



Workflow Resilience for Mission Critical Systems

Mahmoud Abdelgawad^(✉), Indrakshi Ray^(✉), and Tomas Vasquez^(✉)

Department of Computer Science, Colorado State University, Fort Collins, CO 80523, USA
{M.Abdelgawad, Indrakshi.Ray, Tomas.Vasquez}@colostate.edu

Abstract. Mission-critical systems, such as navigational spacecraft and drone surveillance systems, play a crucial role in a nation’s infrastructure. Since these systems are prone to attacks, we must design resilient systems that can withstand attacks. Thus, we need to specify, analyze, and understand where such attacks are possible and how to mitigate them while a mission-critical system is being designed. This paper specifies the mission-critical system as a workflow consisting of atomic tasks connected using various operators. Real-world workflows can be large and complex. Towards this end, we propose using Coloured Petri Nets (CPN), which has tool support for automated analysis. We use a drone surveillance mission example to illustrate our approach. Such an automated approach is practical for verifying and analyzing the resiliency of mission-critical systems.

Keywords: Mission-critical Systems · Workflow · Coloured Petri Nets

1 Introduction

A mission-critical system is one whose failure significantly impacts the mission [8, 16]. Examples of mission-critical systems include navigational systems for a spacecraft and drone surveillance systems for military purposes. These systems are prone to attacks because they can cripple a nation [6]. Mission-critical systems must fulfill survivability requirements so that a mission continues in the face of attacks. Thus, this requires specifying and analyzing a mission before deployment to assess its resilience and gauge what failures can be tolerated.

A mission can be described in the form of a workflow consisting of various tasks connected via different types of control-flow operators. Researchers have addressed workflow resiliency in the context of assigning users to tasks [7, 12–15, 19, 21]. However, active attackers can compromise the capabilities of various entities. The destruction of the capabilities may cause the mission to abort or fulfill only a subset of its objectives. Analyzing resiliency considering attacker actions for mission-critical systems is yet to be explored.

Our work aims to fill this gap. We formally specify a mission in the form of a workflow, the definition of which is adapted from an earlier work [20]. A mission is often complex, and manually analyzing the workflow is tedious and error-prone. We demonstrate how such a workflow can be transformed into a Coloured Petri Net (CPN). CPN has automated tool support [17] that can be used for formal analysis.

Formal analysis may reveal deficiencies in the mission specification. Addressing such deficiencies improves the cyber-resiliency posture of the mission. We demonstrate

our approach using a mission-critical drone surveillance system. We provide a divide-and-conquer approach that helps decompose a complex workflow and shows how to analyze the sub-parts and obtain the analysis results for the total workflow.

The rest of the paper is organized as follows. Section 2 provides the formal definitions for workflow and Coloured Petri Nets. Section 3 describes a motivating example and formally represents the workflow for mission-critical systems. Section 4 defines the transformation rules from workflow to CPN along with the development of the CPN hierarchy model. Section 5 focuses on verifying the CPN so-generated and analyzing resiliency. Section 6 enumerates some related work. Section 7 concludes the paper and points to future directions.

2 Background

2.1 Workflow Definition

Typically, a workflow consists of tasks connected through operators [1–5, 20]. The syntax of the workflow adapted from [20] is defined as follows.

Definition 1 (Workflow). *A workflow is defined recursively as follows.*

$W = t_i \otimes (t | W_1 \otimes W_2 | W_1 \# W_2 | W_1 \& W_2 | \text{if}\{C\} W_1 \text{ else } W_2 | \text{while}\{C\} \{W_1\} W_2) \otimes t_f$
where

- t is a user-defined atomic task.
- t_i and t_f are a unique initial task and a unique final tasks respectively.
- \otimes denotes the sequence operator. $W_1 \otimes W_2$ specifies W_2 is executed after W_1 completes.
- $\#$ denotes the exclusive choice operator. $W_1 \# W_2$ specifies that either W_1 executes or W_2 executes but not both.
- $\&$ denotes the and operator. $W_1 \& W_2$ specifies that both W_1 and W_2 must finish executing before the next task can start.
- $\text{if}\{C\} W_1 \text{ else } W_2$ denotes the conditioning operator. C is a Boolean valued expression. Either W_1 or W_2 execute based on the result of evaluating C but not both.
- $\text{while}\{C\} \{W_1\}$ denotes iteration operator. If C evaluates to true W_1 executes repeatedly until the expression C evaluates to false.

Definition 2 (Simple and Complex Operators). *A simple operator is an operator that imposes one single direct precedence constraint between two tasks. The only simple operator is the sequence operator. All other operators are referred to as complex operators.*

Definition 3 (Simple and Compound Workflows). *A simple workflow is a workflow that consists of at most one single complex operator and finitely many simple operators. All other workflows are compound workflows. A compound workflow can be decomposed into component workflows, each of which may be a simple or a compound one.*

Definition 4 (Entities). *The tasks of a workflow are executed by active entities, referred to as subjects. The entities on which we perform tasks are referred to as objects. An entity has a set of typed variables that represents its attributes.*

Definition 5 (State of an Entity). *The values of the attributes of an entity constitute its state. The state of a subject determines whether it can execute a given task. The state of an object determines whether some task can be performed on it.*

In the remainder of this paper we abstract from objects and only consider critical systems specified as subjects.

Definition 6 (Mission). *A mission is expressed as a workflow with a control-flow, a set of subjects, a subject to task assignment relation, and some initial conditions and objectives. Formally, $\mathcal{M} = (W, S, ST, I, O)$ where W is the control-flow corresponding to the mission, S is a set of subjects, $ST \subseteq S \times \text{Tasks}(W)$ is the set of subject to task assignments, I is the set initial conditions, and O is the set of mission objectives. The subject to task assignment should satisfy the access control policies of the mission. The conditions and objectives are expressed in predicate logic.*

2.2 Coloured Petri Nets (CPN)

A Colored Petri Net (CPN) is a directed bipartite graph, where nodes correspond to places P and transitions T . Arcs A are directed edges from a place to a transition or a transition to a place. The input place of transition is a place for which a directed arc exists between the place and the transition. The set of all input places of a transition $r \in T$ is denoted as $\bullet r$. An output place of transition is a place for which a directed arc exists between the transition and the place. The set of all output places of a transition $r \in T$ is denoted as $r \bullet$. Note that we distinguish between tasks and transitions by using the label r for transitions. CPNs operate on multisets of typed objects called tokens. Places are assigned tokens at initialization. Transitions consume tokens from their input places, perform some action, and output tokens on their output places. Transitions may create and destroy tokens through their executions. The distribution of tokens over the places of the CPN defines the state, referred to as marking, of the CPN. Formally, a Non-Hierarchical CPN is defined [9, 10] as $CPN = (P, T, A, \Sigma, V, C, G, E, I)$, where P , T , A , Σ , and V , are sets of places, transitions, arcs, colors, variables, respectively. C , G , E , I are functions that assign colors to places, guard expressions to transitions, arc expressions to arcs, tokens at initialization expression respectively.

Definition 7 (Simple Workflow CPN). *A simple workflow CPN is a CPN that models a simple workflow. A simple workflow CPN has a unique input place i and a unique output place o .*

Definition 8 (CPN Module). *A CPN Module consists of a CPN, a set of substitution transitions, a set of port places, and a port type assignment function. The set of port places defines the interface through which a module exchanges tokens with other modules. Formally a CPN module is defined as in [10] as: $CPN_M = (CPN, T_{sub}, P_{port}, PT)$ where (i) CPN is a Colored Petri Net (ii) $T_{sub} \subseteq T$ is a set of substitution transitions (iii) $P_{port} \subseteq P$ is a set of port places (iv) $PT : P_{port} \rightarrow \{IN, OUT, IN \setminus OUT\}$*

Definition 9 (Simple CPN Module). *A Simple CPN Module is a CPN module in which the CPN is a Simple Workflow CPN and the interface of the module is defined by a single input port place i' and a single output port place o' .*

Definition 10. Hierarchical CPN A Hierarchical CPN is defined by the authors in [10] as: $CPN_H = (S, SM, PS, FS)$ where (i) S is a finite set of modules. places and transitions must be disjoint from all other modules' places and transitions. For a module $s \in S$ we use the notation P^s to denote the set of places of the module s . Similarly, for each of the other elements of the module s . (ii) $SM \subseteq T_{sub} \times S$ is a relation that maps each substitution transition to a sub- module. Note that [9, 10] defines $SM : T_{sub} \rightarrow S$ as a function. However, it is easier to compute SM if defined as a set. (iii) $PS(t) \subseteq P_{sock}(t) \times P_{port}^{SM(t)}$ is a port-socket relation function that assigns a port-socket relation to each substitution transition. (iv) $FS \subseteq 2^P$ is a set of non-empty fusion sets. A fusion set is a set of places that are functionally equivalent. (v) We additionally define global sets of places, transitions, arcs, colors, and variables as the union of the places, transitions, arcs, colors, and variables of each module. Furthermore, we define global initialization, arc expression, guard expression functions to be consistent with the functions defined for each module. We refer to these global elements by omitting the superscript in the notation.

Definition 11 (Module Hierarchy). A *Module Hierarchy* is a directed graph where each module $s \in S$ is a node; and for any two modules $s_1, s_2 \in S$, there exists a directed arc from s_1 to s_2 if and only if there is a substitution transition in s_1 that is mapped to the module s_2 . As represented in [10], the module hierarchy is formally defined as: $MH = (N_{MH}, A_{MH})$ where (i) $N_{MH} = S$ is the set of nodes, and (ii) $A_{MH} = \{(s_1, r, s_2) \in N_{MH} \times T_{sub} \times N_{MH} \mid r \in T_{sub}^{s_1} \wedge s_2 = SM(r)\}$

A module with no incoming arcs in the module hierarchy is referred to as a prime module.

3 Motivating Example

We shall refer to the surveillance drone mission example to illustrate the transformation framework. Let \mathcal{M}_{drone} be a mission specification that models a drone performing some data collection tasks over a region of interest. The drone has a camera that can take pictures at high and low altitudes and sensors capturing heat signals and radiation levels. The sensors are only accurate at low altitudes. There are three regions of interest: Regions A, B, and C. Region A is where the drone is regularly scheduled to perform surveillance; this region is large but close to the deployment point. Regions B and C are smaller but are further away from the deployment point. The drone is likely to be detected at low altitudes when it is in Region A. Therefore, the drone can only fly at high altitudes from where it can only use its camera. The drone may fly over regions B or C at high or low altitudes. However, due to the lack of visibility over the regions, only the heat and radiation sensors can capture meaningful data. Therefore, the drone can only collect data with its sensors over regions B and C. The mission succeeds if the drone collects data and returns to the deployment point. The control-flow of this mission is described by the task graph in Fig. 1 and given as control-flow expression below:

$$W = Init \otimes Check_Status \otimes Deploy \otimes if(instruction)\{(Fly_to_RegionB \# Fly_to_RegionC) \otimes (Measure_Radiation_level \# Measure_Heat_Signal)\}else\{Fly_to_RegionA \otimes while\{battery_level > 1\} \{Scan_Vehicles \& Scan_Construction\}\} \otimes Return_to_Base \otimes Final$$

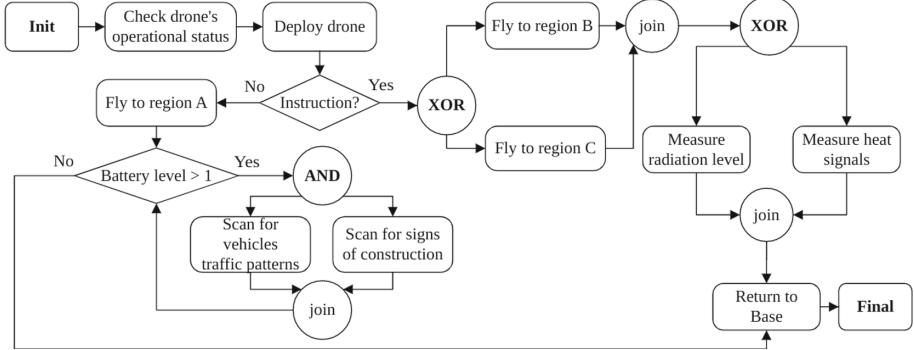


Fig. 1. Workflow of Surveillance Drone Mission.

This mission has a single entity called $drone_1$ of type $Drone$, that is, $\mathcal{S} = \{drone_1 : Drone\}$. The attributes of type $Drone$, denoted as $Drone.Attributes$, are given as: $Drone.Attributes = \{type : string; location : string; fly_enabled : bool; battery_level \in \{1,2,3,4\}; instruction_issued : bool; camera_enabled : bool; sensors_enabled : bool; data_collected : bool\}$

We define functions that return the values of various attributes. For example, $location(drone_1)$ returns the location of $drone_1$. Every task can be performed by $drone_1$. Therefore, the subject task assignment function assigns $drone_1$ to each task. $\mathcal{ST} = \{(drone_1, transition) | t \in Tasks(W)\}$. Let I be a predicate logic formula giving the initialization conditions as:

$$\begin{aligned} I = \exists s \in \mathcal{S} \mid & (type(s) = Drone) \wedge (location(s) = "base") \\ & \wedge (fly_enabled(s) = true) \wedge (battery_level(s) = 4) \wedge (instruction_issued(s) = false) \wedge \\ & (camera_enabled(s) = true) \wedge (sensors_enabled(s) = true) \wedge (data_collected = false) \end{aligned}$$

Let O be a predicate logic formula that defines the mission objectives as:

$O = \exists s \in \mathcal{S} \mid (type(s) = Drone) \wedge (location(s) = "base") \wedge (data_collected(s) = true)$
 If an attribute of the subject s is not explicitly constrained by the initial conditions or the objective, then that attribute can take on any value in its domain. Mission specification is as follows: $\mathcal{M}_{drone} = (W, \mathcal{S}, \mathcal{ST}, I, O)$.

4 Workflow to CPN Transformation Rules

Our approach maps a Mission-Specification $\mathcal{M} = (W, \mathcal{S}, \mathcal{ST}, I, O)$ to a Hierarchical Colored Petri Net $CPN_H = (S, SM, PS, FS)$. The approach follows six processes as shown in Fig. 2. The first two processes deal with the decomposition and simplification of the control-flow. The decomposition procedure partitions the control-flow into a set of disjoint expressions that each model a workflow component. Each component is either a simple or compound workflow. The simplification procedure then iterates each component substituting any nested components with tasks. A task substituting a

nested workflow in another workflow is *substitution task*. A workflow that has had all of its components substituted for tasks is said to be *simplified*. The simplification process outputs a set of simple workflows, a set of substitution tasks for each component, and a substitution relation that tracks which component was substituted by which task. The formalization process extracts from the simple workflow the sets of workflow tasks, substitution tasks, begin and end tasks, and precedence constraints over all tasks. These sets are collectively called a *formalized component*.

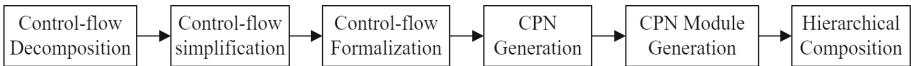


Fig. 2. Workflow to CPN Transformation Framework

The fourth process applies 7 rules to each component that, together with the remaining elements of the mission, are mapped to a Non-Hierarchical Colored Petri Net. The structure of each Non-Hierarchical CPN (places, transitions, and arcs) is semantically equivalent to the component. The fifth process generates a CPN module by adding an interface to each Non-Hierarchical CPN. The sixth process relates each module through relationships between substitution transitions, sub-modules, and port and socket places. The final output is a Hierarchical CPN representing a complete model of the original mission. Table 1 illustrates the notation used to decompose, simplify and formalize the control-flow. Note that the sets $Task(W)$, $Tasks_S(W)$, and $Task_U(W)$ are disjoint subsets of the set of all tasks $Tasks_W$. i.e. $Task(W) \cap Tasks_S(W) \cap Task_U(W) = \emptyset$

Table 1. Control-flow Notation Table

Symbol	Description
$Tasks_W$	Set of all tasks in a workflow
$Tasks(W) \subseteq Tasks_W$	Set of all workflow tasks in W
$Components$	Set of component workflows
$Tasks_S(W) \subseteq Tasks_W$	Set of substitution tasks of W
$Tasks_U(W) \subseteq Tasks_W$	Set of support tasks of W
$Tasks_B(W) \subseteq Tasks_W$	Set of begin tasks of W
$Tasks_E(W) \subseteq Tasks_W$	Set of end tasks of W
$Prec(W)$	The set of precedence constraints in W
$Substitutions$	Relation between component workflows and substitution tasks of W

Transformation Rules. The structure of the Non-Hierarchical CPN is made up of the places P' , transitions T' , and arcs A' . We define the structure of Non-Hierarchically CPNs in a similar manner as a workflow-net or process net [1,2,18]. A workflow-net is a Petri Net with a unique input place i and a unique output place o . A workflow-net has the additional property that when the output place o is connected to the input place i by a transition r' , the resulting *extended net* is strongly connected.

Rule 1: Transition Set Generation: Tasks are modeled as transitions in the Non-Hierarchical CPN. Our algorithm maps each task $t_k \in Tasks_{W'}$ to a unique transition $r_k \in T'$. Let T'_{map} be a relation between the set of tasks and the set of transitions. The pair $(t_k, r_k) \in T'_{map}$ indicates that the transition r_k is mapped from the task t_k . Through the remainder of the paper, we maintain this indexing convention. That is, r_k denotes the transition mapped from the task t_k . Let T'_{sub} , T'_U , T'_B and T'_E be subsets of T' . The set T'_{sub} is the set of substitution transitions such that $T'_{sub} = \{r_s \mid (t_s, r_s) \in T'_{map} \wedge t_s \in Tasks_S(W')\}$. The sets of support transitions T'_U , begin transitions T'_B , and end transitions T'_E are constructed similarly from the respective subsets of tasks. Consider our running example, the sets of transitions for CPN' are mapped from the sets of tasks of W' as follows:

$Tasks'_W \rightarrow T' = \{r_1, r_{s1}, r_{s5}, r_5, r_c\}$, $Tasks_U(W') \rightarrow T'_U = \{r_c\}$, $Tasks_S(W') \rightarrow T'_S = \{r_{s1}, r_{s5}\}$, $Tasks_B(W') \rightarrow T'_B = \{r_1\}$, and $Tasks_E(W') \rightarrow T'_E = \{r_5\}$. The relation between tasks and transitions is $T'_{map} = \{(t_1, r_1), (t_{s1}, r_{s1}), (t_{s5}, r_{s5}), (t_5, r_5), (t_c, r_c)\}$.

Rule 2: Place and Arc Set Generation: The set of places P' is initialized with a unique input place i and a unique output place o . The set of arcs is initialized to the empty set. For each begin transition $r_b \in T'_B$ the algorithm adds a directed arc connecting the input place i and the transition r_b to A' , i.e. $A' \leftarrow (i, r_b)$. Similarly, for each end transition $r_e \in T'_E$, the algorithm adds a directed arc connecting the transition r_b and the input place i to A' , i.e. $A' \leftarrow (r_e, o)$. Consider our running example, $r_1 \in T'_B$ implies that $A' \leftarrow (i, r_1)$ and $r_5 \in T'_E$ implies $A' \leftarrow (r_5, o)$. For each (t_k, t_j) in $Prec(W')$, we create a new place m and add the arcs (r_k, m) and (m, r_j) to A' . That is, $P' \leftarrow m$, and $A' \leftarrow (t_k, m), (m, t_j)$. If one task has two direct successors i.e. $(t_q, t_j), (t_q, t_k) \in Prec(W')$, then this denotes a point at which a split occurs. If a task has two direct predecessors, i.e. $(t_j, t_q), (t_k, t_q) \in Prec(W')$, then this denotes a point at which a join occurs. There are two types of splits and joins. There are or-splits/joins and there are and-splits/joins [18]. Or-splits should always be joined by an or-join. Similarly, an and-split should always be joined by an and-join. The or-split/join only occurs in control-flows that contain the exclusive-or operator or the conditioning operator. When an or-split is imposed by the exclusive-or operator and there is task t_q that directly precedes $t_j \# t_k$, such as $t_q \otimes (t_j \# t_k)$, the transition r_q should output to a single place m that is the input to both r_j and r_k . Then the either transition r_j or r_k will execute by consuming the output of r_q , but not both. We thus leverage the fact that given a place m that is input to two transitions, then either transition may consume the token in the place m . When an or-split is imposed by the conditioning operator, then the split is modeled by the support transition r_c where r_c routes the execution based on the evaluation of a Boolean expression. Therefore, $(t_q, t_j), (t_q, t_k) \in A'$ and an exclusive-or operator in the control-flow implies that $P' \leftarrow m$, $P' \leftarrow m'$, and $A' \leftarrow (r_q, m), (r_q, m'), (m, r_j), (m', r_k)$, however, $m = m'$. Similarly, when an or-join is imposed by the exclusive-or operator or the conditioning operator, and there is task t_q that directly succeeds $t_j \# t_k$, such as $(t_j \# t_k) \otimes t_q$, then the output place of r_j and r_k should be a single place that is the input to r_q . Therefore, $(t_j, t_q), (t_k, t_q) \in A'$ and the conditioning operator or exclusive-or is in the control-flow implies that $P' \leftarrow m$ and $P' \leftarrow m'$, and $A' \leftarrow (r_j, m), (r_k, m'), (m, r_q), (m', r_1)$, however, $m = m'$. The and-split/join only occurs in control-flows that contain the parallel-split operator. The and-split is modeled by the support transition r_g . The and-join is modeled by the support transition r_s . The structure of CPN' is formally described as:

$$P' = \{i, o, p_1, p_2, p_3, p_4\}$$

$$T' = \{r_1, r_{s1}, r_{s5}, r_5, r_c\}$$

$$A' = \{(i, r_1), (r_1, p_1), (p_1, r_c), (r_c, p_2), (r_c, p_3), (p_2, r_{s1}), (p_3, r_{s5}), (r_{s1}, p_4), (r_{s5}, p_4), (p_4, r_5), (r_5, o)\}$$

Thus far we have addressed five of the six well-behaved building blocks proposed by the Workflow Management Coalition [18] to model any control-flow. That is, the and-split, and-join, or-split, or-join, and sequence. The final building block is iteration. In the special case that the simple workflow contains the iteration operator, our algorithm adds two additional arcs to the set of arcs. One directed arc going from the support transition r_{c1} to the output place o models the iteration never executing. One directed arc going from the support transition r_{c2} to the output place of r_{c1} models the iteration continuing. One can see that our model can implement all six well-behaved building blocks proposed by the Workflow Management Coalition [18] and can therefore model any control-flow. Furthermore, we will later show that the manner in which we connect our CPN modules forms well behaved control structures [18].

Rule 3: Colors and Variable Set Generation: The color set of a simple workflow CPN consists of the different types of entities in the Simple CPN model. The types can be primitive types such as *Bool*, *Int*, or *String* or more complex user-defined types. In the drone surveillance example, there is only one subject of type *Drone*. The data type *Drone* is mapped to a record color set labeled *Drone*. Each attribute of the type *Drone* becomes a label in the record color set as:

```
color Drone = record {type : String; location : String; fly_enabled : Bool;
detected : bool; battery_level : Int ∈ {1, 2, 3, 4}; instruction_issued : Bool;
sensors_enabled : Bool; camera_enabled : Bool; data_collected : Bool}
```

Set of colors Σ' is constructed by adding any compound color sets and declaring the primitives that compose them. $\Sigma' = \{Bool, String, Int, Drone\}$ Set of variables V is declared such that there is variable for each color in Σ . $V' = \{instruction : Bool; y : String, i : Int, drone : Drone\}$.

Rule 4: Assigning Colors to Places: Each task can be performed by a drone. Therefore, each place $p \in P'$ is assigned the color $Drone \in \Sigma'$.

Rule 5: Assigning Guard Expressions: For each workflow task $t_k \in Tasks_{W'}$, we assign to the transition r_k a guard expression $G'(r_k)$ equivalent to $Pre(t_k)$. In other words, the guard expression evaluates to true if and only if the conjunction of the preconditions of t_k evaluates to true. The empty set of pre conditions is equivalent to the pre condition that always evaluates to true. Support and substitution tasks always have an empty set of pre conditions, therefore the guards of the respective transitions always evaluate to true. The only tasks of W' that have non-empty sets of pre conditions are t_1 and t_5 such that $Pre(t_1) = \{fly_enabled(drone_1) = true\}$ and $Pre(t_5) = \{fly_enabled(drone_1) = true, battery_level(drone_1) \geq 1\}$. Therefore, the guard expression function is $G'(r_1) = g_1 = fly_enabled(drone) = true$ and $G'(r_5) = g_5 = fly_enabled(drone) = true \wedge battery_level(drone) \geq 1$. Where $drone \in V'$ takes on the value of the instance of $drone_1$ when r_1 and r_2 are enabled. For all other cases in CPN' the guard expression is always true.

Rule 6: Assigning Arc Expressions: Each transition is assigned an input arc expression that evaluates to a token of the same color as the input place of the transition. Each transition is assigned an output arc expression that updates the token received over the input arc, such that the post conditions of the corresponding task evaluate to true. Substitution transitions can neither be enabled nor occur. Therefore, the arc expressions over their connected arcs have no semantic meaning. Support transitions route the execution of the workflow. For sequential and parallel executions, the output arc expressions are identical to the input arc expressions. For support transitions that evaluate a condition, each output arc models a case of the condition. The output arc of these transitions, should output the token received over the input arc only if the case evaluates to true. Otherwise, they output the empty set.

For example, for the arc $(r_c, p_2) \in A'$, is given the arc expression $E'((r_c, p_1)) = c_1 = \text{if}(instruction_issued(drone) = \text{true})\{\text{drone}\} \text{else}\{\text{empty}\}$. The arc (r_5, o) is given the arc expression $E'((r_5, o)) = e_5 = \{(location(drone) = \text{"deployment_point"}) \wedge (battery_level(drone) = battery_level(drone) - 1)\}$. Note that the syntax we use for the arc expressions is based on function notation that is easy to understand. CPN Tools has its own modeling language which we avoid using for simplicity.

Rule 7: Initialization: The initialization function I' sets the initial state of the model by assigning a multiset of tokens to each place $p \in P'$. I must satisfy the initial conditions of the mission. Recall, I calls for a subject s with the following valuation: $(type(s) = \text{Drone}) \wedge (location(s) = \text{"deployment_point"}) \wedge (fly_enabled(s) = \text{true}) \wedge (battery_level(s) = 4) \wedge (instruction_issued(s) = \text{false}) \wedge (camera_enabled(s) = \text{true}) \wedge (sensors_enabled(s) = \text{true}) \wedge (data_collected = \text{false})$

Over the input place i , the initialization function $I'(i)$ evaluates to a multiset of size one that satisfies the initial conditions of the workflow. For every other place, the initialization function evaluates to the empty set of tokens. It is important to note that for our analysis we are considering a single drone in isolation. Therefore, the initialization function of every other CPN in our final hierarchical model will evaluate to the empty set for every place. The result of applying rules 1–7 to our example W' is CPN' as described by Fig. 3. We have omitted writing each arc expression explicitly into the diagram to maintain readability.

We construct a hierarchical CPN (CPN_H) from the set of modules S and their original relationships in the workflow. Recall that a Hierarchical Colored Petri Net $CPN_H = (S, SM, PS, FS)$ consists of a set of modules S , an assignment of substitution transitions to modules SM , a relation between port places and socket places PS -where the pair $(p, p') \in PS$ signifies that p is a socket place of a substitution transition t that has been mapped to a sub module s such that p' is a port place of the same type as p ; and a set of fusion places FS . We have already defined and computed the set of modules S . The set of fusion places is empty e.g. $FS = \emptyset$. Therefore, the task is now to compute SM and PS . The module CPN'_M is the prime module (top level of the hierarchy) and every other module is a sub-module with respect to CPN' . The initial marking M_0 always evaluates to a non-empty set of tokens at the input port $i' \in P'$; and every other place is empty. The place i' has a single output arc connected to the transition r_i . Thus, all execution sequences must begin at r_i . The model has a single final transition r_f

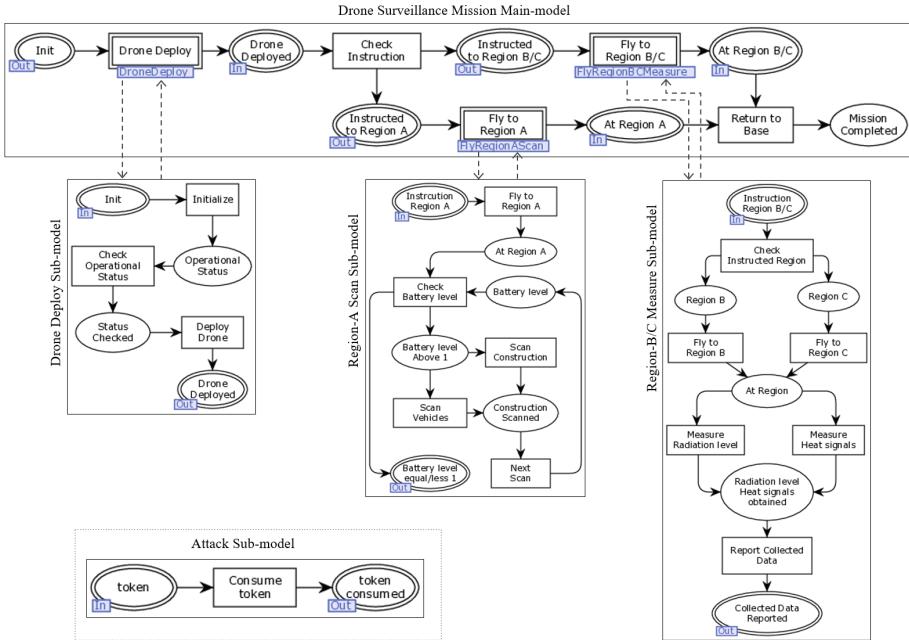


Fig. 3. CPN Models of Surveillance Drone Mission.

with a single output place o' . Therefore, any execution sequence that is complete must terminate with the transition r_f and a token in the place o' . The initial transition $r_i \in T'$ and final transition $r_f \in T'$ model the begin and end tasks of the original workflow.

5 Resiliency Analysis

An attack is an action taken by an adversary that changes the state of a subject in a way that renders it unable to perform a task it has been assigned. An attack scenario consists of the target of the attack. The set of attributes of the target, which are attackable, an attack task, and an integer limit on the number of times the attack may occur. A mission workflow is resilient to an attack scenario if, after the attack occurs, there exists a successful execution from the point at which the attack occurred.

There are various scenarios of workflow resilience, including static, decremental, and dynamic resilience [19] (see Sect. 6). In this paper, we investigate static resilience. Static resilience describes a situation where a subset of users become unavailable before the execution of the workflow. Once a subset of users is made unavailable, no user of this subset may become available again.

The analysis examines the state space of the CPN model with the attack scenario where the drone's camera fails while scanning region A since we know that it is the only location where the drone uses its camera. Table 2 reports 57 nodes and 65 arcs that are strongly connected components (SCC). It also reports that model transitions are fully

executed. However, the state space report has two dead markings (42 and 57). The first dead marking (42) has a drone in the output place and received instruction. The second dead marking (57) has a drone in the output place and did not receive an instruction. Therefore, the mission is guaranteed to terminate in success under an attack scenario.

Table 2. State Space Verification

State Space		SCC Graph		Status
#Nodes	#Arcs	#Nodes	#Arcs	Full
57	65	57	65	Live Transition None
Dead Markings [42, 57]		Dead Transition None		

We then write a program using the CPN-ML programming language [11] to analyze the mission's resilience. We then use the CPN State Space Analysis Tool to evaluate the program functions and return an execution sequence for each outcome where the attack succeeded. The program then backtracks through the shortest execution sequence and finds the node in the state space representing the state where the attack occurred. From that state, it searches for an execution sequence that results in a successful outcome (attack failure). If it finds one, it returns the execution sequence. Otherwise, it returns an empty list. The program also returns information about the set of dead markings and partitions the set into three subsets. The dead markings represent outcomes where the attack did not occur, the attack occurred, and the mission failed, and where the attack occurred, and the mission succeeded.

Table 3 summarizes the resiliency analysis result. The first column describes the scanning iteration of Region A. The first row shows 37 total dead markings (DM), 2 dead markings that represent outcomes where the attack did not occur (DNA), 10 dead markings that represent outcomes where the attack occurred but failed (DAF), 25 dead markings that represent outcomes where the attack occurred and succeeded (DAS). The algorithm found paths that reach only 4 of the 25 DAS markings (RDAS).

Table 3. Static Resilience Analysis Results

Scan Iteration Number	Dead Markings (DM)	Attack Didn't Occur (DNA)	Attack Occurred but Failed (DAF)	Attack Occurred and Succeeded (DAS)	Reachable DAS (RDAS)
1	37	2	10	25	4
2	33	2	10	21	0
3	16	2	9	5	0
4	12	2	10	0	0

The program inspects the set of DNA markings and returns that the mission is correct and valid. The DNA markings represent the original outcomes of the mission. The size of this set should be the same as the original set of dead markings. For the 25

failed outcomes, the program found a successful execution sequence for four outcomes. We compare the successful execution sequence with the failed execution sequence for each of these four outcomes. It leads us to find the state where the execution sequences diverge and, from that state, ensure that the transition and binding element that leads to success always occur. For instance, node 42 corresponds to a failed outcome where the attack happened at node 3. From node 3, our program found a path to node 57 - a successful outcome. We now compare the path from node 3 to node 57 and the path from node 3 to node 42. We find that the execution sequences diverge at node 5. From node 5, proceeding to node 10 or node 9 is possible. The change in state from node 5 to node 10 results from the transition with the variable instruction bound to the value *false*. The change in state from node 5 to node 9 results from the transition with the variable instruction bound to the value *true*. We now know that the attack succeeds if the drone's battery is less than 3 units in the state represented by node 5, where transition *receiveinstruction* is enabled. Transition receive instruction models the workflow routing based on an instruction issued.

In Iteration 2, the state space is re-calculated, and the analysis is repeated. The second row of Table 3 describes the statistics generated by the program's second iteration. The result is four fewer outcomes where the attack succeeded. The program returns a set of nodes representing the states where the attack occurred, and no path to a successful outcome exists. Note that these are not dead markings. From each of these nodes, there is a path to a dead marking that represents the failure of the mission (success of the attack).

In Iteration 3, the state space is re-calculated, and the analysis is repeated. The third row of Table 3 summarizes the statistics generated by the program's second iteration. We expected the set of outcomes where the attack failed to remain at 10, now 9. The reduced size of the failed attack set means the outcome is lost because the drone is initialized with its instruction set to *false*. Thus the attack at node 3, which was found to be resilient, never executes. We focus on the 5 markings where the attack occurred and succeeded. After inspecting all five markings, we find the attack succeeds because it is delivered when the drone has just enough battery to perform one additional pass over Region A. We can eliminate these failed outcomes by increasing the requirement on the loop over Region A from more than one unit of battery to more than two units of battery. It turns out that the drone always has some reserve battery if it is attacked.

In Iteration 4, the state space is re-calculated, and the analysis is repeated. The third row of Table 3 summarizes the statistics generated by the program's second iteration. Since the set of attack successes is empty, we can be confident that the attack can not succeed under the restricted workflow.

In summary, we have shown how to assess a mission's resilience and find the conditions required for the mission to succeed. We have also shown how to restrict a workflow to improve its resilience.

6 Related Work

The literature on workflow resiliency problems introduces solutions to address the unavailability [13, 15, 19]. Our work argues that the workflow resiliency problem can

sometimes be viewed as unavailability and degradation. In other words, attacks do not permanently remove subjects from service; they decrease their capabilities. Consider the drone surveillance example; a failure in the drone's camera affects the termination and success of the workflow. Regarding resilience based on availability, one can assume whether the camera's loss is critical enough to remove the drone from the workflow. The drone should be kept since its other sensors can complete the workflow.

Wang *et al.* [19] introduce three types of resilience, static, decremental, and dynamic resilience. Static resilience refers to a situation in which users become unavailable before the workflow executes, and no users may become available during the execution. Decremental resiliency expresses a situation where users become unavailable before or during the execution of the workflow, and no previously unavailable users may become available during execution, while dynamic resilience describes the situation where a user may become unavailable at any time; a previously unavailable user may become available at any time. The different types of resilience formulations capture various types of attack scenarios.

Mace *et al.* [13, 15] propose a quantitative measure of workflow resiliency. They use a Markov Decision Process (MDP) to model workflow to provide a quantitative measure of resilience. They refer to binary classification, such as returning an execution sequence if one exists and declaring the workflow resilient; or returning false and declaring the workflow not resilient. The authors show that the MDP models give a termination rate and an expected termination step.

7 Conclusion

This paper emphasizes the workflow resiliency of the task degradation problem, specifically for mission-critical cyber systems. We presented a set of rules that formally transforms workflow represented by a mission into Coloured Petri Nets (CPNs). We then solved various analysis problems related to the resiliency of mission-critical such as cyber-attacks. We developed an approach based on formalization rules that address the complexity of mission workflows, simplify them, and transform them into simple CPNs. These simple CPNs are then modulated and combined as a hierarchy CPN model.

We applied the approach to a drone surveillance system as an illustrative example. We used the CPN tools to run verification and reachability analysis. The results showed that the workflow resiliency problem could sometimes be unavailability and degradation. A workflow subject is not permanently removed from service when an attack occurs; it decreases its capabilities. However, the mission can continue, and the workflow can be completed. We have shown how to assess a mission's resilience and find the conditions to succeed. We have also shown how to restrict a workflow to improve its resilience.

Future work will focus on extending the generated model to account for multiple subjects and investigating decremental and dynamic resilience. We will design a set of algorithms corresponding to the transformation rules. Our end goal is to automate the process of verification and resilience analysis of workflows.

Acknowledgements. This work was supported in part by funding from NSF under Award Numbers CNS 1715458, DMS 2123761, CNS 1822118, NIST, ARL, Statnett, AMI, NewPush, and Cyber Risk Research.

References

1. Aalst, W.M.P.: Verification of workflow nets. In: Azéma, P., Balbo, G. (eds.) ICATPN 1997. LNCS, vol. 1248, pp. 407–426. Springer, Heidelberg (1997). https://doi.org/10.1007/3-540-63139-9_48
2. van der Aalst, W.: Structural characterizations of sound workflow nets (1996)
3. Arpinar, I.B., Halici, U., Arpinar, S., Doğanç, A.: Formalization of workflows and correctness issues in the presence of concurrency. *Distrib. Parallel Databases* **7**(2), 199–248 (1999). <https://doi.org/10.1023/A:1008758612291>
4. Bride, H., Kouchnarenko, O., Peureux, F.: Verifying modal workflow specifications using constraint solving. In: Albert, E., Sekerinski, E. (eds.) IFM 2014. LNCS, vol. 8739, pp. 171–186. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10181-1_11
5. Bride, H., Kouchnarenko, O., Peureux, F., Voiron, G.: Workflow nets verification: SMT or CLP? In: ter Beek, M.H., Gnesi, S., Knapp, A. (eds.) FMICS/AVoCS -2016. LNCS, vol. 9933, pp. 39–55. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45943-1_3
6. Chong, J., Pal, P., Atigetchi, M., Rubel, P., Webber, F.: Survivability architecture of a mission critical system: the DPASA example. In: 21st Annual Computer Security Applications Conference (ACSAC 2005), pp. 10–pp. IEEE (2005)
7. Fong, P.W.L.: Results in workflow resiliency: complexity, new formulation, and ASP encoding, pp. 185–196. Association for Computing Machinery, New York (2019)
8. Houliotis, K., Oikonomidis, P., Charchalakis, P., Stipidis, E.: Mission-critical systems design framework. *Adv. Sci. Technol. Eng. Syst. J.* **3**(2), 128–137 (2018)
9. Jensen, K., Kristensen, L., Wells, L.: Coloured Petri Nets and CPN Tools for modelling and validation of concurrent systems. *STTT* **9**, 213–254 (2007). <https://doi.org/10.1007/s10009-007-0038-x>
10. Jensen, K., Kristensen, L.M.: Coloured Petri Nets: Modelling and Validation of Concurrent Systems, 1st edn. Springer, Heidelberg (2009). <https://doi.org/10.1007/b95112>
11. Jensen, K., Kristensen, L.M.: CPN ML programming. In: Jensen, K., Kristensen, L.M. (eds.) Coloured Petri Nets, pp. 43–77. Springer, Heidelberg (2009). https://doi.org/10.1007/b95112_3
12. Mace, J., Morisset, C., van Moorsel, A.: Modelling user availability in workflow resiliency analysis. In: Proceedings of the 2015 Symposium and Bootcamp on the Science of Security, HotSoS 2015, pp. 1–10. ACM (2015)
13. Mace, J.C., Morisset, C., van Moorsel, A.: Quantitative workflow resiliency. In: Kutyłowski, M., Vaidya, J. (eds.) ESORICS 2014. LNCS, vol. 8712, pp. 344–361. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11203-9_20
14. Mace, J.C., Morisset, C., van Moorsel, A.: WRAD: tool support for workflow resiliency analysis and design. In: Crnkovic, I., Troubitsyna, E. (eds.) SERENE 2016. LNCS, vol. 9823, pp. 79–87. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45892-2_6
15. Mace, J.C., Morisset, C., Moorsel, A.: Impact of policy design on workflow resiliency computation time. In: Campos, J., Haverkort, B.R. (eds.) QUEST 2015. LNCS, vol. 9259, pp. 244–259. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-22264-6_16
16. Ponsard, C., Massonet, P., Molderez, J.F., Rifaut, A., van Lamsweerde, A., Van Tran, H.: Early verification and validation of mission critical systems. *Formal Methods Syst. Des.* **30**(3), 233–247 (2007). <https://doi.org/10.1007/s10703-006-0028-8>

17. Ratzer, A.V., et al.: CPN tools for editing, simulating, and analysing coloured Petri nets. In: van der Aalst, W.M.P., Best, E. (eds.) ICATPN 2003. LNCS, vol. 2679, pp. 450–462. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-44919-1_28. <http://cpn-tools.org>
18. van der Aalst, W.M., ter Hofstede, A.H.: Verification of workflow task structures: a Petri-net-baset approach. *Inf. Syst.* **25**(1), 43–69 (2000). [https://doi.org/10.1016/S0306-4379\(00\)00008-9](https://doi.org/10.1016/S0306-4379(00)00008-9). <https://www.sciencedirect.com/science/article/pii/S0306437900000089>
19. Wang, Q., Li, N.: Satisfiability and resiliency in workflow authorization systems. *ACM Trans. Inf. Syst. Secur.* **13**(4), 1–35 (2010)
20. Yang, P., Xie, X., Ray, I., Lu, S.: Satisfiability analysis of workflows with control-flow patterns and authorization constraints. *IEEE Trans. Serv. Comput.* **7**(2), 237–251 (2014). <https://doi.org/10.1109/TSC.2013.31>
21. Zavatteri, M., Viganò, L.: Last man standing: static, decremental and dynamic resiliency via controller synthesis. *J. Comput. Secur.* **27**(3), 343–373 (2019)