

A Hierarchical Coloured Petri Net Model of Fleet Maintenance with Cannibalisation

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.

Keywords

Fleet, Maintenance, Unrestricted Cannibalisation, Cannibalisation Bird, Hierarchical Coloured Petri Net

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 850000 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. Cassidy 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 Cassidy [5] extend the work of Cassidy 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 Cassidy 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 Cassidy [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, 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 [14] 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 [14].

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.

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 [9,11,17], maintenance modelling [6,10,12] and railway bridge asset management [18]. 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.

Figures 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 Figure 1, the transition is enabled and fires after a time delay t , removing tokens from the input places and adding tokens to the output places according to the associated arc weight. By contrast, in Figure 2, the transition is disabled due to the token in the place linked by the inhibitor arc.




Figure 1. Transition enabling and firing.




Figure 2. Inhibitor arc prevents transition firing.

Created by Jensen and Kristensen [7], coloured Petri nets (CPN) extend PN by allowing the behaviour of complex systems, especially concurrent systems, to be concisely specified and modelled. In contrast to a PN, tokens in a CPN are distinguished by *token colours*. The token colour is a label that can represent different data types, such as Boolean, integer or character strings. Each place is assigned a *colour set*, which holds the possible token colours of a specified type. Each arc is assigned an *arc expression*, which determines the amount of tokens of specific colours that are removed from or added to the linked place when a transition fires. Each transition is assigned a *guard*, which is a *Boolean* expression that represents a constraint on its enabling policy.

A transition is enabled only if the marking of each input place satisfies the associated input arc expression and its guard evaluates to be true. Transitions can fire tokens of different colours in different ways and the colour of a token can change after it passes through a transition as decided by a specific modelling condition. Figure 3 shows an instance of transition enabling and firing for a CPN. All places in Figure 3 have the same colour set $CS = \{\text{green}, \text{red}, \text{blue}\}$. cs_1 , cs_2 and cs_3 are colour variables that belong to CS . The arc expression shown as cs_1 , cs_2 or cs_3 indicates one token of a specified colour from the colour set, while $cs_1 + cs_2$ represents two coloured tokens. The transition

guard, $cs_1 = cs_2 \neq cs_3$, indicates that tokens in the two input places can be switched by the transition only if they have the same colour and there is no token with the same colour in the place linked by the inhibitor arc. The delay times associated with the firing of green, red and blue tokens by the transition are respectively t_1 , t_2 and t_3 and $t_1 < t_2 < t_3$. After a delay time t_1 , the transition moves the green tokens from its input places to the output place. After a further delay of $t_2 - t_1$ the transition moves the red tokens. The firing of the transition for the blue tokens is disabled because of the blue token in the input place linked by the inhibitor arc.




Figure 3. Transition fires tokens with different switching time

Volovoi [9] introduces *aging tokens* as an extension to CPNs, which assign a label that acts as a ‘counter’ to each token in a place. A token counter can change continuously while it contributes to the enabling of the transition. The changing of the token counter is interrupted when the transition is disabled. Aging tokens are potentially useful for modelling aircraft operation processes. By using a timed transition to represent the failure process of a component, the token counter can be used to represent the remaining working life (time to failure) of the component if it can only fail during operation. The token counter decreases when the component is operational, a process which stops when the component is not operational. The timed transition fires when the counter reaches 0, indicating that its remaining life has declined to 0 and that it has failed.

A hierarchical CPN (HCPN) is a CPN that is organised as a set of modules, which are linked by a number of interfaces [7]. A module can be a non-hierarchical CPN or a HCPN with its own sub-modules. Module interfaces consist of *port places*, *substitution transitions* and *socket places*. Different modules are connected by exchanging tokens through ports. A module can import tokens via *input ports* and export tokens via *output ports*. A port can be both an input and an output port so that the module can use it to both import and export tokens.

A module and its sub-modules are related by substitution transitions and socket places. A substitution transition, drawn as a rectangular box in this paper, represents an abstract view of a module to its parent module, with the sub-module giving a detailed view of the substitution transition’s function. The places linking a substitution transition to its parent module are called sockets and are the substitution transition’s interfaces.

To achieve a complete hierarchy, it is necessary to identify the *port-socket relation* which shows how the ports of a sub-module are related to the sockets of its substitution transition. When a port and a socket are related, they conceptually become a single *compound place* and hence they always share the same marking. Figure 4 shows an example HCPN with one substitution transition, for which places P_1 and P_2 are respectively an input socket and an output socket. The substitution transition represents a sub-module that is a non-hierarchical CPN with two transitions. The input port and output port of the sub-module CPN are respectively P_1 and P_2 , which are related to the socket places with the same names in the HCPN.

There are clear benefits to employing hierarchies in a CPN [7]. First, for complex systems such as aircraft, a single non-hierarchical net would be very large and hard to construct, interpret and modify. Using modules allows more convenient construction, interpretation and modification, since modules can be considered in isolation, which should be simpler than having to account for the whole net. Secondly, a CPN module can be viewed as a black box that allows modellers to work at different levels of abstraction without being concerned about the details within modules.

The HCPN method is selected as the tool to model fleet cannibalisation because of its capability of providing a graphical, precise and concise representation of the operation and maintenance processes of large, complex fleets. With token colours, any number of platforms in the fleet can be represented by an equivalent number of tokens with different colours in a single place, which means that the size of the HCPN model will not increase with the size of

the modelled fleet. With a modular structure, a HCPN model can be easily modified and extended to consider other factors of fleet operation and maintenance when being applied to practical fleets.




Figure 4. Example hierarchical CPN

3. HCPN Model of Fleet Cannibalisation

In this section a HCPN model of fleet cannibalisation is presented. Cannibalisation is generally applied to fleets of identical platforms such as military vehicles, buses or trains, which are used to perform missions so long as they are mission capable, meaning that component failures on the platforms will not prevent them from performing the tasks required to complete the mission. The model presented here is described in relation to an aircraft fleet but could equally be applied to other, similar fleets with little or no amendment.

Due to the stochastic nature of fleet maintenance processes, calculating fleet performance using an analytical solution of the HCPN model is impossible. However, the HCPN model provides an ideal framework for a Monte Carlo simulation analysis. Such an analysis is demonstrated for an example fleet in Section 4.

3.1 Fleet Description

A fleet is assumed to be comprised of a number of identical and independent platforms, each of which contains a number of different repairable components. Components and platforms have binary states. A platform can be either mission-capable (MC) or non-mission-capable (NMC), while a component can be either working or failed. Components (hence platforms) can only fail during operation. When a mission is called, the fleet manager will assign MC platforms to perform it. The mission aborts if the number of idle, MC platforms in the fleet cannot meet its requirement. Otherwise, the mission starts when the required number of platforms have been assigned. The assigned platforms remain operational until the mission ends. An air mission abort happens when a platform becomes NMC during operation. The platform will immediately be taken out of service and added to the ‘platform maintenance queue’ (PMQ) which holds the NMC platforms that are required to be restored to the MC state.

Fleet maintenance is organised at three levels: O-level, I-level and D-level. NMC platforms are restored at the O-level. If the PMQ contains more than one NMC platform, those platforms are maintained based on a queuing discipline such as random selection, first in first out (FIFO) or last in first out (LIFO). When a platform is selected to be restored, it is first checked whether there are enough spare components to replace its failed components. If so, maintenance technicians take the required spares from stock and install them in the platform to restore it to the MC state. Otherwise, technicians check cannibalisation birds for available working components of the required types if the cannibalisation bird policy is used, removing the required components from birds and installing them in the platform being maintained. Under the unrestricted cannibalisation policy technicians check other NMC platforms to find the required working components. If these exist, unrestricted cannibalisation is performed from the selected NMC platforms instead of cannibalisation birds. Both spare and cannibalised parts may be needed during a platform restoration. Once restored, a platform can be assigned to perform scheduled missions. If there are insufficient spare and cannibalised components to fully restore a NMC platform, the restoration will not begin and the NMC platform will be added to the ‘platform waiting maintenance queue’ (PWMQ) which holds all NMC platforms for which the existing resources do not allow complete restoration.

All removed, failed components at O-level are sent for repair to the I-level repair shop, where there is a certain probability that repair is possible. If repair is possible at I-level, I-level maintenance technicians will return it to the working state. Otherwise, it is sent for repair at the depot. Component repair is assumed to be perfect. Manpower limitations are not considered in this paper, meaning there are always sufficient maintenance personnel to perform the required maintenance actions. The repaired components, forming the new spare part inventory, are returned to the O-level maintenance organisation. If the PWMQ contains NMC platforms when new spares arrive, maintenance technicians will check whether any platform in the PWMQ can be entirely restored using new and existing spares and cannibalisation resources. If so, crews will perform the required spare installations and cannibalisations to restore those platforms to the MC state. Otherwise, NMC platforms remain in the PWMQ until the required resources become available.

3.2 Model Hierarchy

The HCPN cannibalisation model is made up of five modules:

- *Fleet module* - models fleet operation and the management of cannibalisation birds;
- *Platform module* - models platform failures;
- *Component module* - models component failures;
- *Maintenance module* - models fleet three-level maintenance processes under different cannibalisation policies;
- *Master module* - top-level of the HCPN model which connects the other four modules to provide an abstract view of fleet operation and maintenance processes.

The colour sets and colour variables applied in these modules are defined in Tables 1 and 2 respectively. For an arbitrary variable c of colour set C , $c(p)$, $c(ct)$ and $c(a)$ represent a platform, component type and remaining working life, respectively. $c(a)$ is initialised by sampling from a known distribution and decreases continuously when the related component is in operation. A spare component of type ct is represented by $(0, ct, a)$.

Table 1. Colour sets

Colour set	Colours	Meaning
P	$\{1, 2, \dots, N_P\}$	Platform
CT	$\{1, 2, \dots, N_C\}$	Component type
C	$\{(p, ct, a)\}$	Components: $a \geq 0$ is the expected working life of a component; $p=0$ means a spare component.
CRD	$\{(c, c_1)\}$ $\forall c, c_1 \in C, c(ct)=c_1(ct), c(p) \neq 0$ and $c(p) \neq c_1(p)$	Component replacement decision: the failed component c is decided to be replaced by component c_1 . $c_1(p)=0$ means c will be replaced by a new spare component. $c_1(p) \neq 0$ means that the working component c_1 is cannibalised from platform $c_1(p)$ to replace the failed component c in platform $c(p)$.
M	$\{1, 2, \dots\}$	Missions
PM	$\{(p, m)\}$ $\forall p \in P, \forall m \in M$	Platform p is assigned to perform mission m
CAP	$\{'CAB', 'URCA', 'NCA'\}$	Cannibalisation policies: 'CAB' means cannibalisation birds; 'URCA' means unrestricted cannibalisation; 'NCA' means no cannibalisation.

Table 2. Colour variables

Variable	Colour	Meaning
p, p_1	P	A platform
c, c_1, \dots, c_n	C	A component
crd	CRD	A component replacement decision
cap	CAP	A cannibalisation policy
m	M	A mission
pm	PM	A platform assignment decision

3.3 Fleet HCPN

The fleet module, shown in Figure 5, models the mission-oriented fleet operation and the management of cannibalisation birds. It contains four sub-modules:

- Fleet mission scheduling (FMS): models the scheduling of fleet missions;
- Fleet mission assignment (FMA): models the assignment of a required number of idle, MC platforms to the scheduled mission;

- Fleet mission operation (FMO): models the operation of assigned platforms under different cannibalisation policies;
- Fleet cannibalisation bird (FCAB): models the selection of cannibalisation birds and cannibalisation recovery.

Each sub-module is represented by a substitution transition with the same name. All places shown in the fleet module are sockets for these substitution transitions. The meanings and port types of places in the fleet module are listed in Table 3.

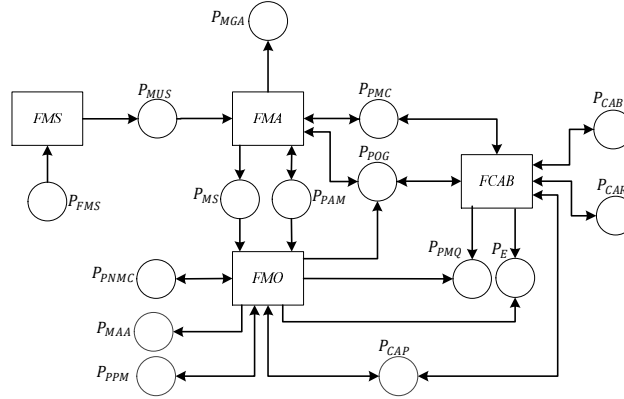


Figure 5. Fleet HCPN module

Table 3. Meanings of places of fleet module

Place	Meaning	Colour Set	Port type
P_{FMS}	Fleet missions	M	-
P_{MUS}	Mission under schedule	M	-
P_{MS}	Mission starts	M	-
P_{PMC}	Platforms that are mission-capable	P	Input/Output
P_{PNMC}	Platforms that are non-mission-capable	P	Input/Output
P_{PAM}	Platforms that are assigned to a mission	PM	-
P_{PPM}	Platforms that are performing missions	PM	Output
P_{POG}	Platforms on ground	P	-
P_{CAP}	Cannibalisation policy	CAP	Input/Output
P_{PMQ}	Platform maintenance queue	P	Output
P_E	Event	-	Output
P_{CAB}	Cannibalisation bird	P	Input/Output
P_{CAR}	Cannibalisation recovery	P	Input/Output
P_{MGA}	Mission ground abort	M	Output
P_{MAA}	Mission air abort	M	Output

The place P_{FMS} represents all the missions that the fleet is required to perform. When a mission is scheduled, the FMS sub-module outputs a token representing that mission to the P_{MUS} . This activates the FMA sub-module, which assigns MC platforms to perform the mission. If there are insufficient idle, MC platforms (represented by tokens in P_{POG} and P_{PMC}) for the mission, it will be aborted and a token output to P_{MGA} which holds all the cancelled missions. Otherwise, the FMA sub-module outputs the assigned platforms to P_{PAM} . The FMO sub-module governs the operations of the assigned platforms. When a mission starts (represented by a token in P_{MS}), tokens representing the assigned platforms will be placed in P_{PPM} , which is output to the component module to enable the failure processes of components on those operating platforms. If an operating platform becomes NMC (represented by a token in P_{PNMC}), a mission air abort event happens with a token being output to P_{MAA} . Meanwhile, the FMO sub-module removes the token representing the NMC platform from P_{PPM} and outputs it to both P_{POG} and P_{PMQ} , indicating that the platform is taken out of service and added to the PMQ. As this happens, a token is placed in P_E if the fleet employs the unrestricted cannibalisation policy, since new cannibalisation resources will be available. P_{CAP} represents the employed cannibalisation policy.

If the cannibalisation bird policy is employed, the FCAB sub-module selects a certain number (one or more) of idle, MC platforms as birds. Tokens representing the selected birds are added to P_{CAB} . Meanwhile a token is output to P_E since new cannibalisation resources are available. Those platforms will not be assigned to perform missions. After being kept as a bird for a certain period of time, a platform is restored to the MC state and another bird selected. The token representing the platform is removed from P_{CAB} and copies of it placed in both P_{CAR} and P_{PMQ} indicating that a cannibalisation recovery has begun.

The FMS sub-module is shown in Figure 6. Places P_{FMS} and P_{MUS} are respectively an input and an output port and are equivalent to the socket places with the same names in the fleet module. The transition T_{FMS} , governs the setting of the scheduled mission start times that are represented by the firing delays of tokens in P_{FMS} .




Figure 6. FMS sub-module.

The FMA sub-module is shown in Figure 7. Places P_{MS} and P_{MGA} are output ports; P_{MUS} is an input port; all other places are input/output ports. Each port place represents the socket place with the same name in the fleet module. Transition T_{FMA} assigns MC platforms when a mission is scheduled. N_m represents the number of MC platforms required by a mission m . After N_m platforms have been assigned to mission m by T_{FMA} , T_{FMA} is inhibited and T_{MS} is enabled. T_{MS} moves the token with colour m from P_{MUS} to P_{MS} immediately, signifying that mission m starts. If there are more than N_m idle, MC platforms when mission m is called, a selection discipline may be applied by the fleet manager in order to select N_m platforms. A mission is often required to start within a certain time of being scheduled. It will be cancelled if the number of assigned platforms cannot meet the mission's requirement during the assignment time window for the mission. When this happens, T_{MGA} removes the token from P_{MUS} and places it in P_{MGA} , stating that a mission ground abort event happens. Meanwhile, the pre-assigned platforms for the aborted mission are returned to the idle state by T_{CMAP} so that they can be assigned to other missions.

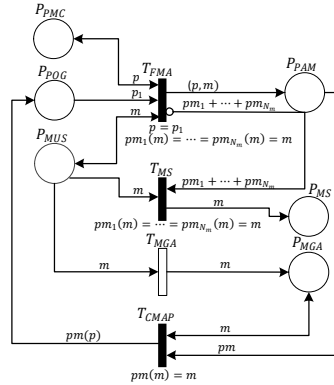


Figure 7. FMA sub-module

The FMO sub-module is shown in Figure 8. Places P_{PAM} and P_{MS} are input ports; P_{POG} , P_{PMQ} , P_{MAA} and P_E are output ports; and P_{CAP} , P_{PNMC} and P_{PPM} are input/output ports. When a scheduled mission m begins (represented by a token with colour m in P_{MS}), the assigned platforms will start to perform it, a situation modelled by T_{PSPM} moving tokens representing those platforms from P_{PAM} to P_{PPM} . This enables T_{PPM} which assigns a delay time (either fixed or randomly sampled from a distribution) to each token in P_{PPM} according to the required operating time of an assigned platform during the mission. If an air abort event happens (a MC platform p becomes NMC during the mission), either T_{MAAB} or T_{MAAU} will take the platform p out of service (P_{POG}), add it to the PMQ (P_{PMQ}) and report the air abort event (P_{MAA}), based on the cannibalisation policy applied (P_{CAP}). T_{MAAB} is responsible for handling the air abort events if the cannibalisation bird or no cannibalisation policy is employed and T_{MAAU} if an unrestricted cannibalisation policy is employed. Since new NMC platforms mean new cannibalisation sources under unrestricted cannibalisation, a token is created in P_E after T_{MAAU} fires. After an assigned platform successfully completes the mission, T_{PPM} fires, switching the platform from P_{PPM} to P_{POG} . A mission ends when all the assigned platforms land, with the related token moved from P_{MS} to P_{ME} by T_{ME} .

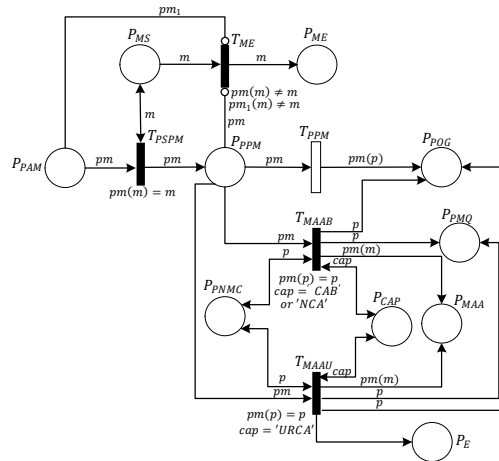


Figure 8. FMO sub-module

The FCAB sub-module is shown in Figure 9. Places P_{POG} , P_{PMC} , P_{CAP} , P_{CAB} and P_{CAR} are input/output ports and P_E and P_{PMQ} are output ports. Under the cannibalisation bird policy, the fleet keeps one or more MC platform grounded as primary cannibalisation sources. In order to avoid platforms staying grounded for too long a period of time, it is assumed that only MC platforms can be selected as cannibalisation birds. The selection of the cannibalisation birds is governed by the transition T_{SCAB} , which places the required number (N_{CAB}) of selected cannibalisation birds in P_{CAB} . After having been used as a cannibalisation bird for a certain period of time, a platform is added by T_{CAR} to the PMQ (P_{PMQ}) to be maintained, after which T_{SCAB} selects a new bird. Meanwhile, any NMC platforms in the PWMQ are checked for the possibility of restoration since new cannibalisation resources became available. After restoration, T_{CARF} returns platforms that have undergone cannibalisation recovery (tokens in P_{CAR}) to P_{POG} .

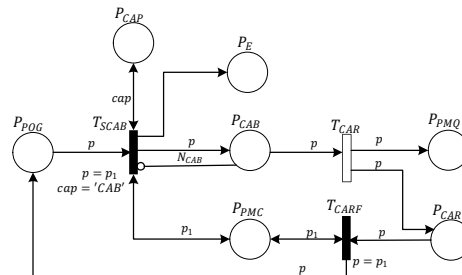


Figure 9. FCAB sub-module.

3.4 Platform Module

The platform module models the process of platform failures, whose failure logic can be represented by fault trees (FTs). A FT models a particular system failure mode, represented by its top event, whose causes are developed using logic gates, intermediate events and basic events, which usually represent component failure [11]. Figure 10 shows an example platform FT with ten basic events, C_1 - C_{10} , four intermediate events, S_1 - S_4 and five logic gates, G_1 - G_5 (G_1 and G_2 are examples of OR and AND gates respectively).

Figure 11 shows the platform module that is equivalent to the example FT shown in Figure 10. All basic events are represented by the place P_{CF} with colour set C while each intermediate or top event is represented by a place with colour P . An AND gate is represented by an immediate transition and each OR gate is converted to a number of immediate transitions equal to the number of converted places of the gate's input events. If component failures lead to the failure of a MC platform p , a token with colour p will be placed in P_{TOP} . This enables T_{PNMC} , which will change the state of the platform to NMC through switching the token with colour p from P_{PMC} to P_{PNMC} . Since cannibalisation birds do not perform missions, the state of MC platforms is immediately changed to NMC by T_{BNMC} once they are selected as cannibalisation birds. When a NMC platform p is restored (P_{PR} contains a token with colour p), its state will be changed to MC by T_{PMC} and the intermediate events of the restored platform will be

cleared by the transitions linking P_{PR} and the places representing the existence of the fault tree intermediate and top events.

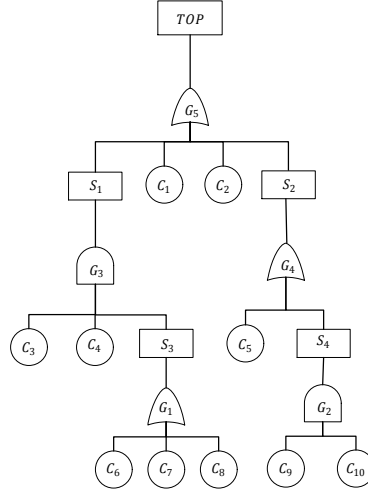


Figure 10. Example platform fault tree.

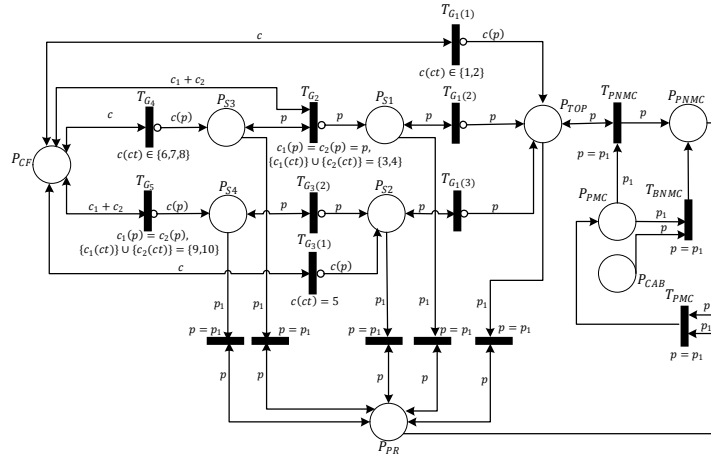


Figure 11. Example platform module.

The different failure modes of a platform can each be represented by a fault tree. A HCPN model (as shown in Figure 11) is needed for each fault tree. Figure 12 shows the platform module for platforms with N failure modes, where each substitution transition represents a module for each failure mode.




Figure 12. Platform module with multiple failure modes.

3.5 Component Module

The component module represents the component failure process. It is assumed that the component time to failure follows a known distribution and a component can fail only during operation of the platform to which it is fitted. Figure 13 shows the component module in which P_{CF} is an output port, P_{CW} is an input port and P_{PPM} is an input/output port.




Figure 13. Component module.

The failure process of a working component c on platform p is activated when the platform starts to perform a mission (P_{PPM} contains a token with colour pm such that $pm(p)=p$). T_{CF} then assigns a delay time equal to $c(a)$, the remaining working life of component c , to the associated token in P_{CW} . $c(a)$ is initialised by randomly sampling from a known probability distribution relating to the component failure time and continuously decreases when T_{CF} is enabled by component c . Since component c can only fail during operation, the component failure process can be interrupted once the mission ends due to its completion or the platform becoming NMC. When this happens, the component remaining life, $c(a)$, stops decreasing. If component c fails before platform p finishes the current mission (P_{PPM} is still marked with token pm when $c(a)$ reaches 0), T_{CF} will switch the token c from P_{CW} to P_{CF} , indicating that component c fails.

3.6 Maintenance Module

The maintenance module, shown in Figure 14, models the process of three-level fleet maintenance considering different cannibalisation policies. It contains three sub-modules: O-level (OM), I-level (IM) and D-level maintenance (DM), and takes P_{SCI} , P_{CAB} , P_{CAR} , P_{CAP} and P_{CW} as input/output ports, P_E , P_{PMQ} and P_{CF} as input ports and P_{PR} as an output port. Places P_{ICMQ} and P_{DCMQ} respectively represent the removed, failed components that have been sent to the I-level and D-level maintenance organisations. P_{CAN} counts the cannibalisation actions performed. P_{SCI} and P_{NS} represent the spare component inventories and repaired components that are returned to the fleet, respectively. Each of the substitution transitions OM , IM and DM represents a sub-module with the same name.

The OM sub-module governs the restoration of NMC platforms (represented by tokens in the place P_{PMQ}) at the O-level maintenance organisation. When a platform is being maintained, maintenance crews will replace its failed components (represented by tokens in P_{CF}) by spares or through cannibalisation. After all failed components have been replaced, the OM sub-module outputs a token representing the restored platform to the place P_{PR} . After removal, failed components are output to the place P_{ICMQ} so that repair can be attempted in the IM sub-module. If a failed component is non-repairable at the I-level maintenance organisation the IM sub-module will output a token representing it to P_{DCMQ} and it will be repaired in the DM sub-module. After repair at either I-level or D-level, repaired components are returned to the fleet through P_{NS} which forms the new spare inventory.




Figure 14. Maintenance module.

3.6.1 O-level Maintenance

The OM sub-module models platform restoration processes including: selection of NMC platforms to maintain, selection of spare or cannibalised components to restore the selected NMC platform, implementation of component replacement decisions, as well as management of spare inventories and cannibalisation resources under different cannibalisation policies. Figure 15 shows the OM sub-module, which is itself made up of five sub-modules:

- Platform restoration queuing (PRQ)
- Component replacement decision (CRD),
- O-level maintenance actions (OMA),
- Cannibalisation inventory control (CAIC),
- Spare inventory control (SIC),

which are shown in Figures 16-20. P_{PMQ} , P_{CF} , P_{NS} and P_E are input ports, P_{PR} and P_{ICMQ} are output ports, P_{CAP} , P_{CAB} , P_{CAR} and P_{CW} are input/output ports. The meanings of places in the OM sub-module are listed in Table 4.

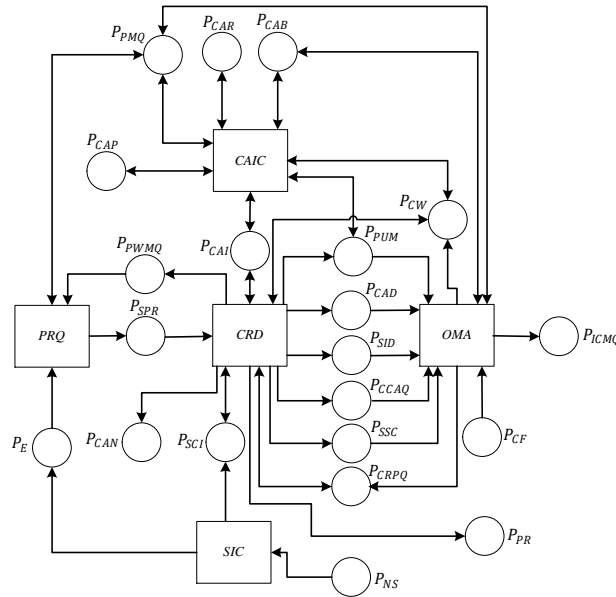


Figure 15. OM sub-module.

The restoration sequence of NMC platforms in the PMQ is decided in the PRQ sub-module based on the queuing discipline specified by the fleet manager. When a platform is selected for restoration, the PRQ sub-module outputs a token representing it to P_{SPR} . The platform is then checked for the possibility of full restoration using the existing spare and cannibalisation resources (respectively represented by tokens in P_{SCI} and P_{CAI}) in the CRD sub-module.

Table 4. Meanings of places of OM sub-module

Place	Meaning	Colour	Place	Meaning	Colour
P_{PUM}	Platforms undergoing maintenance	P	P_{CCAQ}	Component cannibalisation queue (working components that are selected to be cannibalised)	C
P_{CAI}	Component cannibalisation inventory (working components that can be cannibalised)	C	P_{SCI}	Spare component inventory (available spare components in the fleet)	C
P_{PWMQ}	Platform waiting maintenance queue	P	P_{SID}	Spare installation decision	CRD
P_{SPR}	Platform selected for restoration	P	P_{CAD}	Cannibalisation decision	CRD
P_{CRPQ}	Component replacement queue (components to be replaced during maintenance)	C	P_{SSC}	Selected spare component	C

If the selected platform cannot be fully restored, it will be output to P_{PWMQ} to wait for available resources. Otherwise, the CRD sub-module outputs the platform to P_{PUM} , which holds the platforms being maintained. Meanwhile, the

spare replacement decision and specific selected spares will be output to P_{SID} and P_{SSC} respectively. If cannibalisation is also used, the CRD sub-module will add the cannibalisation replacement decision and selected cannibalised working components to P_{CAD} and P_{CCAQ} respectively. The OMA sub-module acts according to the specified component replacement decisions. Failed components are removed and replaced by the selected spare or cannibalised components. Removed, failed components are output to P_{ICMQ} for I-level repair. Once all failed components have been replaced, the platform is restored and output to P_{PR} by the CRD sub-module.

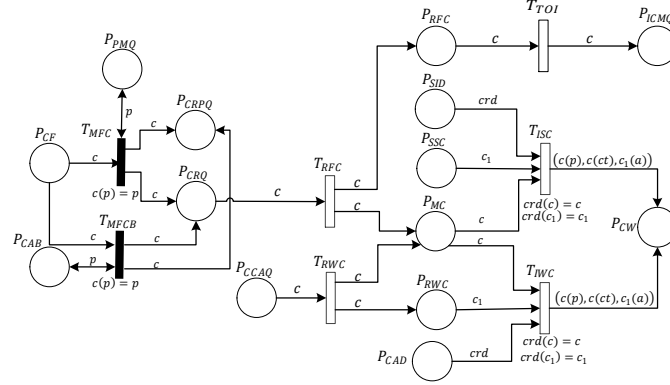


Figure 16 OMA sub-module.




Figure 17. PRQ sub-module.

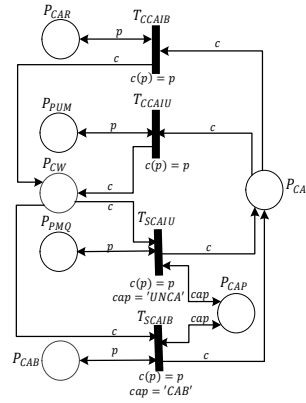


Figure 18. CAIC sub-module.

When repaired components are returned to the fleet (represented by tokens in P_{NS}), the SIC sub-module adds them to P_{SCI} , the spare component inventory. An event token will also be placed in P_E to initialise the restoration check of NMC platforms within the PWMQ in the CRD sub-module.

The available component cannibalisation resources are managed in the CAIC sub-module and differ according to the cannibalisation policy employed. They are working components from cannibalisation birds under the cannibalisation bird policy and those from NMC platforms when unrestricted cannibalisation is employed. If cannibalisation is allowed, the available cannibalisation resources are added to P_{CAI} (the component cannibalisation inventory) and are removed from the inventory when the platforms to which they are fitted are being maintained or undergoing cannibalisation recovery.

$P_{c(ct)}$ represents the probability that a component c can be repaired at the I-level repair shop. When a failed component arrives (a token with colour c is added to P_{ICMQ}), a variable r will be randomly sampled from the range 0 to 1. If $r < P_{c(ct)}$, T_{SICRQ} is enabled and fires immediately, moving the token with colour c from P_{ICMQ} to P_{IFCRQ} , signifying that the component is repairable at the I-level. Otherwise, the token is moved immediately from P_{ICMQ} to P_{FCTD} by T_{SDCRQ} and sent to the depot maintenance queue (P_{DCMQ}) by T_{TID} . The repair of failed components at I-level is governed by the transition T_{IRFC} whose associated firing delays represent component repair times. After repair,

T_{IRFC} outputs the perfectly repaired components to P_{IRC} (represented by tokens with colour $(0, ct, a)$ where a is randomly sampled from a known probability distribution relating to the component failure time). These new spares are then returned to the O-level maintenance organisation by T_{TIO} .

If the fleet maintenance is organised at two levels and there is no I-level maintenance organisation, instead of excluding the IM module from the maintenance module, p_c , the I-level repair probability, is set to zero for all components along with the firing delays of T_{TIO} and T_{TOI} (in the OMA sub-module), so that all failed components are sent to the depot after removal at the O-level maintenance organisation.

3.6.3 D-level Maintenance

The DM sub-module, shown in Figure 22, models depot repairs of failed components that are sent from the I-level maintenance organisation. In this module, P_{DCMQ} is an input port and P_{NS} is an output port. Note that all components are assumed to be repairable at the depot. Failed components that arrive at the depot will first be repaired by T_{DRC} to the ‘as good as new’ state and then sent back to the O-level through the firing of T_{TDO} .




Figure 22. DM sub-module.

3.7 Master Module

The master module, shown in Figure 23, is the top-level module which provides an abstract view of fleet operation and maintenance. This module consists of four substitution transitions representing the HCPN module with the same name and has three main functions.

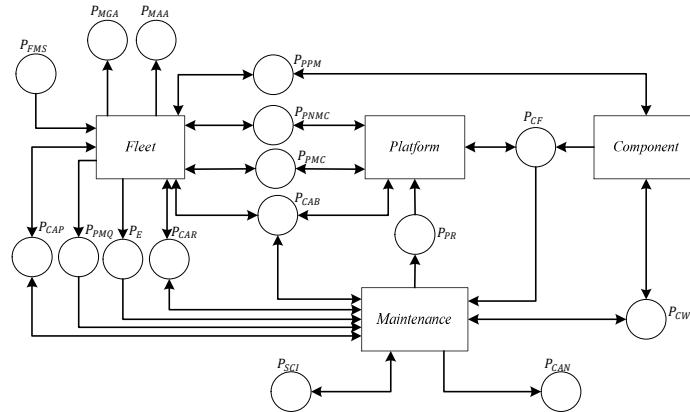


Figure 23. Master module.

The first is to show the token exchanges between fleet, platform, component and maintenance modules through their ports. The second is to initialise the whole model by setting the initial marking of places P_{FMS} , P_{CAP} , P_{PMC} , P_{CW} and P_{SCI} , representing the total fleet missions, the cannibalisation policy, the number of platforms in a fleet, the number and types of components within a platform and the initial amount of spares for each component type. The third is to collect the data which is needed for the calculation of fleet performance measures.

3.8 Model Operation

A HCPN model has been built to describe the mission-oriented fleet operation and three-level maintenance processes under multiple cannibalisation policies. As shown in Figure 24, it is made up of 17 modules that are organised into 4 levels. The master module lies on top, providing an abstract view of fleet operation and maintenance processes and acting as the user interface. Model users can study their fleet of interest by setting the marking of appropriate places in the master module with no need for further modification of the whole HCPN model. The model outputs to the master module the information needed for the calculation of fleet performance measures, which can support fleet managers in making cannibalisation and maintenance related decisions.

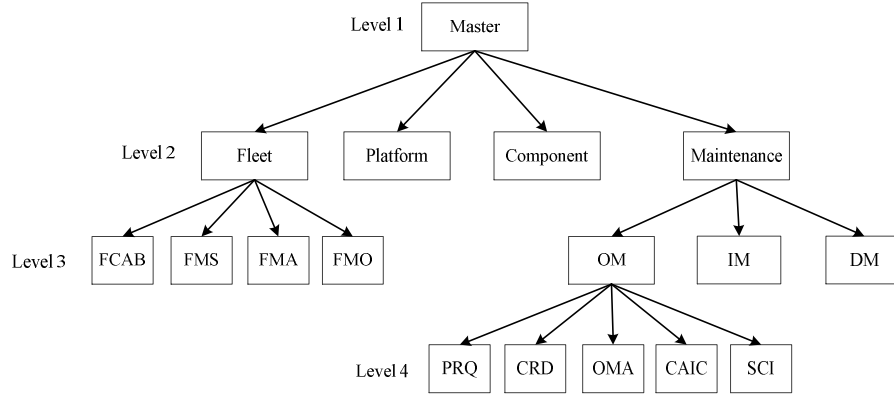


Figure 24. Module hierarchy.

3.8.1. Model Input

To use the HCPN model to investigate a fleet, it must first be initialised through the input of the fleet parameters:

- Fleet operational requirements: the number of missions required each day, the start time, assignment time window and length of each mission, and the discipline of mission assignment (how MC platforms are selected to perform scheduled missions);
- Fleet size: the number of platforms in the fleet;
- Platform structure: the number of components in a platform, platform failure modes and the associated fault trees;
- Component failure characteristics: the probability distribution describing the time to failure of each component type;
- Maintenance policies: the cannibalisation strategy (no cannibalisation, unrestricted cannibalisation or the cannibalisation bird policy), disciplines of platform restoration queuing and the selection of cannibalisation sources, the number of cannibalisation birds simultaneously used and the length of time a platform is kept as a bird under the cannibalisation bird policy;
- Maintenance resources: the number of spares for each type of component;
- Maintenance activities: the removal time, installation time, repair time and repair probability at the I-level maintenance organisation for each type of component, and the transportation time between different maintenance organisations.

The fleet parameters are input by setting the initial marking (the number and colours of tokens) of their related places and transition firing delays. For instance, cannibalisation policy is governed by the marking of P_{CAP} in the master module (Figure 23). When no cannibalisation is allowed, P_{CAP} contains a token with colour $cap='NCA'$, while the token colour is 'URCA' or 'CAB' under the unrestricted and cannibalisation bird policies respectively. The length of time that a platform remains a bird is equal to the firing delay of T_{CAP} in the FCAB sub-module (Figure 9). Tables 5 and 6 respectively list the places and transitions that must be initialised and the associated fleet input parameters represented by the place markings and transition firing delays.

In addition, a different module must be built for each platform failure mode by converting the related fault tree to a HCPN, which further corresponds to a substitution transition in the platform module (Figure 12). The platform restoration queuing and selection of cannibalisation resources are respectively governed by the transition T_{SP} in the PRQ sub-module (Figure 17) and the transition T_{PCAD} in the CRD sub-module (Figure 19), which fire tokens in their input places according to the order determined by the selected disciplines. The mission assignment discipline is governed by the transition T_{FMA} in the FMA sub-module (Figure 7). The selection of cannibalisation birds is governed by T_{SCAB} in the FCAB sub-module (Figure 9). The number of cannibalisation birds, the required number of MC platforms to perform a scheduled mission and the I-level repair rates of failed components are respectively governed by the arc expressions of the inhibitor arc linking P_{CAB} to T_{SCAB} in the FCAB sub-module, the inhibitor arc

linking P_{PAM} to T_{FMA} in the FMA sub-module and the input arc linking P_{ICMQ} to T_{SICRQ} in the IM sub-module (Figure 21).

Table 5. Input parameters represented by place marking

Place	Module	Input parameter
P_{PMS}	Master	The total number of required missions
P_{PMC}	Master	The total number and indices of platforms in the fleet
P_{CW}	Master	The total number and indices of components in platforms
P_{CAP}	Master	The cannibalisation policy
P_{SCI}	Master	The number of spares for each type of component

Table 6. Input parameters represented by transition firing delays

Transition	Module	Input parameter
T_{FMS}	FMS	The start time of each mission
T_{MGA}	FMA	The assignment time window of each mission
T_{PPM}	FMO	The length of each mission
T_{CAR}	FCAB	The length of bird life
T_{CF}	Component	Component failure time
T_{RFC}	OMA	Removal time of failed components
T_{RWC}	OMA	Removal time of cannibalised components
T_{ISC}	OMA	Installation time of spare components
T_{IWC}	OMA	Installation time of cannibalised components
T_{IRFC}	IM	Component repair time at the I-level
T_{IRFC}	DM	Component repair time at the D-level
T_{TOI}	OMA	The transportation time from the O-level to I-level
T_{TIO}	IM	The transportation time from the I-level to O-level
T_{TID}	IM	The transportation time from the I-level to D-level
T_{TDO}	DM	The transportation time from the D-level to O-level

3.8.2 Model Output

After the characteristics of all places, transitions and arcs of the HCPN model haven been identified for the modelled fleet, the initialised model can be used as a framework for a Monte Carlo simulation analysis of the modelled fleet to obtain its performance.

Four indicators are used in this paper to measure the fleet performance: mission capable rate (MCR), mission abort rate (MAR), cannibalisation rate (CAR) and maintenance cost.

The MCR is defined as the proportion of time that platforms in the fleet are capable of performing missions in a given time interval and is given by [15]:

$$MCR = \frac{\text{total MC hours}}{\text{Total possessed hours}} \times 100\% \quad (1)$$

The MC hours and possessed hours are respectively the cumulative time (in hours) that platforms are in the MC state and the total time (in hours) that platforms are possessed by the fleet. MCR is a positive indicator which means that higher values indicate better fleet performance.

The MAR can be used to measure both reliability and maintenance performance. This is a negative indicator (lower values indicate better performance) since it represents the number of missions aborted due to reliability and maintenance problems. Since missions can be aborted from the ground abort or air, the abort rate takes account of both and is defined as [15]:

$$MAR = \frac{\text{Air aborts} + \text{Ground aborts}}{\text{Total missions required}} \times 100\% \quad (2)$$

The mission ground abort rate (MGAR) and mission air abort rate (MAAR) can also be calculated:

$$MGAR = \frac{\text{Ground aborts}}{\text{Total missions required}} \times 100\% \quad (3)$$

$$MAAR = \frac{\text{Air aborts}}{\text{Total missions required}} \times 100\% \quad (4)$$

CAR is a negative measure which indicates the number of cannibalisations performed per 100 missions [15]:

$$CAR = \frac{\text{Total cannibalisations}}{\text{Total missions performed}} \times 100 \quad (5)$$

The number of cannibalisations performed, the number of ground abort events and the number of air abort events during the fleet's lifetime are respectively equal to the number of tokens in P_{CAN} , P_{MGA} and P_{MAA} . A platform's total MC hours is recorded by the length of time the related token stays in P_{PMC} . Maintenance expenditure consists of costs associated with platform failures, technician wages and spare acquisition and holding costs. The total manpower hours for repairs is equivalent to the sum of sampled delay times of T_{IRFC} and T_{DRFC} , while those for removals and installations are respectively obtained by adding the total sampled firing delays of T_{RFC} and T_{RWC} (renewals), and T_{ISC} and T_{IWC} (installations). The fleet MAR (MAAR and MGAR), CAR, MCR and maintenance cost are calculated using the collected information and equations 1-5.

3.8.3 Model Simulation

Monte Carlo simulation uses repeated random sampling and statistical analysis to obtain the required fleet performance [16]. Performing a Monte Carlo simulation needs a computer model to govern the enabling and firing processes of transitions within the HCPN model, which indicates the occurrence of all possible events relating to the fleet operation and maintenance processes such as component failures, repairs, removals and installations. Bespoke C# software was developed to encode the simulation model. The software takes in the required fleet parameters in Excel-file format and uses them to set the relevant place markings, arc expressions and distributions to generate the transition firing delays and structure of the platform module. The program then automatically produces the different modules mentioned above and connects them in order to produce the complete net.

In order to obtain the required fleet performance measures, the software simulates the produced net for a given period of time. During the simulations, a time for each transition to switch can be obtained by randomly sampling from the relevant probability distribution. For example, if a timed transition is governed by a Weibull distribution with scale parameter η and shape parameter β , for which the cumulative distribution function (CDF), $F(t)$, is:

$$F(t) = 1 - \exp \left[-\left(\frac{t}{\eta}\right)^\beta \right] \quad (6)$$

then using a random number x sampled from the $U[0,1]$ distribution, the transition firing time can be calculated using the expression:

$$t = \eta [-\ln(x)]^{1/\beta} \quad (7)$$

Numerical algorithms may be used if a probability distribution does not have an analytical inverse function, for example the central limit theorem can be used to sample from a Normal distribution [17]. The simulation program runs for a specified number (assumed to be N_s) of simulations as follows:

- Step 1. Set the simulation clock to 0;
- Step 2. Find the enabled immediate transitions and switch them;
- Step 3. Repeat step 2 until no more immediate transitions are enabled;
- Step 4. Search through each timed transition and sample firing delays for enabled transitions;
- Step 5. Switch the timed transition with the earliest firing delay and push the simulation clock to the current simulation time;
- Step 6. If the simulation clock reaches the given threshold, the current simulation is complete – go to step 7, else go to step 2;
- Step 7. If the number of simulations completed $< N_s$, go to step 1, else end.

The information for calculating the performance measures is collected during the simulation and statistics of the N_s simulations are output to Excel files to analyse the overall fleet performance. The confidence level of the simulation results depends on the number of simulations. Running a large number of simulations will lead to more

accurate results, but could be time-consuming [18]. The results are converged when carrying out further simulations does not make significant changes to the model results. A large number of simulations may be unnecessary if the results reach a desired level of convergence, a threshold for which can be specified [18]. In this paper, convergence is considered to occur when the changes in the simulation results are less than 1%, which was found to be achieved by running more than 100 simulations.

3.8.4 Model Capability

The HCPN model can concisely model different fleets and provides an ideal framework for Monte Carlo simulation, and hence provides a tool for fleet managers to:

- Estimate and predict fleet performance. Given fleet characteristics including fleet size, platform configuration, operational requirements, component failure characteristics, times relating to maintenance actions, available maintenance resources and the applied maintenance policies, the HCPN models can be used to obtain the fleet performance measures required by fleet managers;
- Select the maintenance strategy. The HCPN model can be adapted to model different aspects of fleet operation and maintenance, allowing fleet managers to study and compare the effects on fleet performance of different maintenance policies such as cannibalisation policies and platform restoration queuing disciplines, and address their strengths and weaknesses without modification of the whole model. This feature ensures that maintenance policies can be easily evaluated and selected by fleet managers according to their maintenance management objectives;
- Optimise the implementation of the selected maintenance strategy. The performance of a maintenance policy such as unrestricted cannibalisation may be affected by a variety of factors such as fleet operational requirements and spare inventories, which can vary with time. With the capability of analysing the effects of fleet factors on the selected maintenance strategy, the HCPN model can help fleet managers determine how best to implement the selected maintenance strategy;
- Manage the maintenance logistics system. Due to the flexibility of the HCPN model, different spare provision scenarios can be easily compared based on their effects on fleet performance, which can help to direct the provision of spare inventories;
- Direct the design of the fleet maintenance system. Since the HCPN model can represent different fleet maintenance systems that vary according to the number of maintenance shops and the distance between them, it can be used to help fleet managers determine the organisation and locations of the available maintenance facilities through the comparisons of fleet performance under different scenarios.

4. Application to a Simple Fleet

To demonstrate the application of the HCPN model, an example fleet is considered which consists of 50 identical and independent platforms, each of which comprises 10 components connected in series. The fleet is assumed to perform fixed, identical, daily missions for 5 years. 45 missions must be performed by the fleet each day. Each mission is performed by one MC platform, starting at 8 am and lasting 16 hours. The assignment time window for each mission is 15 minutes. Thus, the total number of missions that must be performed by the fleet is 82125 in the 5 years.

Table 7 shows the distributions of component failure time, removal/install time and repair time, the I-level repair probability for each type of component and the spare costs. The time taken to transport components between different maintenance organisations is included in the repair time. When a mission is called, if the number of MC platforms in the fleet exceeds the mission requirements, the required number of MC platforms is randomly selected to perform the mission. The random selection discipline is also applied to platform restoration queuing, selection of cannibalisation birds under the cannibalisation bird policy, and selection of NMC platforms to cannibalise from under the unrestricted cannibalisation policy. Each air abort event and ground abort event are assumed to have fixed \$300 and \$150 costs respectively. The hourly wage for component removal/installation is \$30 while that for repair is \$55.

First the effects of the unrestricted cannibalisation and cannibalisation bird policy on fleet performance are evaluated using the reliability- and operation-related performance measures introduced above, along with an

analysis of the strengths and weaknesses of different cannibalisation policies. Then two decision variables relating to the cannibalisation bird policy, the number of birds and the length of bird life, are investigated using the HCPN model.

Table 7. Example fleet parameters

Component	Failure time	I-level Repair probability	Remove time (hour)	I-level Repair time (hour)	D-level Repair time (hour)	Install time (hour)	Spare acquisition cost (\$)	Annual holding cost (\$)
1	Weibull ($\beta=1.3, \eta=2300$)	0.8	Triangle (3,4,5)	Uniform (120,130)	Uniform (150,160)	Triangle (4,5,6)	250000	5000
2	Normal ($\mu=1900, \sigma=23$)	0.7	Uniform (5.5,6.5)	Uniform (75,80)	Uniform (115,125)	Uniform (6,8)	500000	8000
3	Lognormal, ($\mu=7.6, \sigma=0.8$)	0.9	Triangle (4,5,6)	Uniform (135,145)	Uniform (165,175)	Triangle (6,6.5,7)	220000	6500
4	Weibull ($\beta=2.0, \eta=1950$)	0.75	Triangle (3,4.5,5.5)	Triangle (95,100,105)	Triangle (115,120,135)	Triangle (4,5,5.6)	330000	5500
5	Weibull ($\beta=2.5, \eta=2500$)	0.85	Triangle (3.5,4.5,6)	Triangle (120,130,135)	Triangle (170,180,195)	Triangle (3,4,4.5)	200000	3000
6	Weibull ($\beta=1.4, \eta=2100$)	0.8	Uniform (4.5,5.5)	Triangle (80,86,90)	Triangle (120,136,140)	Uniform (4,5)	460000	3800
7	Weibull ($\beta=1.8, \eta=1800$)	0.7	Uniform (5,5.5)	Triangle (130,140,150)	Triangle (160,170,180)	Uniform (5.5,6.5)	350000	7000
8	Weibull ($\beta=1.7, \eta=2300$)	0.9	Uniform (3,3.5)	Triangle (120,128,135)	Triangle (150,158,165)	Uniform (4,4.5)	400000	2000
9	Normal ($\mu=2100, \sigma=18$)	0.75	Triangle (2,3.5,4)	Uniform (150,160)	Uniform (185,195)	Triangle (3.5,4,4.5)	430000	2900
10	Weibull ($\beta=1.5, \eta=2450$)	0.85	Triangle (3,3.5,4)	Uniform (160,168)	Uniform (195,210)	Triangle (3,2,3.8,4,2)	600000	5200

4.1 Evaluating the Effect of Cannibalisation on Fleet Performance

When the cannibalisation bird policy is applied, it is assumed that the fleet only uses one cannibalisation bird at a time and a bird will be used for 30 days before its cannibalisation recovery starts. To evaluate and compare the effects of different cannibalisation policies on fleet performance, six scenarios are studied using the HCPN model:

- Scenario 1: no cannibalisation, no spares;
- Scenario 2: one cannibalisation bird, no spares;
- Scenario 3: no cannibalisation, one spare for each component;
- Scenario 4: one cannibalisation bird, one spare for each component;
- Scenario 5: unrestricted cannibalisation, no spare;
- Scenario 6: unrestricted cannibalisation, one spare for each component.

Figures 25 and 26 respectively show the changes of average fleet MCR and MAAR under different cannibalisation policies with the number of simulation performed when spares are not available. It can be seen that good convergence of the fleet performance measured by either indicator is achieved when more than 100 simulations are performed. The same conclusion is also achieved for other indicators and scenarios studied in this paper. To guarantee the accuracy of simulation results, 200 simulations are carried out for each scenario.

Table 8 summarises the associated fleet performance under the different cannibalisation and spare provisioning scenarios. The simulation results show that, for the modelled fleet, performing cannibalisation can improve the fleet performance no matter whether spares are available or not. When spares are unavailable, using one cannibalisation bird (scenario 2) or unrestricted cannibalisation (scenario 5) can respectively increase the fleet MCR by 20.6% and 24.3% while decreasing the MGAR by 76.3% and 89.4% when compared to scenario 1. However, these respectively bring 14.7% and 19.3% increases in the maintenance cost and the MAAR also respectively increases by 20.3% and 23.7%. When there is one spare for each type of component, using one cannibalisation bird (scenario 4) and unrestricted cannibalisation (scenario 6) cause the fleet MCR to respectively increase by 8.2% and 10.9%, with 66.0% and 84.0% decreases in MGAR when compared to the case where cannibalisation is not allowed (scenario 3). However, the MAAR increases by 7.5% and 9.0% respectively.




Figure 25. Convergence of average fleet MCR under different cannibalisation policies when spares are unavailable

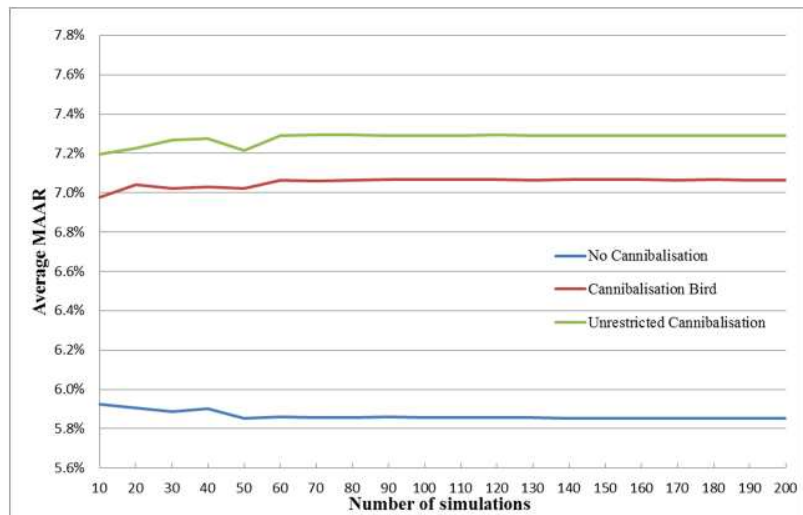


Figure 26. Convergence of average fleet MAAR under different cannibalisation policies when spares are unavailable

Table 8. Fleet performance

Fleet Performance		Scenario 1	Scenario 2	Scenario3	Scenario 4	Scenario 5	Scenario 6
MCR	Average	70.5%	85.0%	80.8%	87.4%	87.6%	89.6%
	Maximum simulated	70.9%	85.7%	81.4%	87.9%	88.0%	89.9%
	Minimum simulated	69.9%	84.4%	80.4%	86.9%	87.3%	89.4%
	Standard deviation	0.19%	0.24%	0.17%	0.18%	0.13%	0.13%
MAR	Average	26.6%	12.0%	16.1%	10.4%	9.5%	8.8%
	Maximum simulated	27.3%	12.6%	16.5%	10.9%	9.7%	9.1%
	Minimum simulated	26.0%	11.2%	15.5%	10.0%	9.1%	8.5%
	Standard deviation	0.24%	0.24%	0.20%	0.17%	0.13%	0.14%
MGAR	Average	20.7%	4.9%	9.4%	3.2%	2.2%	1.5%
	Maximum simulated	21.3%	5.6%	9.7%	3.7%	2.5%	1.8%
	Minimum simulated	20.2%	4.0%	8.9%	2.7%	1.8%	1.2%
	Standard deviation	0.21%	0.25%	0.18%	0.17%	0.13%	0.14%
MAAR	Average	5.9%	7.1%	6.7%	7.2%	7.3%	7.3%
	Maximum simulated	6.0%	7.2%	6.8%	7.3%	7.4%	7.4%
	Minimum simulated	5.8%	6.9%	6.6%	7.1%	7.1%	7.2%
	Standard deviation	0.04%	0.05%	0.04%	0.05%	0.05%	0.05%

CAR	Average	0	0.74	0	0.71	4.02	2.46
	Maximum simulated	0	0.76	0	0.73	4.24	2.62
	Minimum simulated	0	0.72	0	0.69	3.75	2.32
	Standard deviation	0	0.005	0	0.006	0.095	0.069
Cost(\$k)	Average	38437	44094	46100	48406	45874	46012
	Maximum simulated	39250	45029	46936	49194	46488	46583
	Minimum simulated	37863	43225	45377	47604	44765	45461
	Standard deviation	258	341	274	316	343	295

A perhaps unexpected result for the modelled fleet is that using only cannibalisation (scenarios 2 and 5) brings more benefit to fleet performance in terms of MCR, MGAR, MAR and maintenance cost than the case where cannibalisation is not allowed but there exists one spare for each type of component (scenario 3). Compared with scenario 3, using a cannibalisation bird (scenario 2) brings a 4.9% increase in MCR, a 25.5% decrease in the MAR, a 47.9% decrease in the MGAR and a 4.4% decrease in the maintenance cost. While using unrestricted cannibalisation (scenario 5), the fleet MCR increases 8.5% with MAR, MGAR and cost decreasing 41.0%, 76.6% and 0.5%, respectively. The reason why performance is better when using cannibalisation than when using spares might be that one spare for each type of component is not enough to maintain the fleet MCR due to the high frequency of failures (approximately 3 failures per day) and long component repair time (on average over 100 hours), whereas cannibalisation provides a constantly refreshed supply of cannibalisation sources.

The results also suggest that, for the modelled fleet, the unrestricted cannibalisation policy is more effective than the cannibalisation bird policy in improving fleet performance as measured by MCR and MAR, but leads to more cannibalisations being performed. When spares are not used, employing the unrestricted cannibalisation policy can increase the average MCR by 3.1%, and decreases the MAR and MGAR by 20.8% and 55.1% compared to the cannibalisation bird policy. However, this comes with a 440% increase in the CAR. When spares are also available, using the unrestricted cannibalisation policy causes the fleet MCR to increase 2.5%, with the MAR and MGAR respectively decreasing by 15.4% and 53.1% compared to the cannibalisation bird policy. However, the use of cannibalisation is still very high when applying the unrestricted cannibalisation policy, with the average CAR increasing about 250%.

Therefore, for the modelled fleet, it can be concluded that using either unrestricted cannibalisation or the cannibalisation bird policy is beneficial in improving the fleet performance with the MCR increased and mission cancellations reduced. The fleet could generally perform better as measured by the MCR using unrestricted cannibalisation, but fewer cannibalisations would be required when employing the cannibalisation bird policy.

4.2 Optimising the Implementation of the Cannibalisation Bird Policy

The number of cannibalisation birds that the fleet simultaneously keeps and the length of time that a platform is used as a cannibalisation bird may have a significant effect on performance under the cannibalisation bird policy. Thus, when the cannibalisation bird policy is applied it may be necessary for fleet managers to decide the appropriate number of birds to be used at any one time and the length of bird life that maximises the benefits gained by using this policy.

Table 9. Fleet performance under different number of cannibalisation birds without spares

Fleet Performance		1	2	3	4	5
MCR	Average	85.0%	86.0%	87.4%	87.6%	86.5%
	Maximum simulated	85.7%	86.5%	87.7%	87.9%	86.7%
	Minimum simulated	84.4%	85.5%	87.0%	87.4%	86.3%
	Standard deviation	0.24%	0.20%	0.14%	0.10%	0.07%
MAR	Average	12.0%	11.0%	9.7%	9.1%	9.3%
	Maximum simulated	12.6%	11.6%	10.3%	9.4%	9.5%
	Minimum simulated	11.2%	10.6%	9.5%	8.8%	9.1%
	Standard deviation	0.24%	0.21%	0.15%	0.11%	0.10%
MGAR	Average	4.9%	3.9%	2.5%	1.8%	2.0%
	Maximum simulated	5.6%	4.4%	2.9%	2.0%	2.2%

	Minimum simulated	4.0%	3.4%	2.2%	1.5%	1.8%
	Standard deviation	0.25%	0.21%	0.12%	0.09%	0.08%
MAAR	Average	7.1%	7.1%	7.2%	7.3%	7.3%
	Maximum simulated	7.2%	7.2%	7.4%	7.5%	7.4%
	Minimum simulated	6.9%	7.0%	7.1%	7.2%	7.2%
	Standard deviation	0.05%	0.05%	0.06%	0.05%	0.06%
CAR	Average	0.74	1.47	2.14	2.67	3.01
	Maximum simulated	0.76	1.49	2.17	2.72	3.07
	Minimum simulated	0.72	1.46	2.11	2.63	2.95
	Standard deviation	0.005	0.008	0.010	0.016	0.024
Cost(\$k)	Average	44094	44578	45342	45626	45627
	Maximum simulated	45029	45279	46516	46525	46523
	Minimum simulated	43225	43633	44075	44937	44688
	Standard deviation	341	308	379	316	372

First, the appropriate number of cannibalisation birds for the modelled fleet is investigated. 5 scenarios are studied in which the number of birds used differs from 1 to 5. Spares are not used in all scenarios and the bird life is set to be 30 days. The random discipline is applied to platform restoration queuing, the selection of platforms to be kept as birds and the selection of birds to cannibalise from when more than one bird contains the required working components. Note that the scenario with one bird has already been considered and the related fleet performance is shown in Table 8. The results are repeated in Table 9, for clarity, together with the fleet performance under the other 4 scenarios.

The simulation results show that the fleet CAR increases with the number of cannibalisation birds. When 5 cannibalisation birds are used, approximately 3.01 cannibalisations are used every 100 missions, which is over 4 times as many as when only one bird is used. The fleet performance as measured by the other metrics generally increases with the number of cannibalisation birds until the number of birds reaches 4, with the fleet MCR increasing 3.1% and the MGAR and MAR respectively decreasing 63.3% and 24.2%. However, when the number of cannibalisation birds increases to 5, the average fleet MCR decreases 1.3% with the MGAR increasing 11.1%, which means that it is not necessary to use 5 birds since they do not all directly contribute to fleet MCR and are also not allowed to perform missions. Therefore, in order to maximise the positive effects of the cannibalisation bird policy on fleet performance, if spares are not provisioned it might be necessary to keep 4 MC platforms as birds when a bird is used for 30 days.

To further investigate the appropriate length of bird life, 10 scenarios are studied, within which bird life length varies from 5 to 50 days. Four cannibalisation birds are assumed to be used at a time. Table 10 shows the related fleet performance under the 10 scenarios.

Table 10. Simulation results under various bird life

Fleet Performance		5	10	15	20	25	30	35	40	45	50
MCR	Average	86.4%	87.1%	87.3%	87.5%	87.5%	87.6%	87.6%	87.6%	87.5%	87.6%
	Maximum simulated	86.7%	87.4%	87.6%	87.8%	87.8%	87.9%	87.8%	87.9%	87.7%	87.8%
	Minimum simulated	86.1%	86.8%	87.1%	87.2%	87.3%	87.4%	87.3%	87.4%	87.3%	87.5%
	Standard deviation	0.13%	0.12%	0.08%	0.12%	0.11%	0.10%	0.13%	0.11%	0.11%	0.07%
MAR	Average	10.9%	9.8%	9.5%	9.2%	9.2%	9.1%	9.1%	9.0%	9.1%	9.0%
	Maximum simulated	11.2%	10.3%	9.7%	9.6%	9.5%	9.4%	9.4%	9.3%	9.4%	9.2%
	Minimum simulated	10.6%	9.5%	9.2%	8.9%	8.9%	8.8%	8.8%	8.7%	8.9%	8.9%
	Standard deviation	0.15%	0.15%	0.09%	0.14%	0.12%	0.11%	0.14%	0.12%	0.13%	0.09%
MGAR	Average	3.8%	2.6%	2.2%	1.9%	1.8%	1.8%	1.8%	1.7%	1.8%	1.7%
	Maximum simulated	4.1%	3.0%	2.4%	2.2%	2.1%	2.0%	2.1%	1.9%	2.1%	1.9%
	Minimum simulated	3.4%	2.3%	1.9%	1.7%	1.7%	1.5%	1.6%	1.4%	1.7%	1.6%
	Standard deviation	0.15%	0.13%	0.09%	0.11%	0.10%	0.09%	0.12%	0.12%	0.10%	0.07%
MAAR	Average	7.1%	7.2%	7.30%	7.3%	7.3%	7.3%	7.3%	7.3%	7.3%	7.3%
	Maximum simulated	7.3%	7.3%	7.4%	7.5%	7.4%	7.5%	7.5%	7.4%	7.4%	7.4%

	Minimum simulated	7.0%	7.1%	7.2%	7.2%	7.2%	7.2%	7.2%	7.2%	7.2%	7.2%
	Standard deviation	0.06%	0.05%	0.05%	0.06%	0.05%	0.05%	0.05%	0.04%	0.05%	0.05%
CAR	Average	10.55	6.37	4.67	3.76	3.11	2.67	2.37	2.10	1.90	1.73
	Maximum simulated	10.71	6.50	4.77	3.83	3.15	2.72	2.41	2.12	1.92	1.74
	Minimum simulated	10.31	6.24	4.61	3.69	3.07	2.63	2.35	2.08	1.88	1.71
	Standard deviation	0.094	0.058	0.036	0.028	0.021	0.016	0.013	0.010	0.010	0.009
Cost (\$k)	Average	46588	46103	45953	45891	45739	45626	45518	45528	45574	45446
	Maximum simulated	47338	46735	46802	46877	46469	46525	46608	46080	46248	46054
	Minimum simulated	45895	45449	45200	45055	45074	44937	44830	44895	44886	44856
	Standard deviation	342	340	356	385	330	316	332	301	327	301

When four cannibalisation birds are used and the bird life is no more than 50 days, the simulation results show a tendency for the fleet performance to increase as bird life length increases. When the bird life increases from 5 days to 30 days, the average fleet MCR increases 1.4% with a 52.6% decrease in MGAR and a 74.7% decrease in the CAR. This may be because the shorter bird life means more cannibalisation recoveries are performed during the fleet lifetime, which results in more cannibalisation actions and fewer platforms available to be assigned to perform missions. As the bird life increases from 30 to 50 days, the CAR decreases 35.2% and fleet MCR, MGAR and MAAR and maintenance cost remain almost stable. This may be because fewer cannibalisation recoveries are required when a bird is used for a longer time and using four birds at a time can provide sufficient cannibalisation resources. Recall that in order to avoid hangar queens a cannibalisation bird is generally not allowed to be used for more than 30 days. Considering that the fleet has similar performance when the bird life is 30 days or more, for the modelled fleet it may be better not to use a cannibalisation bird for more than 30 days if 4 birds are kept at a time.

5. Summary and Conclusions

A HCPN model of fleet operation and maintenance under multiple cannibalisation policies has been developed. A number of fleet factors are considered including mission-oriented operation, multiple level maintenance organisation, multiple platform failure modes, unrestricted cannibalisation, cannibalisation birds, component repair and spare inventories. The HCPN model consists of a number of modules, which are organised at a number of levels and connected by port places, socket places and substitution transitions. Due to the flexibility of the HCPN model, a variety of fleets, differing according to fleet size, operational requirements, maintenance organisation and maintenance policies, can be directly studied without modification of the model. In addition, due to the modular structure of the HCPN model, it can be easily extended to model practical fleets if further operation and maintenance factors must be considered.

The HCPN model delivers a powerful tool for fleet managers to design, manage and optimise the fleet maintenance system. The use of the HCPN model was demonstrated through its application to example fleets with different cannibalisation policies. The effects of the unrestricted cannibalisation and cannibalisation bird policies on fleet performance were evaluated and compared, according to a number of reliability and operation related indicators including MCR, MAR, MGAR, MAAR, CAR and maintenance cost. Further, two decision variables relating to the cannibalisation bird policy, the number of cannibalisation birds used at any one time and the length of time that a platform is kept as a bird, are investigated to optimise the implementation of the cannibalisation bird policy for the modelled fleet. These case studies show the capability of the HCPN model in helping fleet managers to make maintenance decisions and optimise the maintenance system.

6. References

1. SALMAN, S., C.R. CASSADY, E.A. POHL and S.W. ORMON. Evaluating the Impact of Cannibalisation on Fleet Performance. *Quality and Reliability Engineering International*, 2007, 23(4), 445-457.
2. MYETTE, K.M. *Cannibalisation of the F-14 and S-3A Aircraft: a Viable Logistic Alternative*. M.S. thesis, Naval Postgraduate School, 1981.
3. FISHER, W.W. Issues and Models in Maintenance Systems Incorporating Cannibalisation: a Review. *Information Systems and Operational Research*, 1990, 28(1), 154-164.
4. CASSADY, C. R., S.J. MASON and T.G. YEUNG. *Quantifying the Impact of Aircraft Cannibalisation*. Wright-Patterson Air Force Base: Air Force Research Laboratory, 2006.

5. ORMON, S.W. and C.R. CASSADY. Cannibalization Policies for a Set of Parallel Machines. In: *Proceedings of Annual Reliability and Maintainability Symposium, Jananuary 26-29, 2004*. New York: IEEE, 2004, pp. 540 - 545.
6. SHENG, J. and D.R. PRESCOTT, Using Coloured Petri Nets to Investigate Fleet Cannibalisation. *International Journal of COMADEM*, 2015,18(2), 21-32.
7. JENSEN, K. and L.M. KRISTENSEN. *Coloured Petri Nets: Modelling and Validation of Concurrent Systems*. Heidelberg: Springer, 2009.
8. PETRI, C.A. *Kommunikation mit automaten (Communication with Machines)*. Ph.D. thesis, Darmstadt University of Technology, 1962.
9. VOLOVOI, Vitali. Modelling of System Reliability Petri Nets with Aging Tokens. *Reliability Engineering & System Safety*, 2004, 84(2), 149-161.
10. CLAVAREAU Julien and Pierre-Etienne LABEAU. A Petri Net Based Modelling of Replacement Strategies Under Technological Obsolescence. *Reliability Engineering & System Safety*, 2009, 94(2), 357-369.
11. CHEW, S., S.J. DUNNETT and J.D. ANDREWS. Phased Mission Modelling of Systems with Maintenance-free Operating Periods Using Simulated Petri Nets. *Reliability Engineering & System Safety*, 2008, 93(7), 980-994.
12. ANDREWS John and Claudia FECAROTTI. System Design and Maintenance Modelling for Safety in Extended Life Operation. *Reliability Engineering & System Safety*, 2017, 163, 95-108.
13. BEABOUT, B.A. *Statistical Process Control: An Appliucation in Aircraft Maintenance Management*. M.S. thesis, Air University, 2003.
14. POWELL, M.J. *The Effects of Consolidating F-16 Phase and Cannibalization Aircraft on Key Maintenance Indicators*. M.S. thesis, United States Army Command and General Staff College, 1993
15. DUPUY, M.J., D.E. WESELY and C.S. JENKINS. Airline Fleet Maintenance: Trade-off Analysis of Alternate Aircraft Maintenance Approaches. In: *Proceedings of Systems and Information Engineering Design Symposium, Charlottesville, April 29-29, 2011*. New York: IEEE, 2011, pp.29-34.
16. RAYCHAUDHURI Samik. Introduction to Monte Carlo Simulation. In *Proceedings of Winter Simualtion Conference, Miami, December 7-10, 2008*. New York: IEEE, 91-100.
17. PRESCOTT, Darren Richard. *Safety Modelling for the Time Limited Dispatch of Aircraft*. Ph.D. thesis, Loughborough University, 2007.
18. LE, Bryant Linh Hai. *Modelling Railway Bridge Asset Management*. Ph.D. thesis, University of Nottingham, 2013.

Appendix: Abbreviations

FT	Fault Tree	I-level	Intermediate Level
PN	Petri Nets	D-level	Depot Level
CPN	Coloured Petri Nets	PMQ	Platform Maintenance Queue
HCPN	Hierarchical Coloured Petri Nets	PWMQ	Platform Waiting Maintenance Queue
MC	Mission Capable	FMS	Fleet Mission Scheduling
NMC	Non Mission Capable	FMA	Fleet Mission Assignment
MCR	Mission Capable Rate	FMO	Fleet Mission Operation
MAR	Mission Abort Rate	FCAB	Fleet Cannibalisation Bird
MAAR	Mission Air Abort Rate	PF	Platform Failure
MGAR	Mission Ground Abort Rate	OM	O-level Maintenance
CAR	Cannibalisation Rate	IM	I-level Maintenance
O-level	Organisation Level	CAN	Cannibalisation Number
DM	D-level Maintenance	NS	New Spare
PRQ	Platform Restoration Queuing	SCI	Spare Component Inventory
CRD	Component Replacement Decision	ICMQ	I-level Component Maintenance Queue
OMA	O-level Maintenance Actions	DCMQ	D-level Component Maintenance Queue

CAIC	Cannibalisation Inventory Control
SIC	Spare Inventory Control
FMS	Fleet Missions
MUS	Mission Under Schedule
PMC	Platform Mission Capable
PNMC	Platform Non Mission Capable
PAM	Platform Assigned Mission
PPM	Platform Performing Mission
POG	Platform on Ground
CAP	Cannibalisation Policy
CAB	Cannibalisation Bird
CW	Component Working
CF	Component Failed
MAA	Mission Air Abort
MGA	Mission Ground Abort
MS	Mission Start
ME	Mission End
PR	Platform Restored

SPR	Selected Platform to Repair
CAD	Cannibalisation Decision
SID	Spare Installation Decision
CCAQ	Component Cannibalisation Queue
SSC	Selected Spare Component
CRPQ	Component Replacement Queue
CRQ	Component Repair Queue
PCA	Platform Cannibalised
RFC	Removed Failed Component
RWC	Removed Working Component
CAI	Cannibalisation Inventory
IFCRQ	I-level Failed Component Repair Queue
FCTD	Failed Component to Depot
IRC	I-level Repaired Component
DRC	D-level Repaired Component
PCAD	Primary Cannibalisation Decision
PSID	Primary Spare Installation Decision
PUM	Platform Undergoing Maintenance