

This form is a summary description of the model entitled “LeafsetExtension” proposed for the Model Checking Contest @ Petri Nets. Models can be given in several instances parameterized by scaling parameters. Colored nets can be accompanied by one or many equivalent, unfolded P/T nets. Models are given together with property files (possibly, one per model instance) giving a set of properties to be checked on the model.

Description

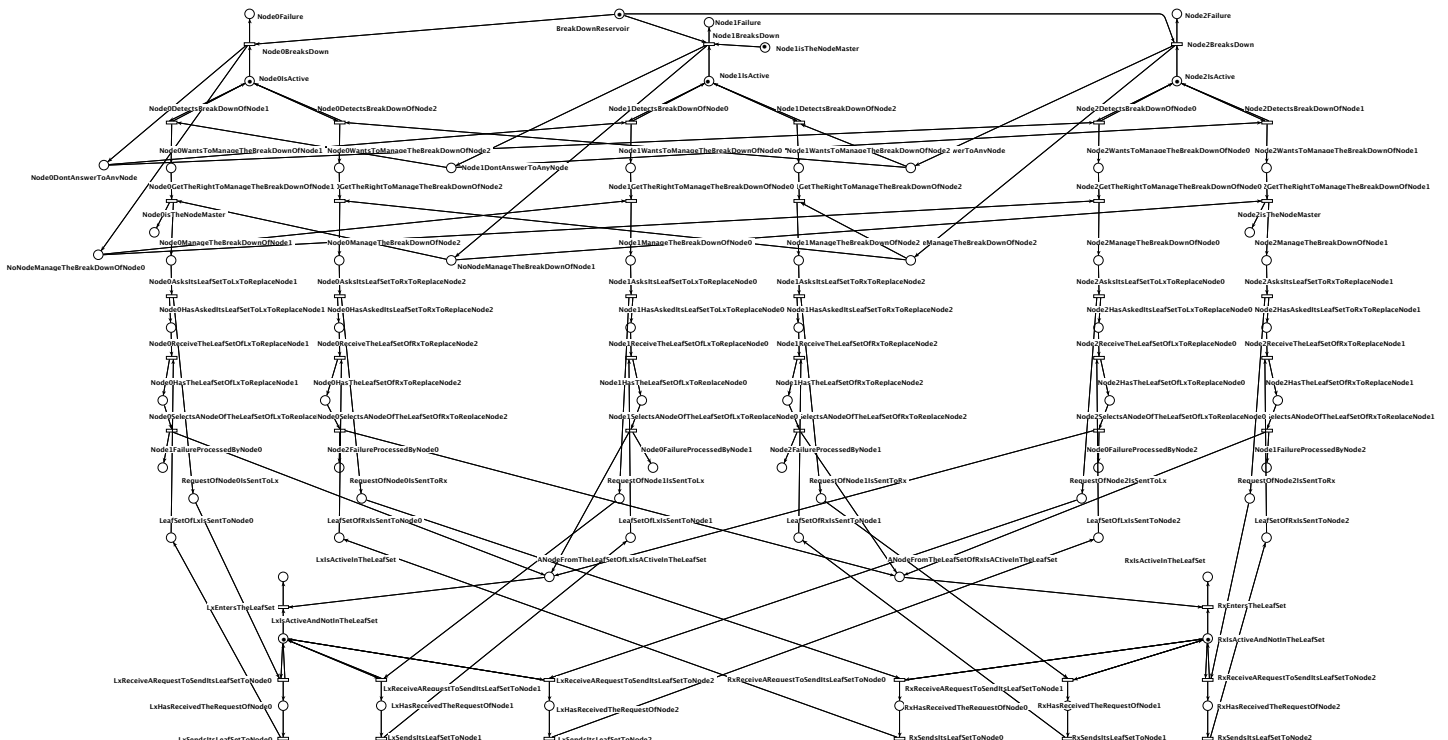
This Petri net models the extension of a LeafSet in a pastry [2] DHT (Distributed Hash Table) where nodes store data among leafset: each value is stored on a “master” node surrounded by $\frac{S}{2}$ nodes on its right and $\frac{S}{2}$ on its left ; each of them managing copies [1].

When a node in the leafset crashes, another one in the leafset detects its absence processes to the leafset extension so that the number of copies in the DHT remains constant (then, the value may be decided at a majority of $\frac{S}{2} + 1$ identical copies). The model focuses on the protocol that extend the leafset when a ode is down only.

We consider the following modeling hypotheses:

- place BreakDownReservoir represents the maximum number of crashes,
- the crash is handled by the node which first detects the breakdown,
- the leafset is extended to the right or to the left, according to the position of the crashed node in the leafset,
- if the master node crashes, then the closest node to the right or to the lefts is selected to be the new master (depending on the side where the node having detected the crash is).

Initially set to support one crash (initial marking of place BreakDownReservoir was set to 1 token), the initial marking of this place has been modified to make its reachability graph significantly more complex (useful for the MCC).



Graphical representation for $S = 2$ and $C = 1$

References

1. S. Legtchenko, S. Monnet, P. Sens, and G. Muller. Churn-resilient replication strategy for peer-to-peer distributed hash-tables. In R. Guerraoui and F. Petit, editors, *Stabilization, Safety, and Security of Distributed Systems*, 11th International Symposium, SSS 2009, Lyon, France, November 3-6, 2009. Proceedings, volume 5873 of *Lecture Notes in Computer Science*, pages 485–499. Springer, 2009
2. A. I. T. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In R. Guerraoui, editor, *Middleware 2001, IFIP/ACM International Conference on Distributed Systems Platforms Heidelberg*, Germany, November 12-16, 2001, Proceedings, volume 2218 of *Lecture Notes in Computer Science*, pages 329–350. Springer, 2001.

Scaling parameter

Parameter name	Parameter description	Chosen parameter values
S, C	S is the size of the leafset, C is the number of possible crashes	(8, 2), (8, 3), (8, 4), (16, 2), (16, 3), (16, 4), (32, 2), (32, 3), (32, 4), (64, 2), (64, 3), (64, 4)

Size of the model

Parameter	Number of places	Number of transitions	Number of arcs
(8, *)	462	409	1 324
(16, *)	1 542	1 449	4 668
(32, *)	5 622	5 449	17 500
(64, *)	21 462	21 129	67 740

Structural properties

ordinary — all arcs have multiplicity one	✓
simple free choice — all transitions sharing a common input place have no other input place	✗ ^(a)
extended free choice — all transitions sharing a common input place have the same input places	✗ ^(b)
state machine — every transition has exactly one input place and exactly one output place	✗ ^(c)
marked graph — every place has exactly one input transition and exactly one output transition	✗ ^(d)
connected — there is an undirected path between every two nodes (places or transitions)	✓ ^(e)
strongly connected — there is a directed path between every two nodes (places or transitions)	✗ ^(f)
source place(s) — one or more places have no input transitions	✓ ^(g)
sink place(s) — one or more places have no output transitions	✓ ^(h)
source transition(s) — one or more transitions have no input places	✗ ⁽ⁱ⁾
sink transitions(s) — one or more transitions have no output places	✗ ^(j)
loop-free — no transition has an input place that is also an output place	✗ ^(k)
conservative — for each transition, the number of input arcs equals the number of output arcs	✗ ^(l)
subconservative — for each transition, the number of input arcs equals or exceeds the number of output arcs	✗ ^(m)

^(a) stated by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

^(b) transitions “Node0BreaksDown” and “Node1BreaksDown” share a common input place “BreakDownReservoir”, but only the former transition has input place “Node0IsActive”.

^(c) stated by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

^(d) stated by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

^(e) stated by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

^(f) stated by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

^(g) there exist 2 source places, e.g., place “BreakDownReservoir”.

^(h) At least place “Node0failure”; confirmed by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

⁽ⁱ⁾ stated by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

^(j) stated by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

^(k) stated by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

^(l) stated by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

^(m) stated by [CÆSAR.BDD](#) version 3.5 on all 12 instances (see all aforementioned parameter values).

nested units — places are structured into hierarchically nested sequential units⁽ⁿ⁾ X

Behavioural properties

safe — in every reachable marking, there is no more than one token on a place X^(o)
 dead place(s) — one or more places have no token in any reachable marking ?^(p)
 dead transition(s) — one or more transitions cannot fire from any reachable marking X^(q)
 deadlock — there exists a reachable marking from which no transition can be fired ✓^(r)
 reversible — from every reachable marking, there is a transition path going back to the initial marking X^(s)
 live — for every transition t , from every reachable marking, one can reach a marking in which t can fire X^(t)

Size of the marking graphs

Parameter	Number of reachable markings	Number of transition firings	Max. number of tokens per place	Max. number of tokens per marking
(8, 2)	172 248 ^(u)	359 351 ^(v)	?	≥ 14 ^(w)
(8, 3)	$> 10\,693\,747$ ^(x)	$> 31\,614\,301$ ^(y)	?	≥ 15 ^(z)
(8, 4)	?	?	?	≥ 16 ^(aa)
(16, 2)	?	?	?	≥ 22 ^(ab)
(16, 3)	?	?	?	≥ 23 ^(ac)
(16, 4)	?	?	?	≥ 24 ^(ad)
(32, 2)	?	?	?	≥ 38 ^(ae)
(32, 3)	?	?	?	≥ 39 ^(af)
(32, 4)	?	?	?	≥ 40 ^(ag)
(64, 2)	?	?	?	≥ 70 ^(ah)
(64, 3)	?	?	?	≥ 71 ^(ai)
(64, 4)	?	?	?	≥ 72 ^(aj)

⁽ⁿ⁾ the definition of Nested-Unit Petri Nets (NUPN) is available from <http://mcc.lip6.fr/nupn.php>

^(o) By construction, initial marking of place BreakDownReservoir > 1 ; confirmed by CÆSAR.BDD version 3.5 on all 12 instances (see all aforementioned parameter values).

^(p) stated by CÆSAR.BDD version 3.5 to be false on 6 instance(s) out of 12, and unknown on the remaining 6 instance(s).

^(q) by construction; confirmed by CÆSAR.BDD version 3.5 on 6 instance(s) out of 12.

^(r) Stated by PROD on April 2021 which is not a surprise, the protocol blocks in several situations corresponding to final configurations.

^(s) By construction, since the model ends.

^(t) By construction, since the model ends.

^(u) Stated by PROD on April 2021.

^(v) Stated by PROD on April 2021.

^(w) lower bound given by the number of initial tokens.

^(x) Stated by PROD on April 2021.

^(y) Stated by PROD on April 2021.

^(z) lower bound given by the number of initial tokens.

^(aa) lower bound given by the number of initial tokens.

^(ab) lower bound given by the number of initial tokens.

^(ac) lower bound given by the number of initial tokens.

^(ad) lower bound given by the number of initial tokens.

^(ae) lower bound given by the number of initial tokens.

^(af) lower bound given by the number of initial tokens.

^(ag) lower bound given by the number of initial tokens.

^(ah) lower bound given by the number of initial tokens.

^(ai) lower bound given by the number of initial tokens.

^(aj) lower bound given by the number of initial tokens.