

*This form is a summary description of the model entitled “Raft” 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

Consensus algorithms ensure the replication of a log among a set of faulty servers, thus enabling the implementation of fault-tolerant services replicated on several machines. Nowadays, such algorithms are crucial for developers of cloud services that must survive data center outages.

The standard consensus algorithm, the famous Paxos designed by L. Lamport, is notoriously complex and thus difficult to implement correctly. Therefore, a new consensus algorithm named Raft has been proposed recently [OO14]. Raft aims at providing the same guarantees as Paxos while being simpler to grasp and thus easier to implement. The rapid adoption of Raft by several companies (Facebook, Hashicorp, CoreOS) illustrates the need for a “simpler Paxos”.

In Raft, time is divided in terms and at most one server can be elected as leader for a given term. The leader is in charge of committing client requests on a majority of servers, and server crashes or network partitions can trigger a leader election for a new term.

The present P/T nets are derived from the formal specification in LNT of the Raft consensus protocol, where servers can crash and restart at any moment. There are several instances for various numbers of Raft servers, ranging from 2 to 10. Each instance models two client requests and allows up to two crashes per server.

This LNT specification was used as a case study for the DLC (*Distributed LNT Compiler*) tool [EL15]. DLC is able to automatically generate for Raft a distributed implementation in C, which can be deployed on a cluster of machines communicating via sockets.

The LNT specification was translated to LOTOS, and then to an interpreted Petri net using the CADP toolbox. Finally, the present P/T net was obtained by stripping out all dataflow-related information (variables, types, assignments, guards, etc.) from the interpreted Petri net, leading to a NUPN (*Nested-Unit Petri Net*) model translated to PNML using the CÆSAR.BDD tool.

## References

[OO14] Diego Ongaro and John Ousterhout. *In Search of an Understandable Consensus Algorithm*. Proceedings of the USENIX Annual Technical Conference (USENIX ATC 14), Philadelphia, PA, USA, June 2014, pages 305–319. USENIX Association. <http://www.usenix.org/conference/atc14/technical-sessions/presentation/ongaro>

[EL15] Hugues Evrard and Frédéric Lang. *Automatic Distributed Code Generation from Formal Models of Asynchronous Concurrent Processes*. Proceedings of the 23rd Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP’15), Turku, Finland. IEEE, 2015. <http://hal.inria.fr/hal-01086522>

See also:

- Raft web site: <http://raftconsensus.github.io>
- Verification based on CADP found a missing transition in a Raft specification: <http://groups.google.com/forum/#!topic/raft-dev/you-wOUx-gnA>

## Scaling parameter

Parameter name	Parameter description	Chosen parameter values
$N$	$N$ : number of Raft servers	2, 3, 4, 5, 6, 7, 8, 9, and 10

## Size of the model

Parameter	Number of places	Number of transitions	Number of arcs	Number of units	HWB code
$N = 2$	28	52	159	7	2-6-18
$N = 3$	53	98	328	12	2-11-33
$N = 4$	88	160	563	19	2-18-54
$N = 5$	133	238	864	28	2-27-81
$N = 6$	188	332	1231	39	2-38-114
$N = 7$	253	442	1664	52	2-51-153
$N = 8$	328	568	2163	67	2-66-198
$N = 9$	413	710	2728	84	2-83-249
$N = 10$	508	868	3359	103	2-102-306

## Structural properties

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

## Behavioural properties

<b>safe</b> — in every reachable marking, there is no more than one token on a place .....	✓ (p)
<b>dead place(s)</b> — one or more places have no token in any reachable marking .....	? (q)
<b>dead transition(s)</b> — one or more transitions cannot fire from any reachable marking .....	? (r)
<b>deadlock</b> — there exists a reachable marking from which no transition can be fired .....	? (s)

(a) stated by [CÆSAR.BDD](#) version 2.2 on all 9 instances.

(b) stated by [CÆSAR.BDD](#) version 2.2 on all 9 instances.

(c) stated by [CÆSAR.BDD](#) version 2.6 on all 9 instances.

(d) stated by [CÆSAR.BDD](#) version 2.2 on all 9 instances.

(e) stated by [CÆSAR.BDD](#) version 2.2 on all 9 instances.

(f) stated by [CÆSAR.BDD](#) version 2.2 on all 9 instances.

(g) from place 1 one cannot reach place 0.

(h) place 0 is a source place.

(i) stated by [CÆSAR.BDD](#) version 2.2 on all 9 instances.

(j) stated by [CÆSAR.BDD](#) version 2.2 on all 9 instances.

(k) stated by [CÆSAR.BDD](#) version 2.2 on all 9 instances.

(l) stated by [CÆSAR.BDD](#) version 2.2 on all 9 instances.

(m) 3 transitions are not conservative, e.g., transition 0.

(n) transition 0 is not subconservative.

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

(p) safe by construction – stated by the [CÆSAR](#) compiler.

(q) stated by [CÆSAR.BDD](#) version 3.3 to be false on 7 instance(s) out of 9, and unknown on the remaining 2 instance(s).

(r) stated by [CÆSAR.BDD](#) version 2.2 to be false on 7 instance(s) out of 9, and unknown on the remaining 2 instance(s).

(s) stated by [CÆSAR.BDD](#) version 2.2 to be false on 5 instance(s) out of 9, and unknown on the remaining 4 instance(s).

**reversible** — from every reachable marking, there is a transition path going back to the initial marking .....?  
**live** — for every transition  $t$ , from every reachable marking, one can reach a marking in which  $t$  can fire .....?

### Size of the marking graphs

Parameter	Number of reach-able markings	Number of tran-sition firings	Max. number of tokens per place	Max. number of tokens per marking
$N = 2$	7381 <sup>(t)</sup>	?	1	6 <sup>(u)</sup>
$N = 3$	3.38196e+07 <sup>(v)</sup>	?	1	11 <sup>(w)</sup>
$N = 4$	2.96586e+12 <sup>(x)</sup>	?	1	18 <sup>(y)</sup>
$N = 5$	5.93795e+18 <sup>(z)</sup>	?	1	27 <sup>(aa)</sup>
$N = 6$	2.90944e+26 <sup>(ab)</sup>	?	1	38 <sup>(ac)</sup>
$N = 7$	$\geq 3.36507e+35$ <sup>(ad)</sup>	?	1	51 <sup>(ae)</sup>
$N = 8$	$\geq 9.21571e+45$ <sup>(af)</sup>	?	1	66 <sup>(ag)</sup>
$N = 9$	$\geq 4.83981e+52$ <sup>(ah)</sup>	?	1	83 <sup>(ai)</sup>
$N = 10$	$\geq 1.84049e+66$ <sup>(aj)</sup>	?	1	102 <sup>(ak)</sup>

---

<sup>(t)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(u)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(v)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(w)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(x)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(y)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(z)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(aa)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(ab)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(ac)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(ad)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(ae)</sup> stated by [CÆSAR.BDD](#) version 2.6.  
<sup>(af)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(ag)</sup> stated by [CÆSAR.BDD](#) version 2.6.  
<sup>(ah)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(ai)</sup> stated by the [CÆSAR](#) compiler.  
<sup>(aj)</sup> stated by [CÆSAR.BDD](#) version 2.2.  
<sup>(ak)</sup> stated by the [CÆSAR](#) compiler.