

On Testing 1-Safe Petri Nets

Guy-Vincent Jourdan, Gregor v. Bochmann

School of Information Technology and Engineering (SITE)

University of Ottawa

800 King Edward Avenue, Ottawa, Ontario, Canada, K1N 6N5

{gvj, bochmann}@site.uottawa.ca

Abstract—Formal models are often considered for software systems specification, and are helpful for verifying that certain properties are respected, or for automatically generating the implementation code corresponding to the model, or again for conformance testing, for the automatic generation of test cases to check an implementation against the formal specification. Variations of Finite State Machine (FSM) models have been mostly used for conformance testing, while the otherwise very popular formal model of Petri Nets is seldom mentioned in this context. In this paper, we look at the question of conformance testing when the model is provided in the form of a 1-safe Petri Net. We provide a general framework for conformance testing, and give algorithms for deriving test cases under different assumptions: Besides the adaptation of methods originally developed for FSMs which lead to exponentially long test sequences, we have identified cases for which polynomial testing algorithms for free-choice Petri nets can be provided. These results are significant when modeling concurrent systems, as exemplified by workflow modeling.

Conformance testing, fault model, 1-safe Petri nets, free-choice Petri nets, automatic test generation

I. INTRODUCTION

Software systems are notoriously incorrect, a problem that only gets worse when dealing with distributed systems. In order to help producing better systems, one common suggestion is to create a formal model of the specification of the system, and then test whether a given implementation *conforms* to the specification, for some suitable definition of conformance. Ideally the tests are automatically generated based on the formal specification. This kind of approach has been mostly researched using Finite State Machines (FSM) as the formal model (see e.g. [1] for a survey). More recently, the same questions have been asked for concurrent systems, for which FSM do not offer a good model. When modeling concurrent systems with FSMs for test generations, *multi-ports* FSM [2,3] and *Partial Order Input/Output Automata* [4] have been used. When modelling concurrent systems in general, a popular formalism among researchers are Petri Nets (see e.g. [5] for a survey). Surprisingly, little has been done in the area of conformance testing of systems specified as Petri Nets, even though Petri Nets have long been used in practice and have for example influenced the design of some UML schema and have been used as a basis for workflow modeling [6]. Other formalisms beside the already mentioned FSMs have been used in the context of testing distributed systems (for example for Labeled

Transition Systems [7] or Message Sequence Charts [8]), but in this context Petri Nets have mostly been used for fault diagnosis (see e.g. [9]). An interesting classification of testing criteria for Petri nets was provided by Zhu and He [10], but without testing algorithms.

In this paper, we investigate the question of automatically testing Petri Nets, to ensure that an implementation of a specification provided as a Petri Net is correct. We focus on *1-safe* Petri Nets, that is, Petri Net for which a place never holds more than one token, and on free-choice Petri Nets. This model is well suited for workflows [6], which is our target application. Thus, the traditional *transitions* of Petri Nets are seen as *tasks* of a workflow. In this context, we provide a precise fault model, capturing what kind of changes could make a candidate implementation not conforming to the specification: some of the specified flow constraints between the tasks may be missing, or some unspecified flow constraints between tasks may be added. free-choice Petri nets are also useful to model for example flows in networks of processors [11]. We provide a precise framework for the conformance question in Section II C. We then provide our main results, our testing algorithms, in Section III. We first show that testing Petri Nets can be reformulated as a special case of FSM testing. This approach makes sense since FSM testing has been so extensively studied. Unfortunately, the complexity of the transformation of a Petri Net into an FSM is exponential in the worst case. We then study some particular cases for *1-safe* free-choice Petri nets, for which a polynomial testing algorithm can be proposed. In Section III B., we provide a polynomial-time algorithm to test implementations that have only missing flow constraints between tasks. We also provide a polynomial-time algorithm for testing implementations with only additional *input flows* between tasks in Section III C. 1). We show in Section III C. 2) that the problem is more difficult for additional *output* flows. We conclude in Section IV. The basic concepts and definitions are introduced in Section II.

II. BASIC CONCEPTS AND ASSUMPTIONS

In this section, we give the definition of Petri Nets and we explain the assumptions that we make about testing environment.

A. Petri nets

Definition 1: Petri Nets, input/outputs, traces, executable sets, markings. A Petri Net is a 4-tuple $N=(P,T,F,M_0)$ where P is a finite set of places, T is a finite set of transitions (or *tasks* in our context), $F \subseteq (P \times T) \cup (T \times P)$ is a set of arcs (the *flows* in our context). A marking is a mapping from P to the natural numbers, indicating the number of tokens in each place. We write M_0 for the initial marking of the net. For a task $t \in T$, we note $\bullet t = \{p \in P | (p,t) \in F\}$ the set of *inputs* of t , and $t \bullet = \{p \in P | (t,p) \in F\}$ the set of *outputs* of t . Similarly, for a place $p \in P$, we note $\bullet p = \{t \in T | (t,p) \in F\}$ and $p \bullet = \{t \in T | (p,t) \in F\}$.

A task t is enabled for execution if all inputs of t contain at least one token. When a task is executed the marking changes as follows: The number of tokens in all inputs of t decreases by one, and the number of tokens in all outputs of t increases by one. A *trace* of a Petri Net is a sequence of tasks that can be executed starting from the initial marking in the order indicated, without executing any other tasks. An *executable set* of tasks is a multiset of tasks whose task elements can be sequenced into a trace. Clearly, for each trace there is a unique executable set, but several traces may correspond to the same executable set, in which case we say that the traces are *equivalent*. A *marking of a trace t* is the marking obtained after the execution of t from the initial marking. We say that t *marks* P if the place P contains at least one token in the marking of t .

Definition 2: Petri Net equivalence. Two Petri Nets $PN_1=(P_1,T_1,F_1,M_{10})$ and $PN_2=(P_2,T_2,F_2,M_{20})$ are said to be equivalent if, $T_1=T_2$ and they have the same set of traces.

Definition 3: Free-choice Petri Nets. A Petri Net $PN=(P,T,F,M_0)$ is *free-choice* if and only if for all places $p \in P$, we have either $|p \bullet| < 2$ or $(\bullet p) = \{p\}$.

Definition 4: 1-Safeness (k -safeness). A marking of a Petri Net is *1-safe* (resp. *k -safe*) if the number of tokens in all places is at most one (resp. k). A Petri Net is k -safe if the initial marking is k -safe and the marking of all traces is k -safe. (Note: We consider in this paper only 1-safe Petri Nets).

Proposition 1: In a 1-safe free-choice Petri Net $PN=(P,T,F,M_0)$, if for some task $t \in T$ and some place $p \in P$ such that $p \in \bullet t$ and $|\bullet t| > 1$, there is no trace that marks $(\bullet t) \setminus p$ but not p then removing the input (p,t) from F defines a Petri Net PN' which is equivalent to PN .

Proof: Let $PN=(P,T,F,M_0)$ be a 1-safe free-choice Petri Net, let $t \in T$ be a task of PN and $p \in P$ a place of PN such that p is in $\bullet t$, $|\bullet t| > 1$ and there is no trace of PN that marks $(\bullet t) \setminus p$ but not p . Let PN' be the Petri Net obtained by removing (p,t) from F in PN . We show that PN' is equivalent to PN . Suppose they were not equivalent, that is,

PN and PN' do not have the same set of traces. Since we have only removed a constraint from PN , clearly every trace of PN is also a trace of PN' , therefore PN' must accept traces that are not accepted by PN . Let Tr be such a trace. Since the only difference between PN and PN' is fewer input flows on t in PN' , necessarily t is in Tr . Two situations might occur: either Tr is not a trace of PN because an occurrence of t cannot fire in PN (i.e. $\bullet t$ is not marked at that point in PN), or another task t' can fire in PN' but not in PN . In the former case, because we have only removed (p,t) , it means that $(\bullet t) \setminus p$ is marked but $\bullet t$ is not, a contradiction with the hypothesis. In the latter case, if t' can fire in PN' but not in PN then $\bullet t'$ must be marked in PN' and not in PN . Again, the only difference being t not consuming a token in p in PN' , it follows that t' can use that token in PN' but not in PN . In other words, $p \in \bullet t'$, but then $|p \bullet| > 1$ and thus $\bullet(p \bullet) = \{p\}$ (PN is free-choice), a contradiction with $|\bullet t| > 1$.

Proposition 2: In a 1-safe Petri Net $PN=(P,T,F,M_0)$, for any tasks $t, t' \in T$ and for any place $p \in P$ such that $p \in \bullet t$ and $p \in \bullet t'$, if there is a trace Tr of PN containing both t and t' , with t occurring before t' in Tr , then there is a task $t'' \in T$ in Tr (possibly $t'' = t'$) that consumes the token put in p by t .

Proof: (immediate from the definition) When Tr comes to executing t' , if the token that must be present in p is not the one put there by t then, because of 1-safety, necessarily the token put by t has already been consumed by some task t'' .

B. Assumptions about the specification, implementation and testing environment

We assume that the specification of the system under test is provided in the form of a 1-safe Petri Net. Moreover, we assume that there is a well identified initial marking (initial state). The goal of our test is to establish the conformance of the implementation to the specification, in the sense of trace equivalence defined in Definition 2.

For the system under test, we make the assumption that can be modeled as a 1-safe Petri Net. In addition, we make one of the following assumptions: (a) that the number of reachable markings in the implementation is not larger than in the specification, or (b) that the number of transitions (tasks) in the implementation is not larger than in the specification. Assumption (a) corresponds to a similar restriction commonly made for conformance testing in respect to specifications in the form of state machines, where one assumes that the number of states of the implementation is not larger than the number of states of the specification. Without such an assumption, one would not be sure that the implementation is completely tested, since some of the (implementation) states might not have been visited. If one wants to avoid this assumption, it is possible to assume that the number of markings in the implementation does not exceed the number of markings in the specification by more than some upper bound k , however, we do not address this question in this paper.

Regarding the testing environment, we assume that the system under test provides an interface which provides the following functions:

- At any time, the tester can determine which transitions are enabled.
- The tester can trigger an enabled transition at will. Moreover, transitions will only be executed when triggered through the interface.
- There is a reliable reset, which bring the system under test into what corresponds to the initial marking.

A good example of applications compatible with these assumptions are Workflow Engines. Workflow processes are usually specified by some specification language closely related to Petri Nets. Workflow engines are used to specify (and then enforce) the flow of tasks that are possible or required to perform for a given activity. In that setting, the set of possible tasks is well known, and a given task can be initiated if and only if the set of tasks that directly precede are finished. In a computer system, it typically means that at any time, the set of tasks that can be performed are enabled (accessible from the user interface) while the tasks that cannot be performed are not visible. It is thus possible to know what tasks are enabled, and choose one of them to be performed. In [6], workflow are modeled using 1-safe Petri nets.

C. Fault Model

We assume that the difference between the system under test and the reference specification can be explained by a certain types of faults, as explained below. We do not make the single-fault assumption, thus the difference between the implementation and the specification may be due to several occurrences of these types of faults.

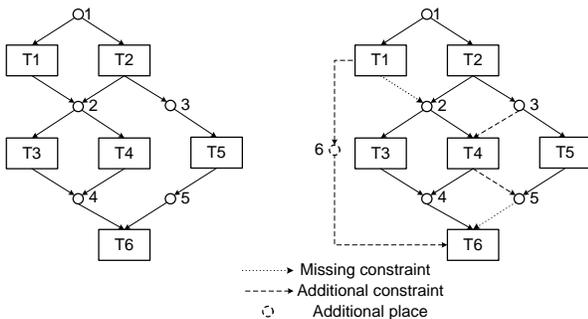


Figure 1: An example of a Petri Net specification (left) and a Petri Net implementation (right) with several types of faults

1. **Missing output flow:** a task does not produce the expected output, that is, does not put the expected token in a given place. In Figure 1, the specification (left) says that task T1 must produce an output into place 2.

However, in the implementation (right), T1 fails to produce this output; this is a missing output flow fault.

2. **Missing input flow:** a task does not require the availability of a token in a given place to fire. In Figure 1, the specification (left) says that task T6 must take an input from place 5. However, in the implementation (right), T6 does not require this input; this is a missing input flow fault.
3. **Additional output flow:** a task produces an output into an existing place, that is, places a token into that place while the specification does not require such output. In Figure 1, the specification (left) says that task T4 does not produce an output into place 5. However, in the implementation (right), T4 produces this output; this is an additional output flow fault.
4. **Additional input flow:** a task requires the availability of a token in a given existing place while the specification does not require such a token. In Figure 1, the specification (left) says that task T4 does not require an input from place 3. However, in the implementation (right), T4 does require this input; this is an additional input flow fault.

In the rest of the paper, we are interested only in faults that actually have an impact on the observable behavior of the system, that is, create a non equivalent Petri Net (in the sense of Definition 2). It may prevent a task to be executed when it should be executable, and/or make a task executable when it should not. It is clear that it is possible to have faults as defined here that create an equivalent system, either because they simply add onto existing constraints, or replace a situation by an equivalent one.

We assume that none of these faults change the basic assumptions regarding the system under test. In particular, each task still has at least one dependency (with the initial state) and cannot fire on its own.

When considering faults of additional output flows, new tokens may “appear” anywhere in the system every time a task is executed. This raises the question of 1-safeness of the faulty implementation. One may assume that the faulty implementation is no longer 1-safe, and thus a place can now hold more than one token. Or one may assume that the implementation is still 1-safe despite possible faults. Another approach is to assume that the implementation’s places cannot hold more than one token and thus one token can “overwrite” another. Finally, and this is our working assumption, one may assume that violation of 1-safeness in the system under test will raise an exception that we will catch, thus detecting the presence of a fault under test.

III. TESTING

A. Using the testing techniques for finite state machines

In this section we consider the use of the testing techniques developed for state machines. We can transform any 1-safe Petri Net into a corresponding state machine where each marking of the Petri Net corresponds to a state of the state machine, and each transition of the Petri Net corresponds to a subset of the state transitions of the state

machine. The state machine can be obtained by the classical marking graph construction method.

Algorithm 1: FSM-based Testing

1. From the initial marking, enumerate every possible marking that can be reached. Call E this set of markings.
2. Create a finite state machine A having $|E|$ states labeled by the elements of E and a transition between states if and only if in the Petri Net it is possible to go from the marking corresponding to the source state to the marking corresponding to the target state by firing one a task. Label the FSM's transition with this task. (An example of this transformation is given in Figure 2).
3. Generate a checking sequence for A using one of the well known techniques of checking sequence construction for finite state machines [1].
4. For each *recognized state* of the implementation, verify that no task is enabled that should not be enabled (by driving the implementation into that state and listing the tasks that are enabled)

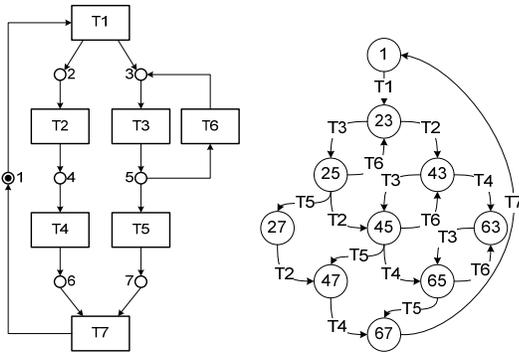


Figure 2: A simple Petri Net and the corresponding finite state machine

Algorithm 1 will produce an exhaustive verification of the implementation.

Proposition 3: *Under the assumption that the number of markings for the implementation is not larger than for the specification, Algorithm 1 detects every possible combination of faults (in the fault model) resulting in an implementation that is not equivalent to the specification.*

Proof: *If a combination of faults results in a non-equivalent implementation, this means that one of the following cases occurs:*

1. *A given marking of the specification cannot be reached in the implementation. This will be caught by the state recognition portion of the checking sequence algorithm.*

2. *From a given marking, a task that should be enabled is not, or it is enabled but executing it leads to the wrong marking. This will be caught by the transition verification portion of the checking sequence algorithm.*
3. *From a given marking, a task that should not be enabled, is. This is typically not tested by checking sequences algorithms, but this problem will be detected by Step 4 of Figure 2. Indeed, in our settings, we are able to list all tasks enabled in a given marking. The checking sequence algorithm will give us a means to know how to put the implementation into a state corresponding to a given marking of the specification, thus Step 4 will detect any addition transition.*

Unfortunately, the number of markings may be exponentially larger than the size of the original Petri Net. In the following, we consider situations where more efficient testing methods can be used.

B. Testing free-choice Petri nets for missing flow faults

In some cases, it is possible to have more efficient testing algorithms. One such case is when the only possible faults are of type missing flow, that is, missing input flow or missing output flow in a free-choice Petri net. In this case, it is possible to check each input and output flow individually.

Intuitively, the principles are the following: for detecting a missing input flow, we note that a task T normally requires all of its inputs to be marked to be enabled. For example, in Figure 3, left, the task T is not enabled because even though places $(\bullet T)p$ are all marked, place p is not. However, if the input flow from p to T is missing in the implementation (center), then is the same situation, with places $(\bullet T)p$ marked but place p not marked, T becomes enabled. The testing algorithm will thus, for all tasks T and all places p in $\bullet T$, mark all places in $(\bullet T)p$ and check if T is enabled. However, as shown in Figure 3 (right), in some case the missing input flow may not be detected; this may happen when at least one place p' in $(\bullet T)p$ was not successfully marked, because of some other (missing output) faults in the implementation.

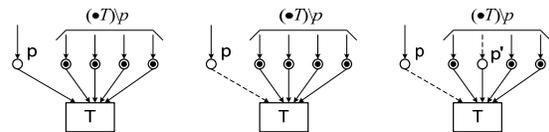


Figure 3: Missing input flow. When $(\bullet T)p$ are marked but not p , T is not enabled in the specification (left), but it is enabled in the implementation (center); however, some additional faults may interfere (right).

The idea for testing for missing output flows is the following: if a task T outputs a token into a place p , and a task T' requires an input from p to be enabled, then marking all the places in $(\bullet T')$ by executing T to mark p (among other

transitions) will enable T' (Figure 4, left). If T is missing the output flow towards p , then T' will not be enabled after attempting to mark all the places in $(\bullet T')p$ because p will not actually be marked (Figure 4, center). Again, the situation can be complicated by another fault, such as a missing input flow between p and T' , in which case T' will be enabled despite the missing output (Figure 4, right).

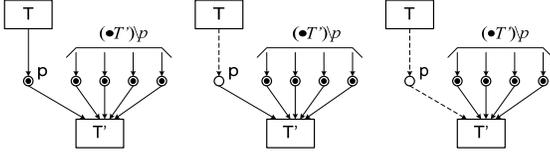


Figure 4: Missing output flow. Correct situation (left), missing output flow of T (center), and interference by a missing input flow to T' .

The problems of the interference between several faults, as indicated in the right of Figure 3 and Figure 4, will be addressed indirectly by the proof of Proposition 4 which will show that the verification for missing input might fail because of another missing output, and the verification for missing output might fail because of a missing input, but they cannot both fail at the same time.

Formally, the algorithm for testing for missing input flows is the following:

Algorithm 2: Testing for missing input flows

1. For all task T
2. For all place p in $(\bullet T)$
3. If there is a trace S that marks $(\bullet T)\backslash p$ but not p
4. Reset the system and execute S
5. Verify *NOT-ENABLED*(T)

The runtime of Algorithm 2 is clearly polynomial in respect to the size of the given Petri Net. As we will see in the proof of Proposition 4, this algorithm is not sufficient on its own, since it can miss some missing input flows, when combined with missing output flows. It must be run in combination with Algorithm 3 which tests for missing output flows:

Algorithm 3: Testing for missing output flows

1. For all task T
2. For all place p in $(T\bullet)$
3. For all task T' in $(p\bullet)$
4. If there is trace S that contains T and marks $(\bullet T')$
5. Reset the system and execute S
6. Verify *ENABLED*(T')

It is also clear that the runtime of Algorithm 3 is polynomial in respect to the size of the given Petri Net.

Proposition 4 shows that running both algorithms is enough to detect all missing flow faults in the absence of other types of faults. To guarantee the detection of these faults, we must assume that the implementation only contains faults of these types.

Proposition 4: *Executing both Algorithm 2 and 3 will detect any faulty implementation of a free-choice Petri net specification that has only missing input and/or output flow faults.*

Proof: *If an implementation is not faulty, then clearly neither algorithm will detect a fault.*

Assume that there is a task T with a missing input flow from a place p . If there is no trace S that marks $(\bullet T)\backslash p$ but not p , then Proposition 1 shows that the input flow is unnecessary and the resulting Petri Net is in fact equivalent to the original one. We therefore assume that there is such a trace S . If after executing S successfully $(\bullet T)\backslash p$ is indeed marked, then T will be enabled and the algorithm will detect the missing input constraint. If after executing S T is not enabled, that means that $(\bullet T)\backslash p$ is in fact not marked. The problem cannot be another missing input for T , since it would mean fewer constraints on S , not more, and wouldn't prevent T to be enabled. Thus, the only remaining option is that some task T' in $\bullet((\bullet T)\backslash p)$ did not mark the expected place in $(\bullet T)\backslash p$ when executing S , that is, T' has a missing output flow. In conclusion, Algorithm 2 detects missing input flows, except when the task missing the input constraint has another input flow which is not missing but for which the task that was to put the token has a missing output flow.

Assume now that there is a task T with a missing output flow to a place p . If this output flow is not redundant, then there is a task T' that has p as an input flow and that will consume the token placed there by T . Such a T' will be found by Algorithm 3. Because of the missing output flow, normally T' will not be enabled after executing S and the algorithm will catch the missing flow. However, it is still possible for T' to be enabled, if it is missing the input flow from p , too. In conclusion, Algorithm 3 detects missing output flow, except when the missing output flow is to a place that has an input flow that is missing too.

To conclude this proof, we need to point out that each algorithm works, except in one situation; but the situation that defaults Algorithm 2 is different from the one that defaults Algorithm 3. In the case of Algorithm 3, we need a place that has lost both an input and an output flow, while in the case of Algorithm 2, we need a place that has lost an output flow but must have kept its input flow. Thus, by running both algorithms, we are guaranteed to detect all problems, Algorithm 3 catching the problems missed by Algorithm 2, and vice versa.

C. Testing for additional constraint faults

Testing for additional flow faults is more difficult than testing for missing flow faults because it may involve tasks that are independent of each other according to the specification. Moreover, the consequence of this type of faults can be intermittent, in that an additional output flow fault may be cancelled by an additional input flow fault, leaving only a short window of opportunity (during the execution of the test trace) to detect the fault. In the case of additional input flow faults without other types of faults, we can still propose a polynomial algorithm, but we cannot check for additional output flow faults in polynomial time, even if no other types of faults are present.

1) Testing for additional input flow faults

An additional input flow fault can have two different effects: it may prevent a task from being executed when it should be executable according to the specification, because the task expects an input that was not specified, or it may prevent some other task from executing because the task with the additional input flow has unexpectedly consumed the token. Figure 5 illustrates the situation: task T has an additional input flow from place p (left). In the case where $\bullet T$ is marked, but p is not (center), T is not enabled, although it should. If p is marked too (right), then T can fire, but then T' cannot anymore, even though it should be enabled.

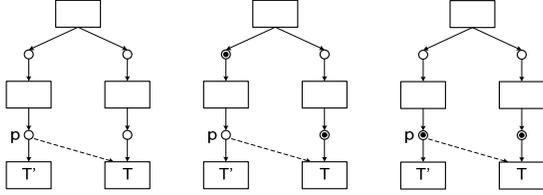


Figure 5: Additional input flow fault: T has an additional input flow from p (left). This may prevent T from firing (center), or, when T fires, T' becomes disabled (right).

A more general description is illustrated in Figure 6: in order to mark $\bullet T$, some trace is executed (the dashed zone in the figure). While producing this trace, some other places will also become marked (for example p') while other places are unmarked (for example p). This normal situation is shown on the left. If T has an additional input constraint from p (center), then after executing the same trace, the faulty $\bullet T$ will not be enabled. If T has an additional input constraint from p' (right), then the faulty $\bullet T$ is still marked after generating the trace, so T is enabled, however, if it fires it will disable T' .

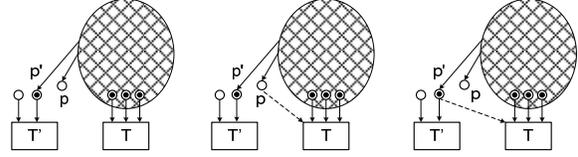


Figure 6: Additional input flow fault (see explanation above)

Consequently, in order to detect these faults, in the absence of any other type of faults, we use an algorithm that works in two phases: first, a trace is executed that should mark $\bullet T$ and verifies that T is indeed enabled. This shows that either T does not have an additional input constraint, or that the additional input place happens to be marked as well. So the second phase checks that the places that are marked by this trace (and that are not part of $\bullet T$) are not being unmarked by the firing of T .

Algorithm 4 gives the details.

Algorithm 4: Testing for additional input flows

1. For all task T
2. Find a trace S that marks $\bullet T$
3. Reset the system and execute S
4. Verify $ENABLED(T)$
5. For all place p not in $\bullet T$ which is marked by S
6. If there is a trace S' containing T and another task T' in $p\bullet$ ($T \neq T'$)
7. Reset the system and verify that S' can be executed
8. Else if there is a trace S'' marking $\bullet T$ but not p
9. Reset the system and execute S''
10. Verify $ENABLED(T)$

Proposition 5: Executing

Algorithm 4 will detect any faulty implementation of a free-choice Petri net specification that has only additional input flow faults.

Proof: If a task T has an additional input from a place p and that fault is not caught at line 4, it necessarily means that p is marked by trace S , and expected to be so because we consider only additional input flow faults. Such a case will be dealt with by lines 5 through 10. If the fault has an impact (i.e. if it leads to a wrong behavior of the implementation), then there must be a task T' in $p\bullet$ and a trace containing both T and T' that is executable according to the specification but not in the implementation. Again, because we consider only additional input flow faults, any trace containing both T and T' will fail, since when the second task is executed, the token in p has already been consumed as often as it has been set, thus the second task

will not be enabled. Lines 6 and 7 of the algorithm ensure that such a trace will be found and run, therefore a fault with an impact when p is marked will be caught. Finally, the fault might have no impact when p is marked, but if there is another way to enable T without marking p , via some trace S'' then T would not be enabled when S'' is executed. Line 9 and 10 of the algorithm address this case.

2) Testing for additional output flow faults

The fault of an additional output flow to an existing place might enable a task when that task should not be enabled. Detecting this type of faults is more difficult than additional input flows, and we cannot do it in polynomial time even in the absence of other types of fault.

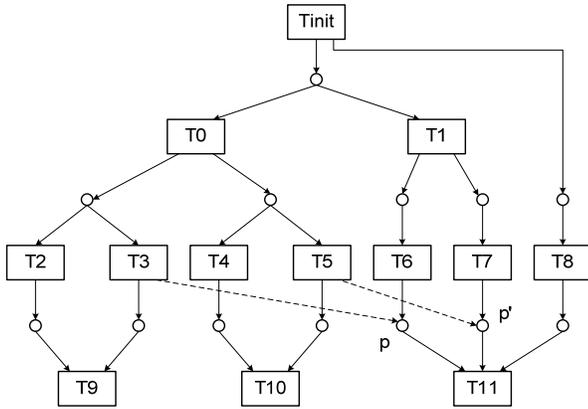


Figure 7: Additional output flow faults: the additional output flows from $T3$ to p and from $T5$ to p' can be detected only if these two transitions are included in the trace

When considering additional output flows, additional tokens may be placed anywhere in the system every time a task is executed. Thus, any trace may now mark any set of places anywhere in the net, so we cannot have any strategy beyond trying everything. This is illustrated in Figure 7, with two additional output flows, one from $T3$ to p and one from $T5$ to p' . Neither $T3$ nor $T5$ are prerequisite to $T11$, so to exhibit the problem requires executing tasks that are unrelated to the problem's location, namely $T3$ and $T5$. Both $T3$ and $T5$ are branching from a choice, and of the 4 possible combinations of choices, there is only one combination that leads to the detection of the problem. For detecting an arbitrary set of additional output flow faults, it is therefore necessary to execute traces for all possible combination of choices.

Because of these difficulties, we cannot suggest a polynomial algorithm for these types of faults.

IV. CONCLUSION

In this paper, we look at the question of conformance testing when the model is provided in the form of a 1-safe Petri Net. We first provide a general framework for testing whether an implementation conforms to a specification which is given in the form of a 1-safe Petri Nets. The types of errors that we consider in this paper include faults of missing or additional flows (inputs to, or outputs from transitions). We provide a general, but inefficient algorithm for testing these faults, derived from methods originally developed for FSMs. We then identify special types of faults for which polynomial testing algorithms can be provided.

This paper is an initial step towards a fully developed Petri Net testing.

ACKNOWLEDGMENT

This work has been supported in part by grants from the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

- [1] Lee D, Yannakakis M (1996) Principles and methods of testing finite state machines – a survey. *Proceedings of the IEEE*, 84(8):1089–1123.
- [2] G. Luo, R. Dssouli, G. v. Bochmann, P. Ventakaram and A. Ghedamsi, Generating synchronizable test sequences based on finite state machines with distributed ports, *IFIP Sixth International Workshop on Protocol Test Systems*, Pau, France, September 1993, pp. 53-68.
- [3] J. Chen, R. Hieron and H. Ural, Resolving observability problems in distributed test architecture, *FORTE 2005, LNCS 3731*, 2005, pp. 219-232.
- [4] G. v. Bochmann, S. Haar, C. Jard and G.-V. Jourdan, Testing Systems Specified as Partial Order Input/Output Automata. *TestCom 2008, LNCS 5047*, pages 169-183, Springer 2008.
- [5] T. Murata (1989) Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4): 541–580.
- [6] W. van der Aalst, T. Weijters, L. Maruster, (2004) Workflow mining: discovering process models from event logs. *IEEE Transactions on Knowledge and Data Engineering*, 16(9): 1128–1142.
- [7] P. Bhateja, P. Gastin, and M. Mukund, A Fresh Look at Testing for Asynchronous Communication. *ATVA 2006, LNCS 4218*, pages 369-383, Springer 2006.
- [8] P. Bhateja, P. Gastin, M. Mukund and K. Narayan Kumar, Local Testing of Message Sequence Charts Is Difficult. *Fundamentals of Computation Theory, 2007, LNCS 4639*, pages 76-87, Springer 2007.
- [9] S. Haar (2008) Law and Partial Order. Nonsequential Behaviour and Probability in Asynchronous Systems. *Habilitation à diriger les recherches, INRIA*. <http://www.lsv.ens-cachan.fr/~haar/HDR.pdf>.
- [10] H. Zhu and X. He, (2002) A methodology of testing high-level Petri nets. *Information and Software Technology*, 44(8): 473-489
- [11] J. Desel and J. Esparza (1995) Free choice Petri nets *Cambridge Tracts In Theoretical Computer Science*; Vol. 40. ISBN:0-521-46519-2.