bar{F}$.

The survival time distribution is stationary, that is:

$$\bar{F}(t, T) = \bar{F}(T - t), \quad (4.1)$$

and the mean demand function $M(t)$ is equal to the integral of its derivative:

$$M(t) = \int_{-\infty}^t m(s) ds. \quad (4.2)$$

Under these conditions Eq. (2.1) can be usefully rewritten in several different ways. From (2.1),

$$\lambda(T) = \int_{-\infty}^T \bar{F}(t, T) dM(t),$$

$$\lambda(T) = M(T) - \int_{-\infty}^T F(T - t) dM(t), \quad (4.3)$$

$$\lambda(T) = M(T) - F * M(T), \quad (4.4)$$

where "*" denotes the functional convolution of F and M; since convolution is commutative,

$$\lambda(T) = M(T) - \int_{-\infty}^T M(T-t)dF(t). \quad (4.5)$$

If R denotes the random survival time, then (4.5) may be rewritten

$$\lambda(T) = M(T) - \text{Ex}[M(T-R)], \quad (4.6)$$

where Ex denotes the expected value with respect to the distribution F.

Using the density $m(s)$, (4.3) may be written

$$\lambda(T) = M(T) - \int_{-\infty}^T F(T-t)m(t)dt. \quad (4.7)$$

Equations (4.3), (4.4) and (4.5), expressing $\lambda(T)$ in terms of the convolution of M and F, are useful from the spectral theory point of view. They tell us, for instance, how the frequency components or power spectra of λ compare with the spectra of M and F and provide a way of quantifying the extent of changes in λ resulting from changes in M. Equation (4.6) has a very intuitive interpretation: $\text{Ex}[M(T-R)]$ can be interpreted as the expected number of arrivals that are no longer surviving; thus (4.6) states that the expected number of survivors at time T is the expected number of arrivals before T minus the expected number of deaths before T.

Equation (4.7) enables us to look at some special survival time distributions. Suppose R is degenerate at 10—i.e., survival time is constant and equals 10. Then,

$$\bar{F}(t) = \begin{cases} 1, & 0 < t < 10 \\ 0, & \text{otherwise.} \end{cases}$$

and suppose that $m(t)$ ¹ has a graph that looks like:




Fig. 1--The intensity function $m(t)$

Then $\lambda(t)$ evaluated at, say, $t = 12$ is given by:

$$\begin{aligned}\lambda(12) &= \int_0^{12} \bar{F}(12 - x)m(x)dx \\ &= \int_2^{12} m(x)dx\end{aligned}$$

$\lambda(12)$ is what we would see if we averaged the $m(t)$ curve over the window $\bar{F}$. That is, following the definition of a Rieman integral, we would vertically slice these curves and multiply their coordinates,

¹This same $m(t)$ --composed of a flat section, a sinusoidal section, and a section with exponential decay--has been used in Figs. 1-8.

then sum the products over the window $\bar{F}$. For $\lambda(27)$, the right-hand side of the window slides out to 27.

Consider another window. Suppose survival time is an exponential r.v. with mean 10:

$$F(t) = 1 - e^{-t/10}, \quad 0 \leq t,$$




Fig. 2--The exponential distribution function

Hence, for fixed $T = 22$, the window $\bar{F}(22_0 - x)$ looks like




Fig. 3--The window $F(22-t)$

and using the same $m(t)$




Fig. 4-- $\lambda(T)$ Corresponding to an exponential repair time having mean and standard deviation equal to 10 days

The mean size of the population at time T_0 , $\lambda(T_0)$, is given by

$$\lambda(T_0) = \int_0^{T_0} \bar{F}(T_0 - x)m(x)dx$$

and can be visualized as averaging $m(t)$ over the window $\bar{F}$. For the intensity function m given in Fig. 1 and exponential survival times, $\lambda(T)$ is given in Fig. 4.

In the case where $\bar{F}$ was degenerate, the area under $\bar{F}$ —that is, the "size" of the window—was equal to μ . Integrating by parts gives the well-known result that this is always the case for nonnegative random variables:

$$\text{Area under } \bar{F} = \int_0^{\infty} \bar{F}(t) dt = \int_0^{\infty} [1 - F(t)] dt = \mu.$$

Thus, if the mean lifetime increases, $m(t)$ is seen through a larger window and at any point the mean of the surviving population will increase accordingly.

All windows corresponding to different lifetime distributions having the same mean have the same "size," but clearly they don't all have the same shape. Furthermore, windows that are "long and low" will result in more averaging of $m(t)$ than shorter windows of the same size; as a result, for a fixed arrival density $m(t)$, they will give a smoother mean function $\lambda(t)$. Because $\bar{F}(t) = 1 - F(t)$, $\bar{F}$ will be "long and low" if $F(t)$ does not quickly converge to 1 as t gets large. Thus, the degree of smoothing that we get when we calculate $\lambda(t)$ from the arrival density $m(t)$ is largely a function of the size of the tail of the distribution F , and to a degree this can be characterized by the second moment or variance σ^2 of F .

In the sequel, several numerical examples suggest that, from a practical point of view, correctly estimating μ and σ^2 for the survival process is what really matters; the resulting $\lambda(t)$ function is quite insensitive to the choice of a particular F having these first two moments.

Just as increasing the variance σ^2 tends to give increasing smoothness to $\lambda(t)$, if $\sigma^2 = 0$ --that is, the survival time is constant--there should in some sense be a limiting case $\lambda_0(t)$ that is more peaked than the graphs of $\lambda(t)$ functions corresponding to other survival distributions having the same fixed mean R_0 .

Suppose that $\sigma^2 = 0$ --that is, survival times are degenerate and equal to some value R_0 . Then if $m(s)$ has a unimodal peak that can be "straddled" by the window $\bar{F}$, $\lambda_0(t)$ will have a local maximum at a value t_0 such that $m(t_0) = m(t_0 - R_0)$:




Fig. 5--Finding the local maximum of λ_0

(If $m(t_0) \neq m(t_0 - R_0)$, then $\lambda_0(t_0)$ could be marginally increased by sliding the window slightly to the right or left.)

If m is unimodal and λ_0 peaks in the vicinity of t_0 , then $m(t - R_0)$ is increasing for t in a neighborhood of t_0 , and if $m(t - R_0)$ is increasing, then $M(t) = \int_0^t m(s)ds$ is convex for t in the neighborhood of $t_0 - R_0$. If, instead of being degenerate, the lifetime R varies around its mean R_0 but its range is restricted so that $M(t_0 - R)$ is convex throughout the range of R , then using the convexity of M and Jensen's inequality (Chung, 1974, p. 47):

$$\text{Ex } M(t_0 - R) > M(t_0 - R_0).$$

Thus, if $\lambda(t)$ is the mean population size determined by the distribution of R (recall that $\lambda_0(t)$ is the mean population size determined by a degenerate R), (4.6) implies:

$$\lambda_0(t_0) = M(t_0) - M(t_0 - R_0) \quad (4.8)$$

$$> M(t_0) - \text{Ex } M(t_0 - R) = \lambda(t_0).$$

In fact, the inequality (4.8) does not depend on $\lambda_0(t_0)$ being the maximum of λ_0 ; it holds for any t_0 such that $m(t_0 - R)$ is increasing throughout the range of R or, equivalently, such that $M(t_0 - R)$ is convex throughout the range of R . We have proved the worst case approximation theorem.²

WORST CASE APPROXIMATION THEOREM

Let Γ be a class of survival time distribution functions having the same mean R_0 and let $\lambda_0(t)$ correspond to a constant survival time as defined above. If the mean density function $m(t_0 - R)$ is, for some fixed t_0 , increasing throughout the range of R for all F in Γ , then

$$\lambda_0(t_0) > \lambda(t_0) \quad (4.9)$$

for any other λ corresponding to F in Γ .

We have shown that, for values t such that $m(t - R_0)$ is increasing, the $\lambda_0(t)$ given by the degenerate survival time distribution is uniformly larger than other $\lambda(t)$ corresponding to survival time distributions with the same mean and suitably restricted range. For this reason, $\lambda_0(t)$ could be used in a worst case analysis to get an upper bound on the maximum population size if nothing is known about the survival time distribution except its mean.

It also follows by the same arguments that, if $m(t_0 - R)$ is decreasing throughout the range of R , then M is concave and the inequalities reverse:

²Although well understood and applied in Rand publications, this theorem has not previously been formally stated or proved in the literature to the best of my knowledge. An alternative proof of a similar theorem has been given by Donald Gavor of the Naval Postgraduate School, Monterey, California (private communication).

$$\lambda_0(t_0) < \lambda(t_0). \quad (4.10)$$

In other words, following the earlier conjecture that $\lambda_0(t)$ is the most peaked or least smooth of the λ functions corresponding to survival distribution with mean R_0 , subject to restrictions on the range of the life time, following a peak in $m(t)$, $\lambda_0(t)$ will have a higher peak than other $\lambda(t)$ functions corresponding to survival distributions with the same mean. Similarly, following a valley in $m(t)$, $\lambda_0(t)$ will have a lower valley than other $\lambda(t)$ functions. Thus, in keeping with the observation that increasing σ^2 increases the "smoothness" of $\lambda(t)$, the $\lambda_0(t)$ we get when we set $\sigma = 0$ is more peaked than other λ functions corresponding to lifetime distributions having the same mean and restricted range.

If in the range of interest the intensity function $m(\cdot)$ is essentially linear (and hence $M(\cdot)$ is essentially quadratic), then it follows from the Taylor series expansions

$$m(t_0 - R) \approx m(t_0 - R_0) - (R - R_0)m'(t_0 - R_0) + (R - R_0)^2 m''(t_0 - R_0)/2$$

and Eq. (4.6) that

$$\lambda(t_0) - \lambda_0(t_0) \approx \sigma^2 m''(t_0 - R_0)/2 \quad (4.11)$$

In this case, (4.9) or (4.10) holds, depending on whether $m(\cdot)$ is increasing or decreasing. Moreover, in this case, the error involved in assuming that repair times are constant (Eq. (4.11)) is equal to a constant multiple times the actual variance of the repair times.

Returning to the Air Force inventory system, it is unfortunate that the use of the function $\lambda(t)$ to forecast spares requirements has been slow to receive widespread acceptance. In part, this may be the result of a concern that the computation is inextricably linked with a plan or scenario that is in reality likely to be wrong, hence the results of the computation will be misleading.

This gives rise to several similar questions that I address in some detail:

(1) If spares requirements are computed from one scenario, but the war dictates another, will spares be adequate?

(2) Conversely, suppose that spares requirements are computed from one scenario, but the war will probably be fought differently. Will the errant calculation require an excessive and unnecessary expenditure?

(3) If several different war plans are used for planning, will they differ greatly in their spares requirements?

The answer to all three questions, with suitable caveats, is No. Moreover, if there are large differences--if for instance Operations or Planning is working from a war plan that is totally unreasonable from the point of view of Supply--then the lack of reality of such a plan should be understood and the plan, or the supplies available, should be altered as necessary.

THE CAVEATS--PROBABLE PERTURBATIONS IN WAR PLANS

Generally speaking, war plans are statements of goals that provide direction to the relevant units of a major command. As goals, these plans should not be unrealistically demanding, nor should they demand much less than can be achieved. Given this generalization, it follows that in war a unit will probably not fly significantly more sorties over an extended period of time than called for in the war plans formulated by that major command.

It is even less likely that the sum of the sorties flown by the units in a theatre over an extended period of time will significantly exceed the number that is called for in the plans. Several factors other than supply augur against high sortie rates for extended periods of time. It is to be expected that the early days of high flying activity will see high attrition and battle damage to aircraft. The resulting reduction in the number of Mission Capable (MC) aircraft will inhibit subsequent sortie production. In addition, maintenance personnel will be pressed to generate large numbers of sorties and the MC rate may subsequently suffer from lack of personnel, even if a generous supply of parts is available.

None of these constraints will prevent short spurts of flying activity from exceeding the war plan. But how will these short spurts influence the spare parts requirements? As shown above, the smoothing induced by looking at the curve of expected demands through an "averaging window" minimizes the effect of these spurts on the parts requirements, especially if the surge is short compared with the size of the window or, equivalently, compared with the mean length of time required to repair or resupply. In summary, excursions from the war plan will generally not cause significant shortfalls in the spares requirements.

Will a dynamic requirements calculation cause substantial overstocking if flying levels fall below the war plan? Not if the "shortfall" in flying is short-lived. In this case, as in the case of short-lived surges, the smoothing achieved by looking at these spurts through the averaging window minimizes the effect of these perturbations. If the shortfall in flying is not short-lived, requirements for that unit will be overstated. If the loss in sorties is made up by other bases flying the same type of aircraft, the surplus of spare parts may be needed and used at those bases. If a long-term loss in sorties is not made up, the spares requirements *will be* overstated, but this "flaw" is not peculiar to the dynamic calculation; the same problem exists with any spares requirements calculation.

What plan should be used if there are several conflicting war plans? There are possible differences in plans outlined above that will cause inconsequential differences in spares requirements. If the plans have major differences in their effect on the spares requirements, then the planners should know and understand that, and it should not be obscured by calculating requirements on the basis of steady-state calculations only remotely related to the war plan. Different plans being considered should be evaluated for their effects on supply. If the differences are small, the choice is unimportant, from the point of view of supply. If the differences are substantial and sufficient spares are not available to support an ambitious plan, then the command involved should understand that.

LACK OF SENSITIVITY TO THE CHOICE OF REPAIR TIME DISTRIBUTION

In the steady-state calculation less information is needed about repair times. Palm's Theorem says that we do not need to know the distribution of repair times, only the mean repair time. It was mentioned that the dynamic spares calculation is sensitive to the size of the averaging window--that is, to the mean repair time. The steady-state calculation is no less sensitive to changes in the mean repair time.

The difference between the steady state and dynamic models is in the sensitivity to the higher moments. The steady-state calculation is insensitive to (in fact independent of) all higher moments. the degree of smoothing in the dynamic calculation is influenced by the size of the tail of the repair time distribution. In Figs. 4, 6, and 7 I have plotted the mean of the number of units in the repair pipeline, by day, using the same flying schedule but three very different repair time distributions having the same mean and the same variance. From the point of view of requirements calculations, the curves are nearly the same. The peaks occur at the same time, and the up-slope parts of the curves that dictate when the extra spares may be needed are almost identical. The large differences in the tails of the distributions cause differences in the curves in the days after the peak, but they would not affect a requirements calculation. For comparison the more peaked $\lambda_0(t)$ corresponding to constant 10 day repair times has been plotted in Fig. 8.

Such examples lead us to believe that if it is possible to estimate the mean and variance of the repair time distribution, it is also possible to use just about any distribution having these moments and the results will be the same.

Fortunately, with the data currently available in the Air Force supply system, we can estimate the repair time distribution for every reparable part and in so doing closely estimate the repair time variance. To estimate the repair time distribution, recall that the current system provides percent base repair (PBR), repair cycle time (RCT), and order ship time (OST) for all items. If this information is used and repair time is assumed equal to RCT with probability PBR a




Fig. 6-- $\lambda(T)$ Corresponding to a discrete repair time taking two values and having mean and standard deviation equal to 10 days (Repair time equals 0 or 20 days with probabilities $1/2, 1/2$)




Fig. 7-- $\lambda(T)$ Corresponding to a discrete repair time taking four values and having mean and standard deviation equal to 10 days (Repair time equals 0, 10, 20, or 30 days with probabilities $3/8, 3/8, 1/8, 1/8$)




Fig. 8-- $\lambda(T)$ Corresponding to a constant repair time of ten days

current system provides percent base repair (PBR), repair cycle time (RCT), and order ship time (OST) for all items. If this information is used and repair time is assumed equal to RCT with probability PBR and to OST with probability $(1 - \text{PBR})$, we have a simple approximation of repair time distribution. In view of the discussion on the degree of smoothness and peakedness, it can be seen that this estimate is slightly conservative--that is, its λ function will be slightly more peaked than if we knew the true repair time distribution.

In summary, minor excursions from the war plans will not invalidate a dynamic requirements calculation. There are unlikely to be any differences between the real and the planned demand pattern that will cause the dynamic calculation to substantially underestimate the parts requirements. Although additional data are needed in the dynamic view (some details of the repair time distribution), these data are currently available in the base supply system in adequate detail.

Appendix

NONHOMOGENEOUS FORM OF PALM'S THEOREM

NOTATION

Let $X(t)$, $t \in (-\infty, +\infty)$ be the random number of arrivals up to and including time t .

Let $M(t) = \text{Ex}(X(t))$, $t \in (-\infty, +\infty)$

Let $A(k)$, $k = 1, 2, \dots$, be the random time of the k th arrival; that is, $A(k) = t$ if $X(t) = k$ and $X(s) < k$ for all $s < t$.

Let $S(k)$, $k = 1, 2, \dots$, be the random survival time of the k th arrival.

Let $\bar{F}(t, T) = \Pr(A(k) + S(k) > T \mid A(k) = t)$ if $T < T$
 $= 0$ otherwise

With this notation we make the following assumptions:

- (1) $M(t)$ is finite for all finite t .
- (2) $\{X(t), t \in (-\infty, +\infty)\}$ is a Poisson process with mean function $M(\cdot)$; that is, the process is separable (Doob, 1959) and if $s < t$ then $X(t) - X(s)$ is a Poisson r.v. with mean $M(t) - M(s)$.
- (3) $S(k)$ is independent of $S(n)$ for $n \neq k$; $n, k = 1, 2, \dots$
- (4) $S(k)$ is independent of $A(n)$ for $n \neq k$; $n, k = 1, 2, \dots$

In the sequel these assumptions will be abbreviated as follows:

$\{X(t)\}$ is a Poisson process with mean function $M(\cdot)$, and the survival process is independent of the arrival process.

NONHOMOGENEOUS FORM OF PALM'S THEOREM

Under the above assumptions if $F(t, T)$ is, for fixed T , measurable in t , then the number of survivors at time T is a Poisson r.v. with mean

$$\lambda(T) = \int \bar{F}(t, T) dM(t).$$

Note that the function $M(\cdot)$ is, by definition, nondecreasing and right continuous. We have not assumed, as is common in the literature,

that $M(\cdot)$ is continuous. The added generality allows us to treat processes where arrivals may be bunched at discrete times as well as continuous time models and combinations. Combinations of continuous time and discrete time processes occur in practical production problems where overnight arrivals are serviced at the beginning of the day shift, and day shift arrivals are serviced in near real-time.

It follows from the above definition and properties of the Poisson distribution function that a Poisson process has independent increments.

The assumption that $\bar{F}(t, T)$ is "measurable" is about the weakest assumption that can be made and still have the integral $\int \bar{F}(t, T) dM(t)$ make sense. The measurability condition is met, for instance, if $\bar{F}(\cdot, T)$ has no more than a countable number of discontinuities.

If the arrival process is compound Poisson then the general result still holds: The number of survivors is a compound Poisson r.v. I have assiduously avoided treating this more general problem for several reasons:

- (1) If the compounding distribution changes in time, the compounding distribution for the number of survivors is quite complicated.
- (2) Renewal processes are uncommon where the arrival process is Poisson with a continuous mean function but the arrivals are nontrivial bunches.
- (3) The practice of assuming an arrival process is compound Poisson when the data exhibit a variance-to-mean ratio greater than 1, as is often advocated in the literature, has little legitimate justification. If the variance-to-mean ratio is *significantly* larger than 1 (whatever that means), the arrival process is more often nonhomogeneous than is compound Poisson; that is, as a result of extraneous influences, some cells should be expected to have more observations than others. In inventory applications, treating the arrival process as a compound Poisson when it is in fact nonhomogeneous Poisson can be shown to have a significant effect on the optimal stocking plan (Lu, 1977).

To prove the dynamic or nonhomogeneous form of Palm's Theorem calls for six lemmas. The first three are well-known measure theo-

retic results included here for completeness. Lemma 4 (Kallenberg, 1976, p. 52) establishes the measure theoretic underpinnings of the theorem. Lemmas 5 and 6 extend the ideas of the simple proof of Palm's Theorem given in Section II to the dynamic case.

Lemma 1 (Halmos, 1950)

$M(\cdot)$ determines a measure on the ring R generated by left open, right closed bounded intervals of the form $(r, s] = \{t: r < t \leq s\}$. $M[\cdot]$ is defined by the requirement that

$$M[(r, s]] = M(s) - M(r).$$

Lemma 2 (Halmos, 1950)

$M[\cdot]$ has a unique extension (also called $M[\cdot]$) to the σ -ring S of Lebesgue measurable sets generated by R .

Lemma 3

A realization of the random process $\{X(t)\}$ has a unique extension to a random measure $X[\cdot]$ on S with the property that $X[E]$, $E \in S$, is the number of arrivals in E :

$$X[E] = \sum_k 1_E[A(k)]$$

Lemma 4 (Kallenberg, 1976)

For $E \in S$, $X[E]$ is a Poisson r.v. with mean $M[E]$.

The next lemma is a restatement of the fact (Feller, 1957, p. 160) that if a Poisson arrival process is randomly censored--arrivals are either detected or undetected as the result of an independent random selection--the detected arrivals are also a Poisson process. This almost trivial result is key to Palm's Theorem.

Lemma 5

If $\{X(T)\}$ is a Poisson arrival process with mean function $M(\cdot)$ and, independent of the arrival process, each arrival before time T

survives until T with probability p ; then if $E \in S$, the number of arrivals in E before T that survive until T is a Poisson r.v. with mean

$$p \cdot M[E \cap \{t: t \leq T\}].$$

The next lemma establishes the dynamic form of Palm's Theorem for a class of survival time distributions that have little practical utility but happen to be dense in the class of all measurable survival time distributions. It is a straightforward application of lemma 6.

Lemma 6

If $\bar{F}(t, T)$ is, for fixed T , a simple function of t (constant on finitely many bounded sets $E_i \in S$ and zero elsewhere) and the survival process is independent of the arrival process; then the size of the population at time T is a Poisson r.v. with mean

$$\int \bar{F}(t, T) dM[t].$$

Proof of lemma 6. If $\bar{F}(\cdot, T)$ is constant on E_1 , say $\bar{F}(t, T) = p$, $p > 0$, for all $t \in E_1$, then by lemma 5, the number of arrivals at t , $t \in E$, that survive until T is Poisson with mean

$$pM[E_1 \cap (-\infty, T)] = \int_{E_1} \bar{F}(t, T) dM[t].$$

Because $\bar{F}$ is zero except on finitely many disjoint sets E_i , the result follows by summing over i .

Proof of the dynamic form of Palm's Theorem. Given an $[X(k), F(t, T)]$ arrival-survival process, the process can be "squeezed" between two ancillary arrival-survival processes, which are constructed so that they satisfy the requirements of lemma 6.

To simplify the notation, consider T fixed and define a "counting function" B .

$$B(X(k)) = \begin{cases} 1 & \text{if } X(k) < T, \text{ and } X(k) \text{ survives until } T \\ 0 & \text{otherwise.} \end{cases}$$

Then, independent of the arrival process, for any large n , define two random processes $D^n[X(k)]$ and $C^n[X(k)]$ as follows:

For each $X(k)$, let i , $0 \leq i \leq n$, be defined so that

$$\frac{i}{n} < \bar{F}[X(k), T] < \frac{i+1}{n}.$$

If $B[X(k)] = 1$, let

$$D^n[X(k)] = \begin{cases} 1 & \text{with probability } \frac{i/n}{\bar{F}[X(k), T]} \\ 0 & \text{with probability } 1 - \frac{i/n}{\bar{F}[X(k), T]} \end{cases}$$

and if $B[X(k)] = 0$, let $D^n[X(k)] = 0$.

Then $D^n[X(k)] \leq B[X(k)]$ and $\sum_k D^n[X(k)]$ can be considered the number of survivors at T of an arrival-survival process where the survival distribution $\bar{G}^n(t, T)$ is equal to i/n on the sets $E_i = \{t: i/n < \bar{F}(t, T) < (i+1)/n\}$, $i = 1, \dots, n$ and zero elsewhere. Thus, by lemma 6, $\sum_k D^n[X(k)]$ is a Poisson r.v. with mean

$$\int \bar{G}^n(t, T) dM[t].$$

Similarly, we may define a $C^n[X(k)]$ process that bounds $B[X(k)]$ on the other side. For each $X(k)$, let i be as before.

If $B[X(k)] = 0$, and $i = n$, let $C^n[X(k)] = 1$. For $i < n$ let

$$C^n[X(k)] = \begin{cases} 1 & \text{with probability } \frac{(i+1)/n - \bar{F}[X(k), T]}{1 - \bar{F}[X(k), T]} \\ 0 & \text{with probability } 1 - \frac{(i+1)/n - \bar{F}[X(k), T]}{1 - \bar{F}[X(k), T]} \end{cases}$$

If $B[X(k)] = 1$, then let $C^n[X(k)] = 1$. Then $B[X(k)] \leq C^n[X(k)]$ and the survival distribution $\bar{H}^n(t, T)$ for the C^n process is equal to $\min[(i+1)/n, 1]$ on E_i ($i = 0, \dots, n$) and zero elsewhere. Thus, by lemma 6, $\Sigma C^n[X(k)]$ is a Poisson r.v. with mean

$$\int \bar{H}^n(t, T) dM[t].$$

To finish the proof note that for large n and any $m = 0, 1, \dots$,

$$\Pr(\Sigma D^n[X(k)] < m) > \Pr(\Sigma B[X(k)] < m) > \Pr(\Sigma C^n[X(k)] < m)$$

and the terms on the left and right are cumulative Poisson probabilities with means

$$\int \bar{G}^n(t, T) dM[t] \quad \text{and} \quad \int \bar{H}^n(t, T) dM[t],$$

respectively. It follows from the construction of D^n and C^n that

$$\bar{G}^n(t, T) \leq \bar{F}(t, T) \leq \bar{H}^n(t, T) \quad \text{and} \quad 0 \leq H^n(t, T) - G^n(t, T) \leq 1/n.$$

Thus, the difference between the means of the Poisson random variables mentioned above is no more than $M(T)/n$. It follows that $\Sigma B[X(k)]$ is a Poisson r.v. with mean $\int \bar{F}(t, T) dM[t]$, as was to be shown.

BIBLIOGRAPHY

- Carter, G. M., and J. E. Rolph, *New York City Fire Alarm Production Models II. Alarm Rates*, R-1215-NYC, The Rand Corporation, January 1975.
- Chung, K. L., *A Course in Probability Theory*, Academic Press, New York 1974.
- Crawford, G. B., *WRSK/BLSS Analysis, The Plans-Oriented Requirements Model*, Hq Pacific Air Forces/OA, March 1977.
- Doob, J. L., *Stochastic Processes*, Wiley, New York, 1953.
- Feller, W., *An Introduction to Probability Theory and Its Applications*, Vol. I, Wiley, New York, 2nd edition, 1957.
- Halmos, P. R., *Measure Theory*, D. Van Nostrand, Princeton University Press, 1950.
- Hillestad, R. J., and M. J. Carrillo, "Models and Techniques for Recoverable Item Stockage When Demands and the Repair Process Are Nonstationary--Part I: Performance Measurement," The Rand Corporation, N-1482-AF, May 1980.
- Kallenberg, O., *Random Measures*, Academic Press, New York, 1976.
- Lu, J. Y., Unpublished research, The Rand Corporation, November 1977.
- Lu, J. Y., R. Brooks, and C. Gillen, *Alternative Measures of Supply Performance: Fills, Backorders, Operationalrate, NORS*, The Rand Corporation, RM-6094-FM, August 1969.
- Muckstadt, J. A., *Comparative Adequacy of Steady State Versus Dynamic Models for Calculating Stockage Requirements*, The Rand Corporation, R-2636-AF, November 1980.
- Munroe, M. E., *Introduction to Measure Theory and Integration*, Addison-Wesley, Cambridge, 1953.
- Palm, C., "Intensitatsschwankungen im Fernsprechverkehr," *Ericsson Technics*, Stockholm, No. 44, 1943, pp. 1-189, esp. p. 47.
- Rice, D. B., *Defense Resource Management Study--Case Studies of Logistic Support Alternatives*, Office of the Secretary of Defense, February 1979, pp. 85-108.
- Takacs, L., *Introduction to the Theory of Queues*, Oxford University Press, New York, 1962.