

This form is a summary description of the model entitled "GPU Forward Progress" 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 model is joint work by Google UK, Princeton University, and University of California, Santa Cruz. It describes how a GPU computation may ensure termination in the absence of inter-workgroup forward progress guarantees. GPU programming generally lacks forward progress guarantees (needed for synchronization) between thread workgroups, such that application relying on inter-workgroup synchronization have to implement a slow fallback path to ensure termination [1]. However, recent work shows that GPUs typically offer semi-fair forward progress models [2], such as linear occupancy bound execution (LOBE), which pave the way to reliable, portable inter-workgroup synchronization.

This collection of P/T nets was derived from an LNT model of the GPU Forward Progress. Each instance was first translated to LOTOS, and then to an interpreted Petri net using the CADP toolbox. Finally, a P/T net was obtained by stripping out all data-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.

Each instance of the model is parameterized by the number N of threads. Each instance is also parameterized by its version V, which specifies how the NUPN has been produced from the LOTOS specification. V is either equal to "a" if the NUPN has been generated *after* applying all the structural and data-flow optimizations of the CÆSAR compiler for LOTOS, or to "b" if the NUPN has been generated *before* these optimizations.

References

[1] Raph Levien. *Prefix sum on Vulkan*. April 2020. https://raphlinus.github.io/gpu/2020/04/30/prefix-sum.html# forward-progress.

[2] Tyler Sorensen, Hugues Evrard and Alastair F. Donaldson. *GPU schedulers: how fair is fair enough?* Proceedings of the 29th International Conference on Concurrency Theory (CONCUR'18), Beijing, China. LIPIcs 118, pages 23:1–23:17, Schloss Dagstuhl, 2018. Available from http://drops.dagstuhl.de/opus/volltexte/2018/9561

Scaling parameter

Parameter name	Parameter description	Chosen parameter values	
(N,V)	N is the number of times each component is replicated and V is the version defined	$\{4, 8, 12, 16, 20, 24, 28, 32, 36, 40\} \times \{a, b\}$	
	above		

Parameter	Number of	Number of	Number of	Number of	HWB code
	places	transitions	arcs	units	
N = 4, V = a	24	29	101	7	2-6-18
N = 4, V = b	112	117	277	11	6-6-35
N = 8, V = a	40	49	181	11	2-10-30
N = 8, V = b	188	197	477	19	10-10-59
N = 12, V = a	56	69	261	15	2-14-42
N = 12, V = b	264	277	677	27	14-14-83
N = 16, V = a	72	89	341	19	2-18-54
N = 16, V = b	340	357	877	35	18-18-107
N = 20, V = a	88	109	421	23	2-22-66
N = 20, V = b	416	437	1077	43	22-22-131
N = 24, V = a	104	129	501	27	2-26-78
N = 24, V = b	492	517	1277	51	26-26-155
N = 28, V = a	120	149	581	31	2-30-90
N = 28, V = b	568	597	1477	59	30-30-179
N = 32, V = a	136	169	661	35	2-34-102
N = 32, V = b	644	677	1677	67	34-34-203
N = 36, V = a	152	189	741	39	2-38-114
N = 36, V = b	720	757	1877	75	38-38-227
N = 40, V = a	168	209	821	43	2-42-126
N = 40, V = b	796	837	2077	83	42-42-251

Size of the model

Structural properties

ordinary — all arcs have multiplicity one	🗸
simple free choice — all transitions sharing a common input place have no other input place	X (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	X (c)
marked graph — every place has exactly one input transition and exactly one output transition	X(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)	X (f)
source $place(s)$ — one or more places have no input transitions	🖌 (g)
sink place(s) — one or more places have no output transitions	X (h)
source transition(s) — one or more transitions have no input places	🗡 (i)
sink transitions(s) — one or more transitions have no output places \dots	🖌 (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	X(l)
subconservative — for each transition, the number of input arcs equals or exceeds the number of output arcs.	X (m)
nested units — places are structured into hierarchically nested sequential units ⁽ⁿ⁾	· · · · · · · · · · · · · · · · · · ·

- ^(a) stated by CÆSAR.BDD version 3.5 on all 20 instances (10 values of $N \times 2$ values of V).
- ^(b) stated by CÆSAR.BDD version 3.5 on all 20 instances (10 values of $N \times 2$ values of V).
- ^(c) stated by CÆSAR.BDD version 3.5 on all 20 instances (10 values of $N \times 2$ values of V).
- ^(d) stated by CÆSAR.BDD version 3.5 on all 20 instances (10 values of $N \times 2$ values of V).
- (e) stated by CÆSAR.BDD version 3.5 on all 20 instances (10 values of $N \times 2$ values of V).
- $\stackrel{(f)}{=}$ from place 1 one cannot reach place 0.

- ⁽ⁱ⁾ stated by CÆSAR.BDD version 3.5 on all 20 instances (10 values of $N \times 2$ values of V).
- ^(j) transition "exit" is a sink transition.

⁽¹⁾ stated by CÆSAR.BDD version 3.5 on all 20 instances (10 values of $N \times 2$ values of V).

 $^{^{(}g)}$ place 0 is a source place.

^(h) stated by CÆSAR.BDD version 3.5 on all 20 instances (10 values of $N \times 2$ values of V).

^(k) stated by CÆSAR.BDD version 3.5 to be true on 10 instance(s) out of 20, and false on the remaining 10 instance(s).

^(m) stated by CÆSAR.BDD version 3.5 on all 20 instances (10 values of $N \times 2$ values of V).

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

Behavioural properties

safe — in every reachable marking, there is no more than one token on a place	🖌 (o)
dead place(s) — one or more places have no token in any reachable marking	? <mark>(p)</mark>
dead transition(s) — one or more transitions cannot fire from any reachable marking	. ? (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	.? ^(s)
live — for every transition t, from every reachable marking, one can reach a marking in which t can fire	.? ^(t)

Size of the marking graphs

Parameter	Number of reach-	Number of tran-	Max. number of	Max. number of
	able markings	sition firings	tokens per place	tokens per marking
N = 4, V = a	1373 ^(u)	?	1	6
N = 4, V = b	$8.28271e+06^{(v)}$?	1	6
N = 8, V = a	380093 ^(w)	?	1	10
N = 8, V = b	9.14e+11 ^(x)	?	1	10
N = 12, V = a	$9.96004e + 07^{(y)}$?	1	14
N = 12, V = b	$1.00265e+17^{(z)}$?	1	14
N = 16, V = a	$2.56837e+10^{(aa)}$?	1	18
N = 16, V = b	$1.08891e + 22^{(ab)}$?	1	18
N = 20, V = a	$6.5901e+12^{(ac)}$?	1	22
N = 20, V = b	$1.17348e + 27^{(ad)}$?	1	22
N = 24, V = a	1.68829e + 15 (ae)	?	1	26
N = 24, V = b	$\geq 1.22622e + 32^{\text{(af)}}$?	$1^{(ag)}$	26
N = 28, V = a	4.323e+17 ^(ah)	?	1	30
N = 28, V = b	$\geq 3.89477 e + 36$ ^(ai)	?	$1^{(aj)}$	30
N = 32, V = a	$1.10677e + 20^{(ak)}$?	1	34
N = 32, V = b	$\geq 1.90864e + 40^{(al)}$?	1 ^(am)	34
N = 36, V = a	$2.83339e+22^{(an)}$?	1	38
N = 36, V = b	$\geq 1.30486e + 44^{\text{(ao)}}$?	$1^{(ap)}$	38
N = 40, V = a	$7.25353e+24^{(aq)}$?	1	42
N = 40, V = b	$\geq 5.29597 e + 47^{(ar)}$?	$1^{(as)}$	42

^(o) safe by construction – stated by the CÆSAR compiler.

(q) stated by CÆSAR.BDD version 3.5 to be false on 15 instance(s) out of 20, and unknown on the remaining 5 instance(s).

^(r) stated by CÆSAR.BDD version 3.5 to be true on 15 instance(s) out of 20, and unknown on the remaining 5 instance(s).

(s) stated by CÆSAR.BDD version 3.5 to be false on 15 instance(s) out of 20, and unknown on the remaining 5 instance(s).

^(t) stated by CÆSAR.BDD version 3.5 to be false on 15 instance(s) out of 20, and unknown on the remaining 5 instance(s).

- $^{(\mathrm{u})}$ stated by CÆSAR.BDD version 3.5.
- (v) stated by CÆSAR.BDD version 3.5.
- (w) stated by CÆSAR.BDD version 3.5.
- (x) stated by CÆSAR.BDD version 3.5.
- (y) stated by CÆSAR.BDD version 3.5.
- ^(z) stated by CÆSAR.BDD version 3.5. ^(aa) stated by CÆSAR.BDD version 3.5.
- (ab) stated by CÆSAR.BDD version 3.5.
- (ac) stated by CÆSAR.BDD version 3.5.
- (ad) stated by CÆSAR.BDD version 3.5.
- $^{\rm (ae)}$ stated by CÆSAR.BDD version 3.5.
- (af) stated by CÆSAR.BDD version 3.5.
- (ag) stated by the CÆSAR compiler.
- (ah) stated by CÆSAR.BDD version 3.5.
 (ai) stated by CÆSAR.BDD version 3.5.
- ^(aj) stated by the CÆSAR compiler.
- (ak) stated by CÆSAR.BDD version 3.5.
- ^(al) stated by CÆSAR.BDD version 3.5.
- (am) stated by the CÆSAR compiler.

^(p) stated by CÆSAR.BDD version 3.5 to be false on 15 instance(s) out of 20, and unknown on the remaining 5 instance(s).

 $\stackrel{\mathrm{since}}{\mathrm{MCC}}$ 2021

- (ao) stated by CÆSAR.BDD version 3.5.
- ^(ap) stated by the CÆSAR compiler. ^(aq) stated by CÆSAR.BDD version 3.5.
- (ar) stated by CÆSAR.BDD version 3.5.

⁽an) stated by CÆSAR.BDD version 3.5.

^(as) stated by the CÆSAR compiler.