Reliability Engineering and System Safety 000 (2017) 1–16

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

A hierarchical coloured Petri net model of fleet maintenance with cannibalisation

Jingyu Sheng [a], Darren Prescott [b,*]

[a] 93582 Troops, Chinese Air Force, Shouzhou, Shanxi 038300, China
[b] Resilience Engineering Research Group, University of Nottingham, Nottingham, NG7 2RD, UK

ARTICLE INFO

Keywords:
Fleet
Maintenance
Unrestricted cannibalisation
Cannibalisation bird
Hierarchical coloured Petri net

ABSTRACT

Cannibalisation refers to a maintenance action where an unserviceable part in an inoperative platform is replaced by a serviceable part of the same type from another platform. It helps a fleet meet operational requirements when spares are in short supply but leads to more maintenance tasks to be carried out. In practice, cannibalisation may be performed in an unrestricted manner, or through the use of cannibalisation birds. A cannibalisation bird is a platform which is selected as the primary source of cannibalisation, while any inoperative platform can be a cannibalisation source under the unrestricted policy. In order to aid fleet managers in making cannibalisation-related decisions, this paper presents a hierarchical coloured Petri net (HCPN) model of a fleet operation and maintenance process which considers mission-oriented operation, multiple level maintenance, multiple cannibalisation policies (no cannibalisation, unrestricted cannibalisation and cannibalisation bird), maintenance scheduling and spare inventory management. The model is applied to an example fleet to compare the effects of different cannibalisation policies on fleet performance using a number of performance measures related to reliability and maintenance and to optimise the number of cannibalisation birds used and the length of time that a platform is taken as a cannibalisation bird for the fleet.

1. Introduction

Cannibalisation is a maintenance activity that involves removing serviceable parts from one platform to replace failed parts in other platforms when the required spares are unavailable. It can restore non-mission-capable (NMC) platforms to the mission-capable (MC) state in a short time when the fleet spare inventory cannot meet the demand for spares. Cannibalisation is widely used in military aviation, with approximately 850,000 documented cannibalisations performed in the U.S. Air Force and Navy from 1996 to 2001 [1]. Fleet cannibalisation can be either unrestricted or follow a cannibalisation bird policy. Under unrestricted cannibalisation, all NMC platforms in the fleet can be cannibalised from and any of their working components used in the place of spares when spares are unavailable. By contrast, when the cannibalisation bird policy is implemented, only designated platforms can be used as sources of cannibalised spare parts. When components of other aircraft fail and there are no spares in stock, maintenance crews remove the required components from the cannibalisation bird and install them in the destination platform to make it mission capable. Thus, any part shortages are consolidated into the cannibalisation bird. If a cannibalised platform remains in the NMC state for too long, it is referred to as a hangar queen [2]. When an aircraft becomes a hangar queen in the U.S. Air Force, it is required to be reported up the chain of command and poorly affects the wing's rating [2].

Current research on cannibalisation mainly focuses on evaluating the impact of unrestricted cannibalisation on fleet performance. Fisher [3] provides an overall review of the issues present in cannibalisation decisions and the related models. Cassady et al. [4] develop a discrete-event simulation model to quantify the effect of unrestricted cannibalisation on fleet readiness (the average number of MC platforms) and manpower working hours. Each platform in the fleet is assumed to consist of two components connected in series. Ormon and Cassady [5] extend the work of Cassady et al. [4] by taking the discrete-event simulation model as a decision-support tool to optimise cannibalisation policies and spare part inventories. The platforms studied are the same as those studied by Cassady et al. [4]. Fleet readiness and maintenance cost are used to measure fleet performance and two optimisation models are built. The first aims to maximise fleet readiness subject to maintenance budget constraints and the second minimises maintenance cost subject to fleet readiness limitations. Salman et al. [1] extend the work of Ormon and Cassady [5] by studying more complex fleets in which component failure times follow Weibull distributions. Sheng and Prescott [6] build a coloured Petri net (CPN) model for the unrestricted fleet cannibalisation process considering multiple fleet maintenance factors including spares,

* Corresponding author.
E-mail address: darren.prescott@nottingham.ac.uk (D. Prescott).

http://dx.doi.org/10.1016/j.ress.2017.05.043
Available online xxx
0951-8320/

repair, platform failure logic and queuing of platforms prior to their restoration. However, none of these cannibalisation models involves the cannibalisation bird policy.

The cannibalisation bird policy is widely used in U.S. military aircraft fleets [4]. Generally, there is only one cannibalisation bird in an aircraft fleet and the cannibalisation bird will not be used to perform missions even when in the fully operative state. After having been kept as a bird for a pre-determined period of time (usually less than 30 days), an aircraft will be restored to a MC, flyable status in a process named cannibalisation recovery and another aircraft will be selected as the new cannibalisation bird. In the literature, there are very few studies of cannibalisation bird programmes. Based on cannibalisation practices at the U.S. Hill Air Force Base, Cassady et al. [4] establish an aircraft-level simulation model to study a cannibalisation bird program. When spare requests cannot be satisfied, a cannibalisation bird is checked for the required components. Without determining whether or not those components are available, they use a mathematical function to model the usefulness of the cannibalisation bird, a feature that decreases while it remains on the ground. Powell [13] introduces the cannibalisation bird practice in an F-16 aircraft fleet deployed at the U.S. Hill Air Force Base and summarises the effect of the cannibalisation bird policy on fleet performance according to the real data collected. However, no mathematical models of the implementation of cannibalisation bird policies are provided [13].

The fleet maintenance process is a complex system which involves various activities, maintenance policies, maintenance organisations and the management of maintenance resources, many of which could affect the desirability of cannibalisation. Fleet maintenance is often organised into three levels: organisation- (O-), intermediate- (I-) and depot-level (D-), or sometimes only O-level and D-level. NMC platforms are maintained at the O-level maintenance organisation where their failed components are replaced by spares or through cannibalisation. Removed, failed components are sent to the I-level maintenance organisation for repair or to the depot in case of two-level maintenance. Failed components that cannot be repaired at the I-level will be sent to the depot. After repair at I-level or D-level, previously-failed components are returned to the fleet as new spares. The use of cannibalisation is affected by the capability to repair failed components at the I-level. Since the depot can be a considerable distance from the fleet's operational base, cannibalisation may be more desirable if there is no I-level maintenance organisation or the I-level maintenance organisation is incapable of repairing many failed components. A fleet cannibalisation model should aim to involve those maintenance factors.

Fleet operation and maintenance are tightly related. The fleet maintenance tasks come about because of the fleet operations and the efficiency of fleet maintenance organisations can significantly affect the operational performance of the fleet. Cannibalisation may be necessary due to weaknesses in fleet logistic and maintenance systems and fleet managers may choose to perform cannibalisation to avoid the risk of missing mission sorties due to stock shortages. However, operational factors such as mission scheduling or fleet assignment are rarely considered during cannibalisation. The effectiveness of cannibalisation is also affected by decision variables such as the selection discipline of cannibalisation sources, the number of birds and the length of time that a platform remains a bird. Designating more aircraft to be cannibalisation birds or using shorter bird lives means there will be more cannibalisation resources but fewer MC (mission capable) platforms and a higher burden of cannibalisation recovery. Fewer birds or longer bird lives lead to less cannibalisation recovery effort but fewer available cannibalisation resources and a higher risk of failures of previously functioning components in cannibalisation birds. Therefore, the fleet maintenance model should also involve fleet mission-oriented operations and be capable of investigating cannibalisation decision variables.

Very few indicators of overall fleet performance are used to measure the influence of cannibalisation, with the most commonly-used being fleet readiness (average number of MC platforms), maintenance cost and consumed manpower hours. Reliability and operational parameters that might better indicate fleet performance are rarely studied. Therefore, benefits may result from modelling both fleet operation and maintenance and using reliability- and operation-related measures to ascertain the impact of cannibalisation on fleet performance.

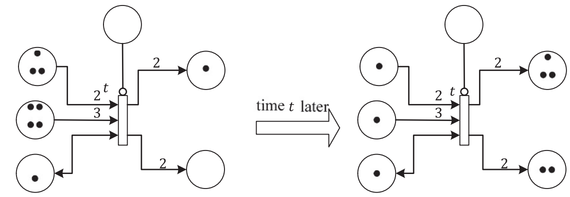

Fig. 1. Transition enabling and firing.

The first objective of this paper is to model fleet mission-oriented operation and multi-level maintenance processes under three different cannibalisation policies: no cannibalisation, unrestricted cannibalisation and the cannibalisation bird policy. A novel hierarchical CPN (HCPN) model is presented, which allows the analysis of any of these policies for fleets with different size (number of platforms), operational requirements, maintenance organisations and maintenance policies without modification of the HCPN model. The second objective of this paper is to demonstrate how the HCPN model can be applied to example fleets, and used to evaluate and compare the effects of different cannibalisation policies on fleet performance and to investigate the effect on cannibalisation of factors such as the number of cannibalisation birds used at one time and the length a platform is kept as a bird. The model allows the study of reliability- and operation-related measures relating to fleet performance, namely the mission abort rate (MAR), mission capable rate (MCR) and cannibalisation rate (CAR), in addition to the maintenance cost. The application of the model demonstrates its potential application as a maintenance decision support tool for fleets.

2. Petri nets

Petri nets (PN), first introduced by Carl Adam Petri in 1962 [8], are powerful graphical and mathematical tools for modelling complex, dynamic systems. Since they can be constructed at various levels of abstraction and designed hierarchically, PN have been successfully applied in many fields, including reliability assessment [9,11,16], maintenance modelling [6,10,12] and railway bridge asset management [17]. A PN is a directed graph with two types of nodes: *places*, shown as circles; and *transitions*, drawn as bars. The nodes are connected by *arcs*, which link a place to a transition or vice-versa. A place can have a discrete number of *tokens* and the distribution of tokens within places defines the PN marking, which represents the state of the modelled system at any point in time. The system state changes as tokens are moved between, created in or removed from places as transitions fire. In order to fire, a transition must first be enabled, which occurs when each input place contains a number of tokens that is no less than the weight of the arc linking the place to the transition. An immediate transition, represented by a solid bar, fires as soon as it is enabled. A timed transition, represented by a hollow bar, fires after a certain delay has elapsed. This firing delay can either be fixed or randomly sampled from a known probability distribution as the transition is enabled. When a transition fires, a number of tokens equivalent to the associated arc weight is removed from each input place and added to each output place.

Figs. 1 and 2 demonstrate the enabling and firing of a timed transition. The transition consists of three input places and three output places where the place linked with a double-headed arc is both an input and output. The place linked by an inhibitor arc, drawn with a circle at its head instead of an arrow, prohibits the firing of a transition if the number of tokens within it is no less than the weight of the inhibitor arc. In Fig. 1, the transition is enabled and fires after a time delay t, removing