Michele TORTELLI
Dario ROSSI
Emilio LEONARDI

Model-Graft: Accurate, Scalable and Flexible Analysis of Cache Networks

23rd ICNRG Meeting (Interim) – 14/15-01-2016
michele.tortelli@telecom-paristech.fr

MOTIVATION

Simulation is the primary tool for performance evaluation of cache networks.

Some phenomena only appear @ scale…

4-level Binary Tree
Constant C/M = 0.1%, Zipf's α = 1, IRM

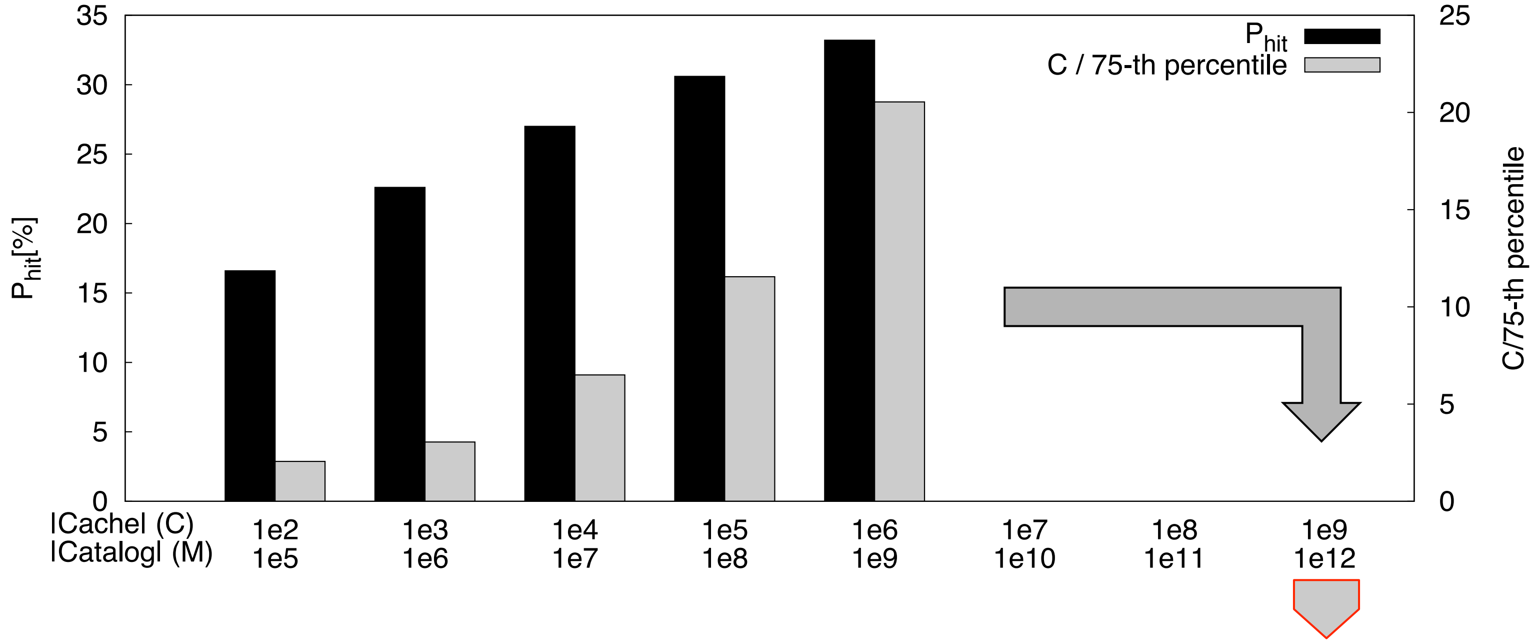

Web catalog size estimate [1] (optimistic)

[1] K. Pentikousis et al., Information-centric networking: Evaluation methodology, Internet Draft - (Oct. 2015)

SIMULATION LIMITATIONS

People have limited CPU & Memory budgets

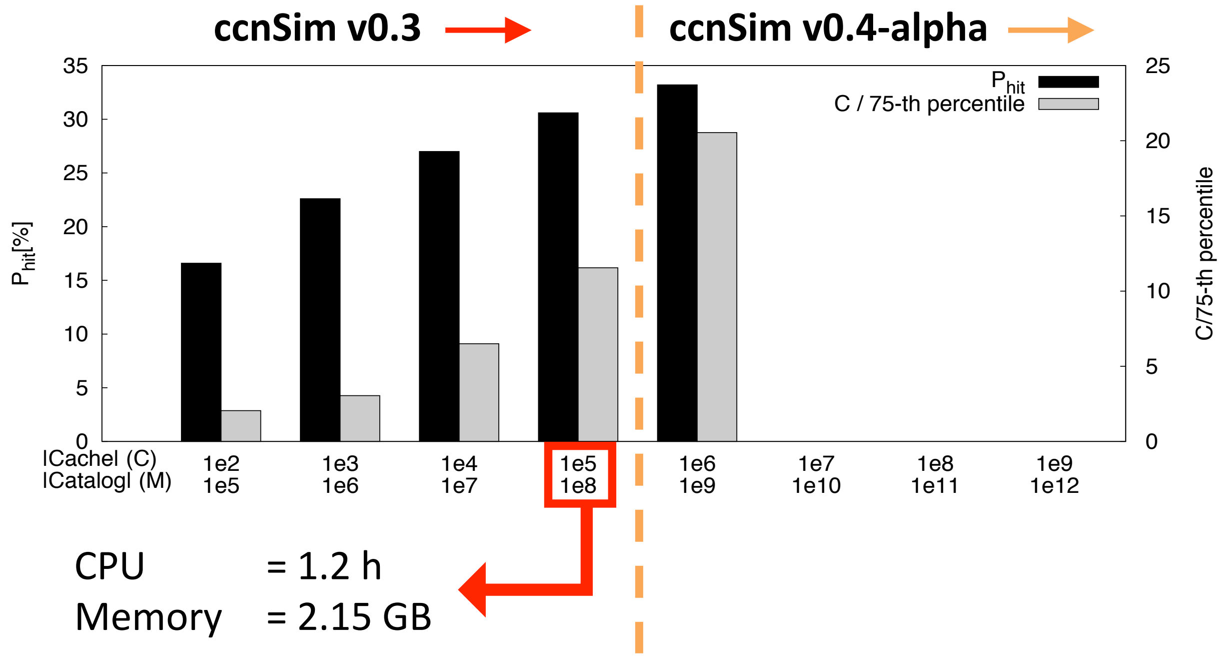

CPU = 1.2 h
Memory = 2.15 GB

SIMULATION LIMITATIONS

People have limited CPU & Memory budgets

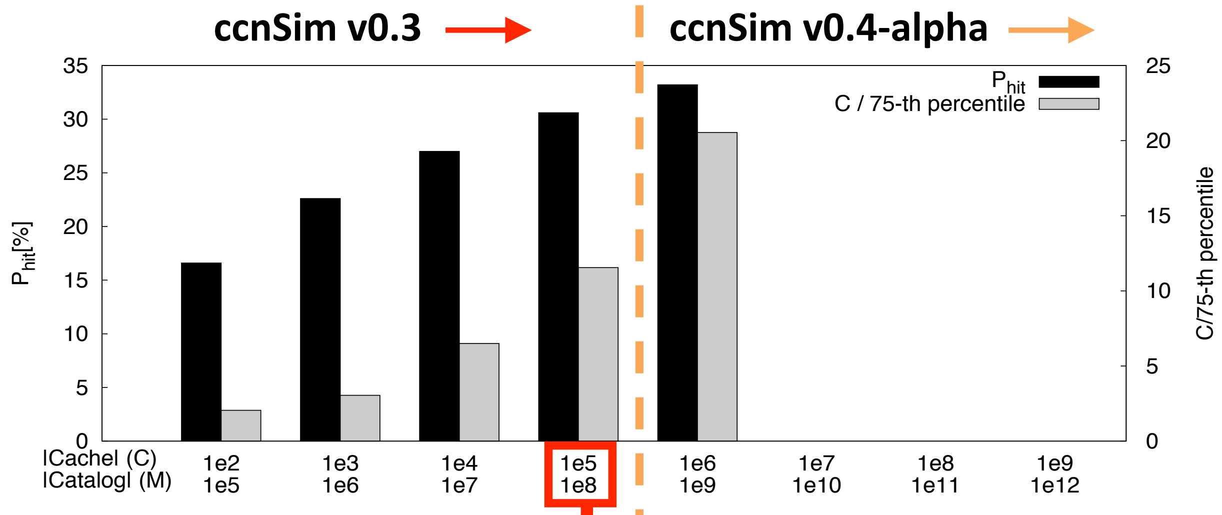

CPU = 1.2 h
Memory = 2.15 GB

CPU = 1.1 h
Memory = 673 MB

Simulation Limitations

People have limited CPU & Memory budgets

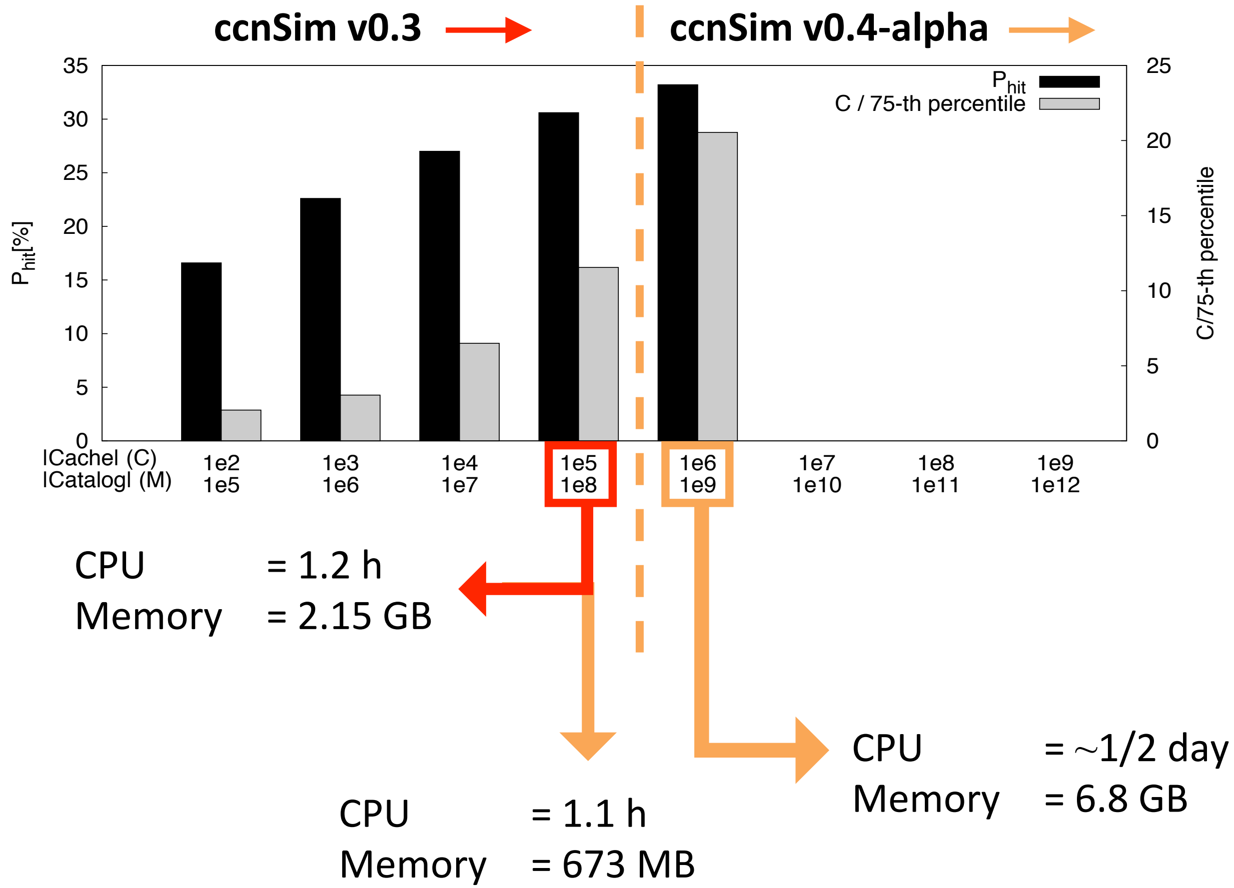

SIMULATION LIMITATIONS

People have limited CPU & Memory budgets

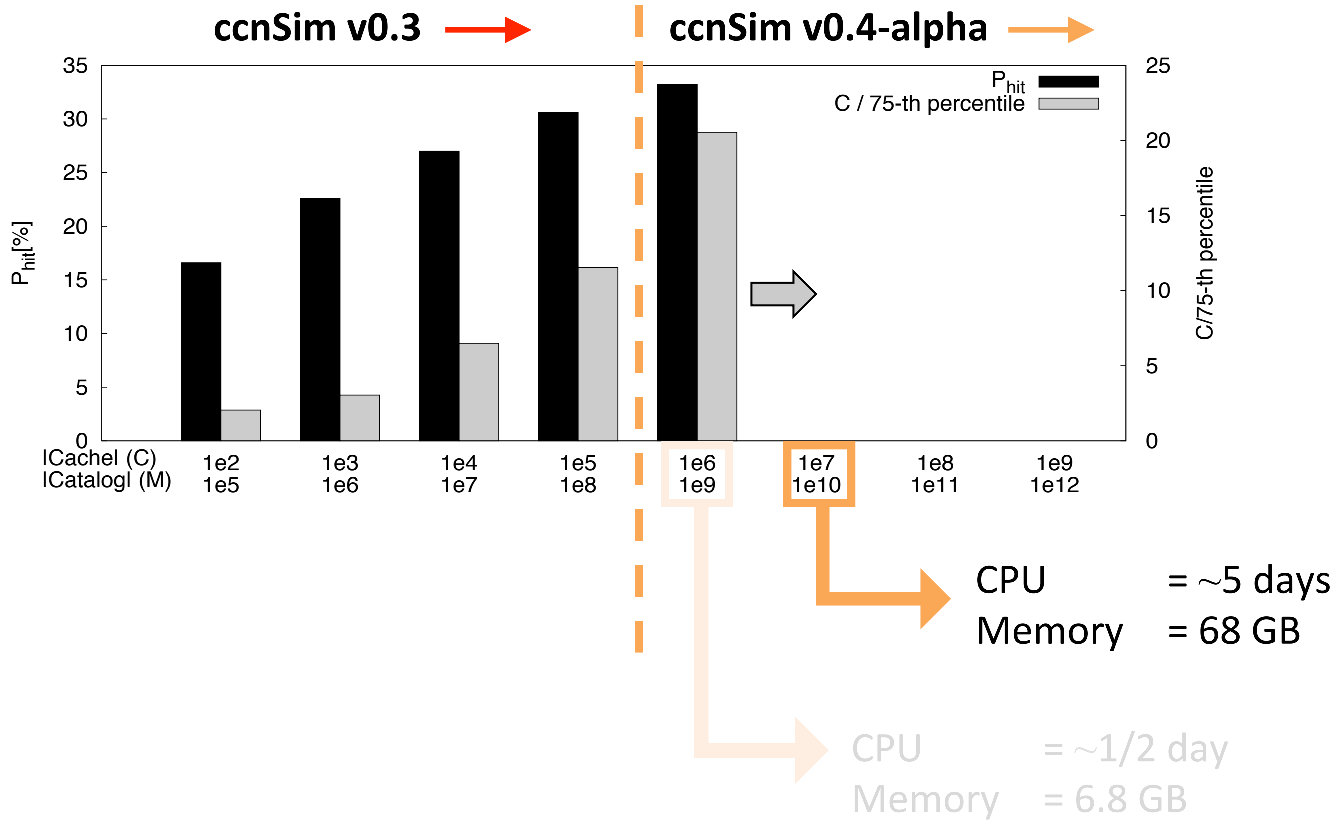

One order of magnitude more becomes **resource expensive**

SIMULATION LIMITATIONS

People have limited CPU & Memory budgets

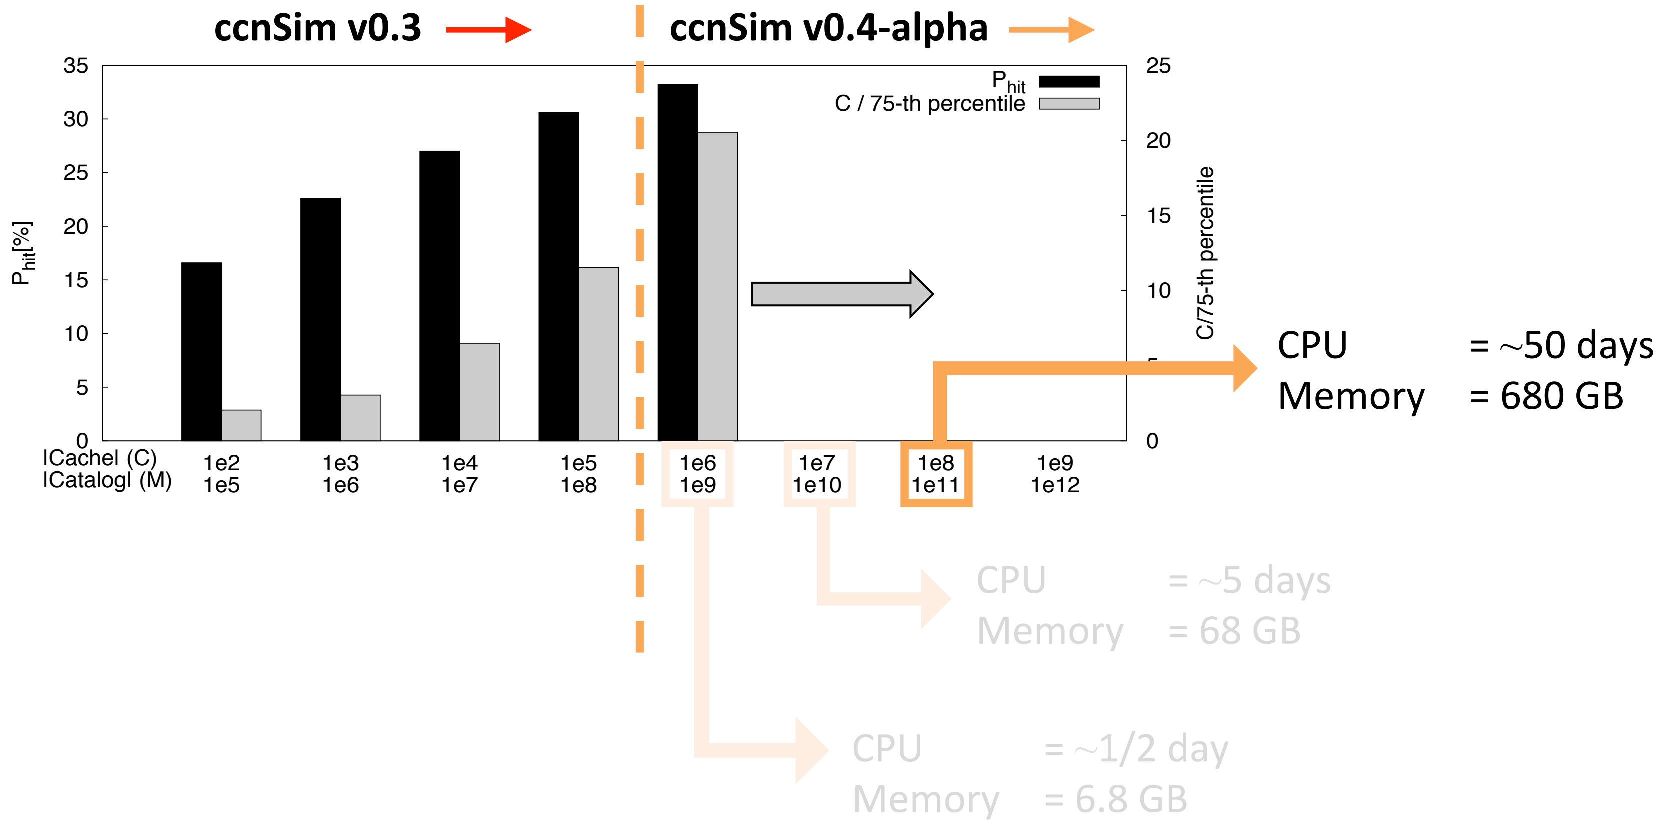

Two orders of magnitude more become **really cumbersome**

SIMULATION LIMITATIONS

People have limited CPU & Memory budgets

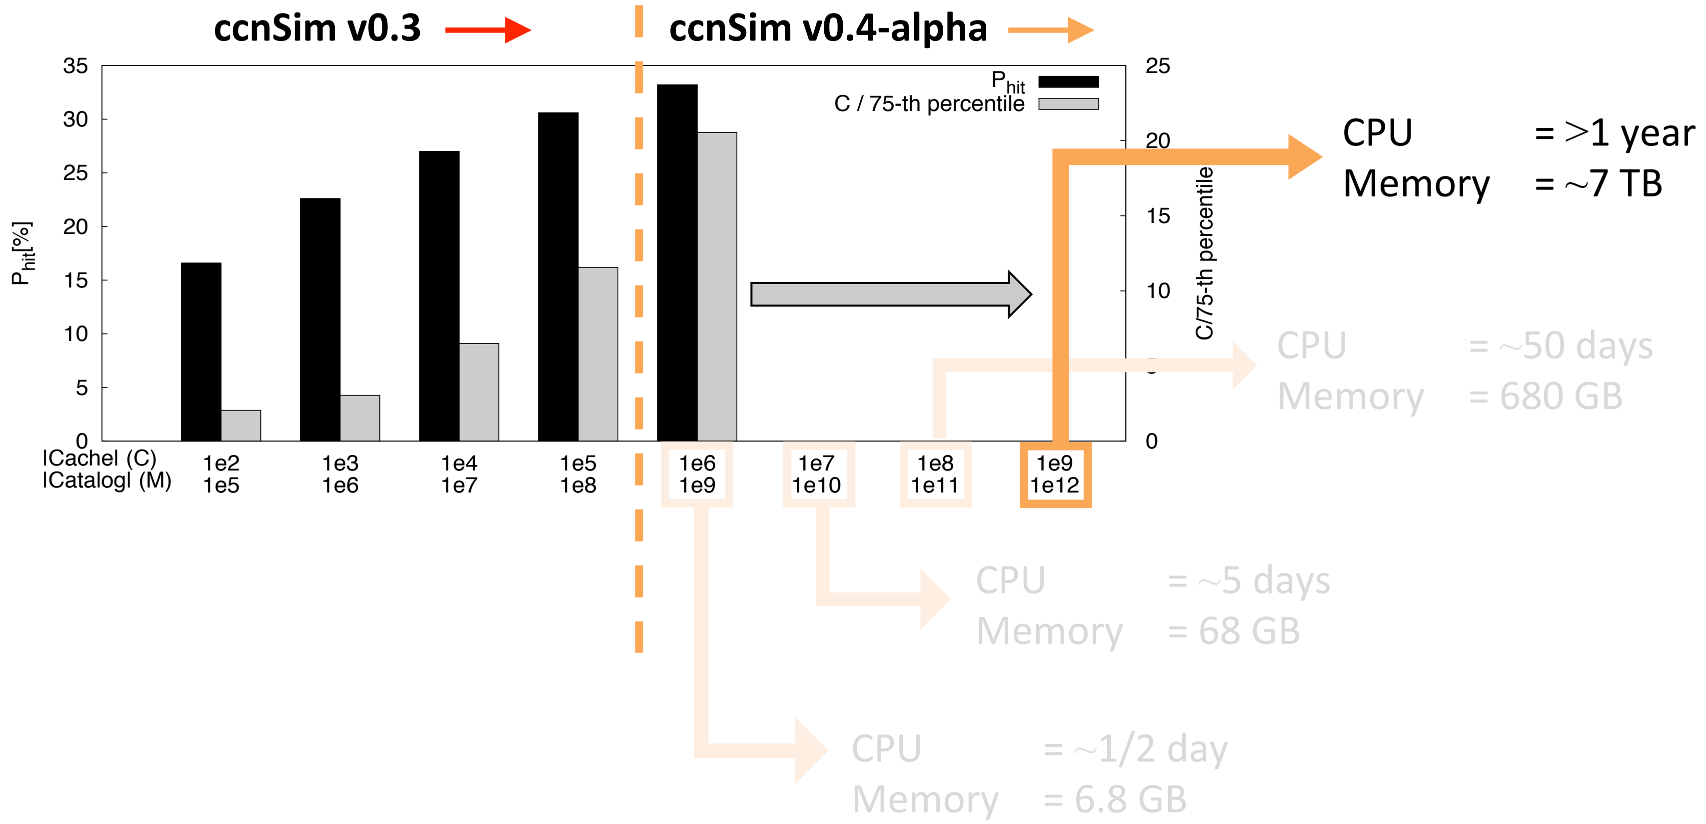

Three orders of magnitude more become **unfeasible**

SIMULATION LIMITATIONS

People have limited CPU & Memory budgets

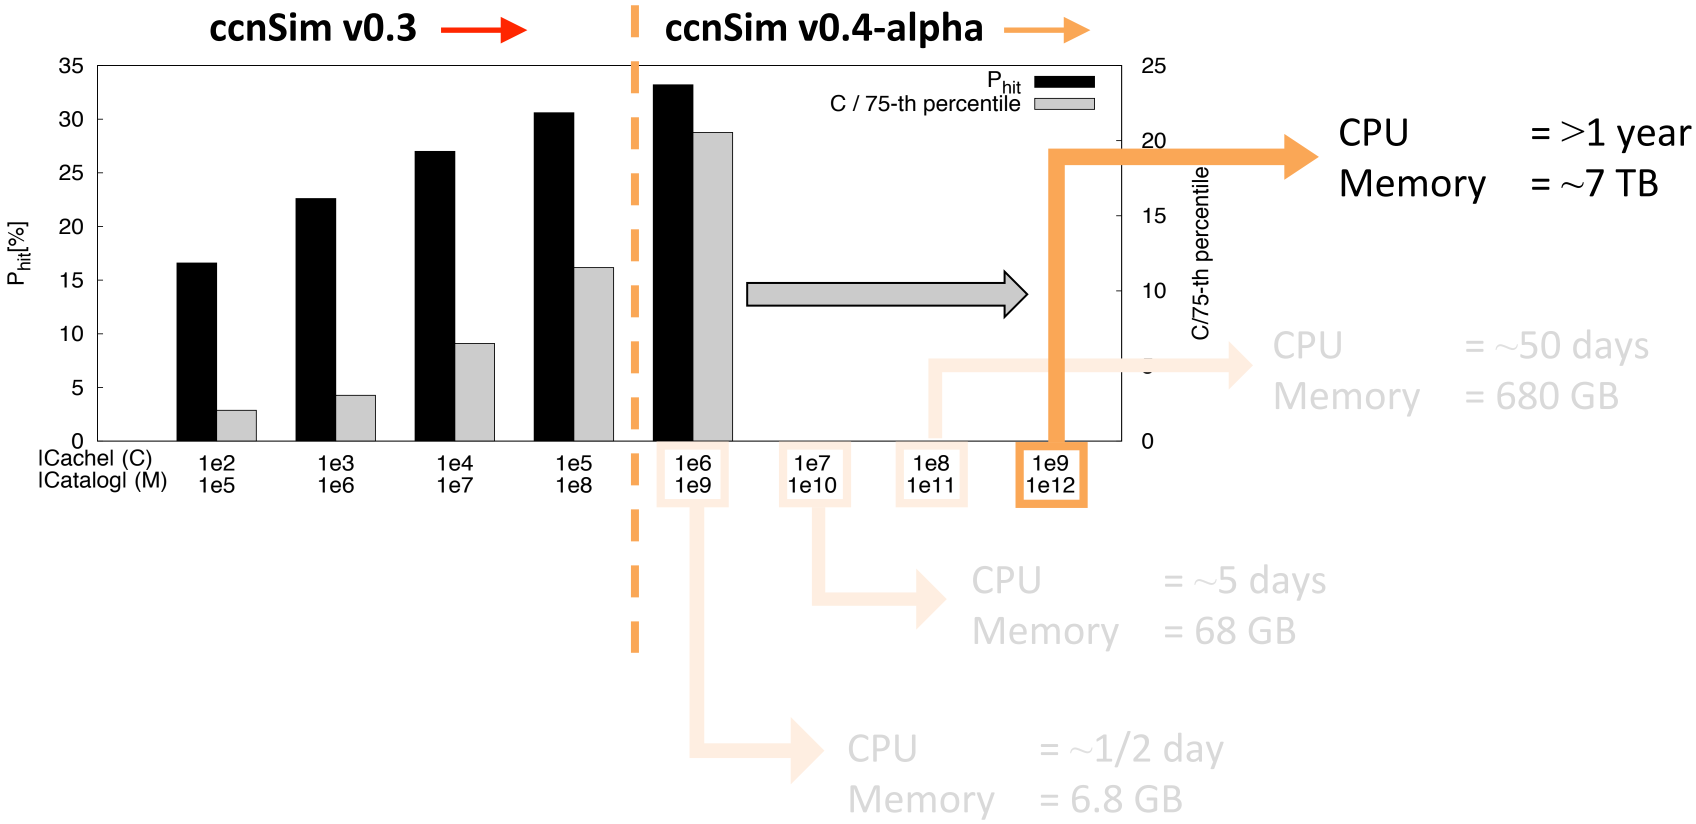

- **Event-driven Simulation** → Require careful instrumentation
  Massive computing power at large scale
  Inefficient (wasting time and memory with expected events)

IDEA

		GAIN
Time to Live (TTL) caches	+	CPU & Memory
Downscaling M, C, and R with factor Δ	+	CPU & Memory
Rejection Inversion Sampling	+	Memory
Error correction with feedback loop	=	Stability & Accuracy

Downscaled MonteCarlo TTL-based (MC-TTL) Simulation

IDEA

		GAIN
Time to Live (TTL) caches	+	CPU & Memory
Downscaling M, C, and R with factor Δ	+	CPU & Memory
Rejection Inversion Sampling	+	Memory
Error correction with feedback loop	=	Stability & Accuracy

Downscaled MonteCarlo TTL-based (MC-TTL) Simulation

IDEA

		GAIN
Time to Live (TTL) caches	+	CPU & Memory
Downscaling M, C, and R with factor Δ	+	CPU & Memory
Rejection Inversion Sampling	+	Memory
Error correction with feedback loop	=	Stability & Accuracy

Downscaled MonteCarlo TTL-based (MC-TTL) Simulation

IDEA

		GAIN
Time to Live (TTL) caches	+	CPU & Memory
Downscaling M, C, and R with factor Δ	+	CPU & Memory
Rejection Inversion Sampling	+	Memory
Error correction with feedback loop	=	Stability & Accuracy

Downscaled MonteCarlo
TTL-based (MC-TTL) Simulation

IDEA

		GAIN
Time to Live (TTL) caches	+	CPU & Memory
Downscaling M, C, and R with factor Δ	+	CPU & Memory
Rejection Inversion Sampling	+	Memory
Error correction with feedback loop	=	Stability & Accuracy

Downscaled MonteCarlo
TTL-based (MC-TTL) Simulation

Ex. with Δ=1e5

~100x CPU & Memory gain
~2% Accuracy

UNDER THE HOOD

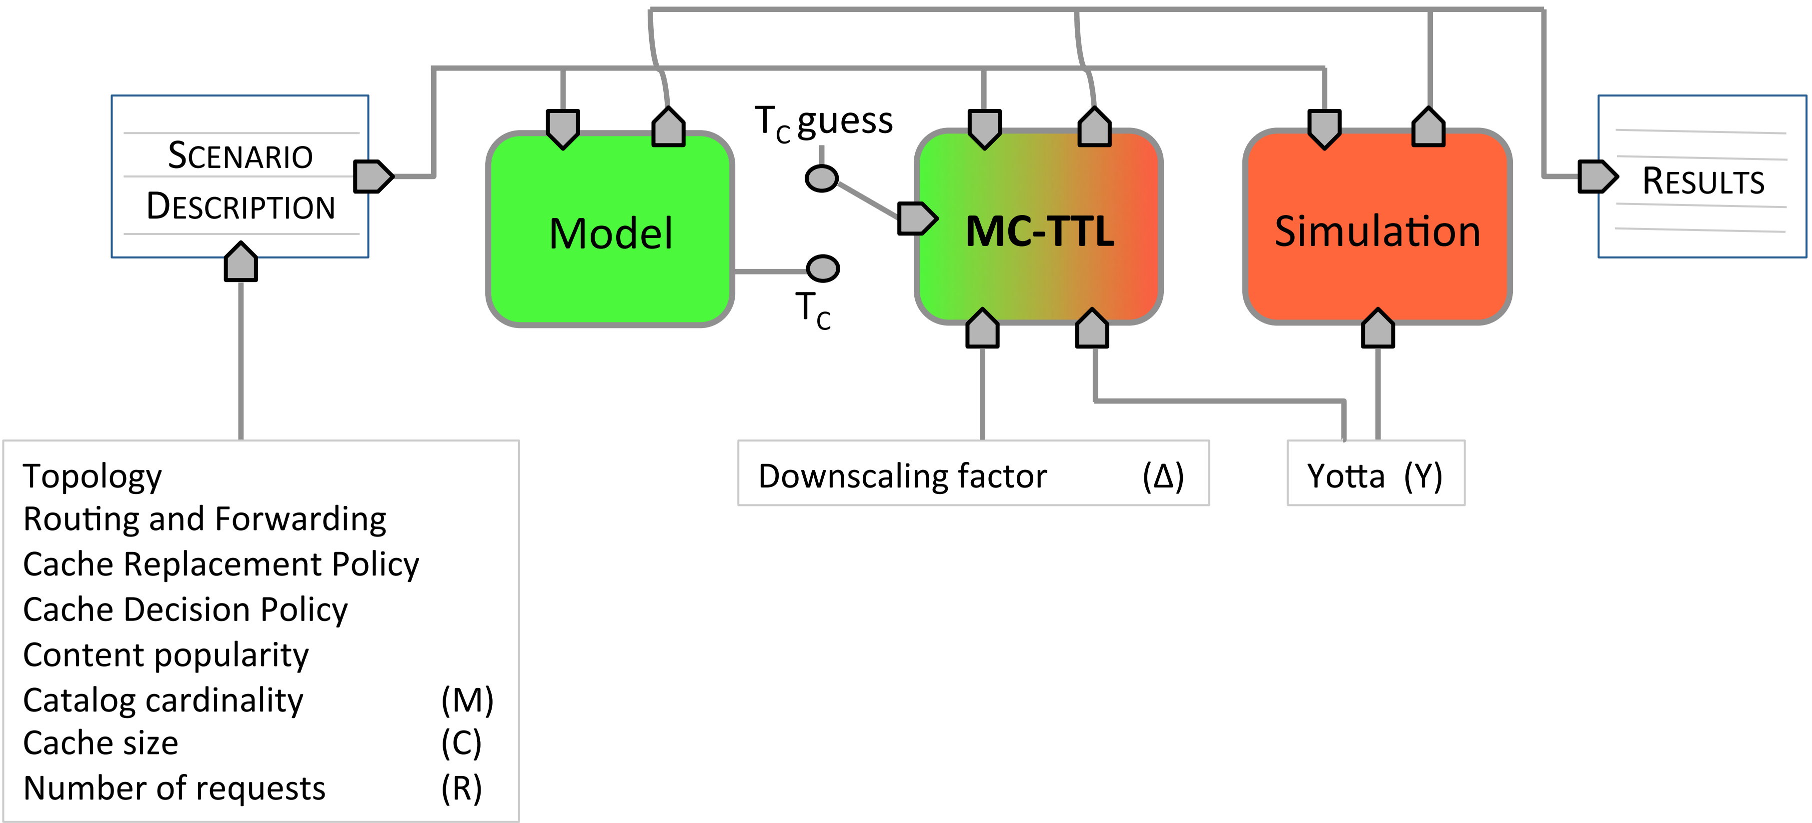

UNDER THE HOOD

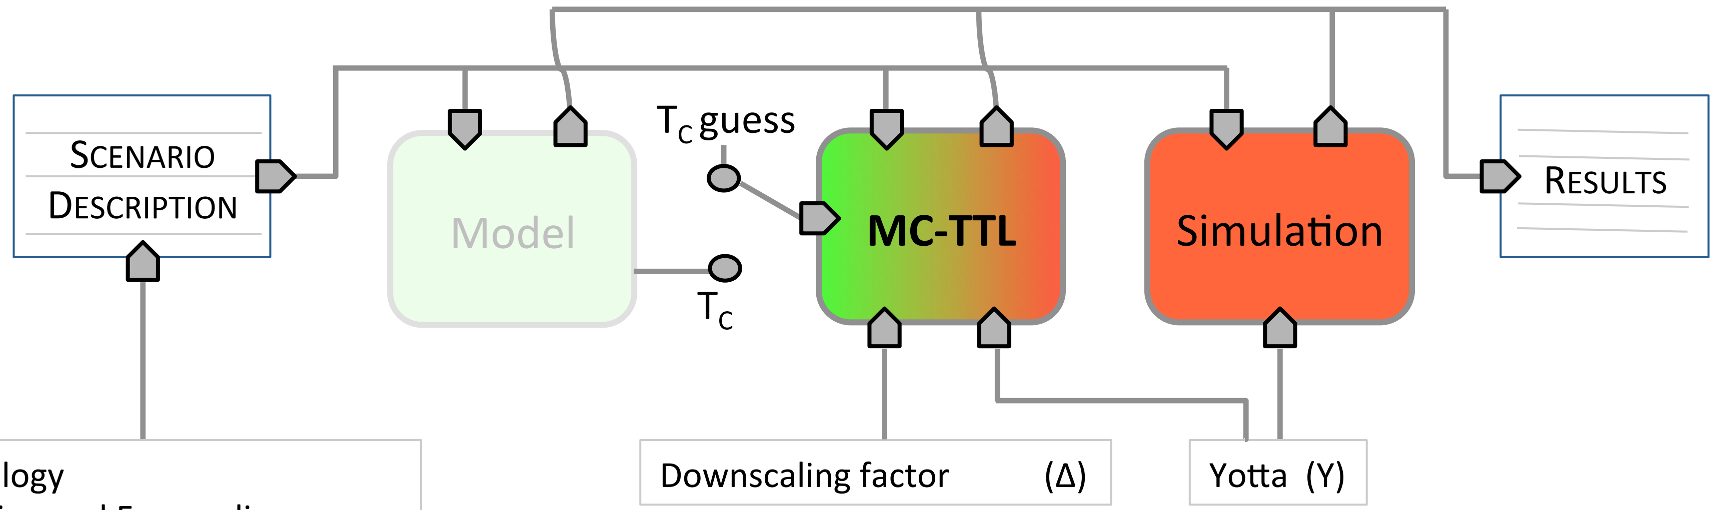

Topology		
Routing and Forwarding		
Cache Replacement Policy		
Cache Decision Policy		
Content popularity		
Catalog cardinality	(M)	
Cache size	(C)	
Number of requests	(R)	

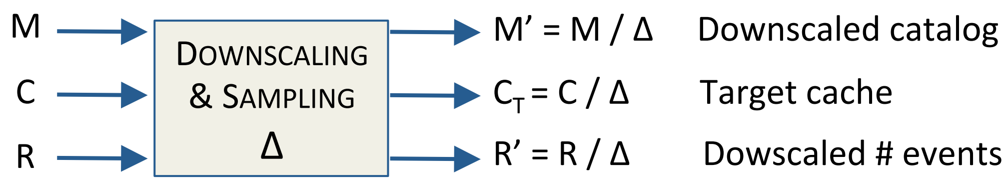

Backup slides for technical details

RESULTS I – VERY LARGE SCENARIO

4-level Binary Tree

|Cache|	-	C = 1e6
|Catalog|	-	M = 1e9
# Requests	-	R = 1e9

Downscaling Factor - Δ = 1e5

Cache Decision Policy		P_{hit}	CPU	Gain	Mem [MB]	Gain
LCE	Simulation	33.2	11.4 h	**160x**	6371	**168x**
	MC-TTL	31.4	256 s		38	
FIX0.1	Simulation	35.4	7.3 h	**90x**	6404	**168x**
	MC-TTL	34.0	291 s		38	
2-LRU	Simulation	37.0	10.8 h	**97x**	8894	**234x**
	MC-TTL	36.1	402 s		38	

RESULTS II – WEB SCALE SCENARIO

Memory Model Fitting (Simulator)

$$\text{Mem} = (1.65 \cdot 10^{-4}) \cdot N \cdot C + 4 \cdot 10^{-6} \cdot M + 19.83 \quad [\text{MB}]$$

(ccnSim v0.4-alpha with rejection inversion sampling)

- 1 cache entry ≈ 165 Bytes
- 1 catalog entry = 4 Bytes
- Fix cost ≈ 19.83 MB

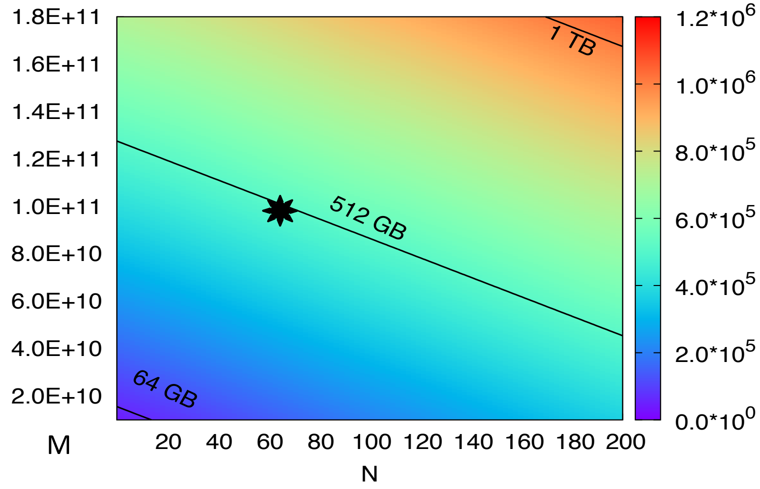

✱ CDN-like: N~60 - M=1e11 - C=1e7

			Technique				
			MC-TTL			Simulation (Est)	
			Mem[MB]	CPU	Cycles	Mem[MB]	CPU
4-level Binary Tree	Parameters	M=R=1e10 - C=1e6 - Δ=1e5	**45**	**0.4 h**	1	**70000**	**4.7 days**
CDN-like (N=67)	Parameters	M=R=1e11 - C=1e7 - Δ=1e6	**31**	**12.5 h**	3	**520000**	**~50 days**

CONCLUSIONS

- Extreme scalability, general methodology
  Implementable in every simulator (ndnSIM, Icarus, …)

- Available in ccnSim v0.4-alpha
  (http://perso.telecom-paristech.fr/~drossi/ccnSim)

- Technical Report (slightly old)
  M. Tortelli, D. Rossi and E. Leonardi,
  Model-Graft: Accurate, Scalable and Flexible Analysis of Cache Networks
  Tech. Rep. [CCN-TR15], Telecom ParisTech, 2015.

THANK YOU !

QUESTIONS ?

BACKUP SLIDES

CHE'S APPROXIMATION

$T_C(m) = T_C$

Cache eviction time for content **m** (i.e., interval of time after which 'm' is evicted)

Assumed to be **CONSTANT**, independent from the content

- $$p_{in}(\lambda_m, T_C) = 1 - e^{-\lambda_m T_C}$$
- $$C = \mathbb{E}[C] = \mathbb{E}\sum_m [1_{\{m \text{ in cache at t}\}}] = \sum_m p_{in}(\lambda_m, T_C)$$
- $$p_{hit} = \sum_m \lambda_m p_{in}(\lambda_m, T_C) / \sum_m \lambda_m$$

UNDER THE HOOD

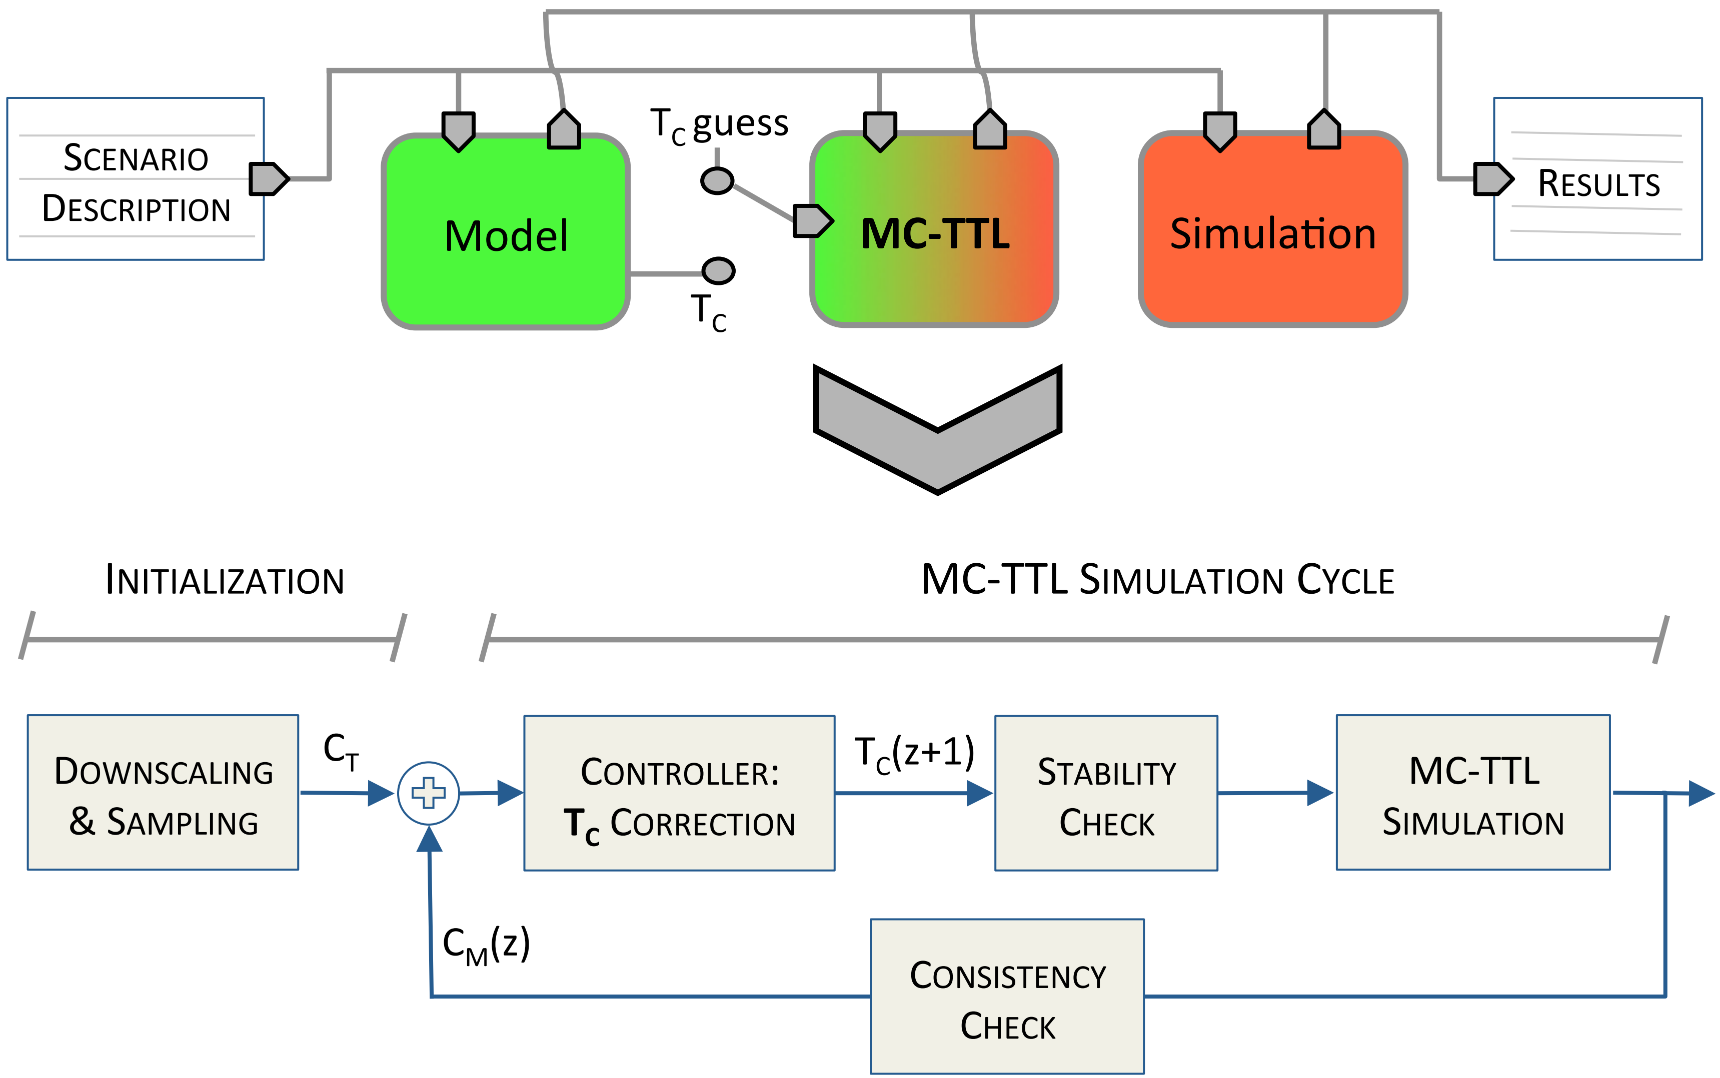

DOWNSCALING & SAMPLING

M → DOWNSCALING & SAMPLING Δ → $M' = M / \Delta$ Downscaled catalog

C → → $C_T = C / \Delta$ Target cache

R → → $R' = R / \Delta$ Dowscaled # events

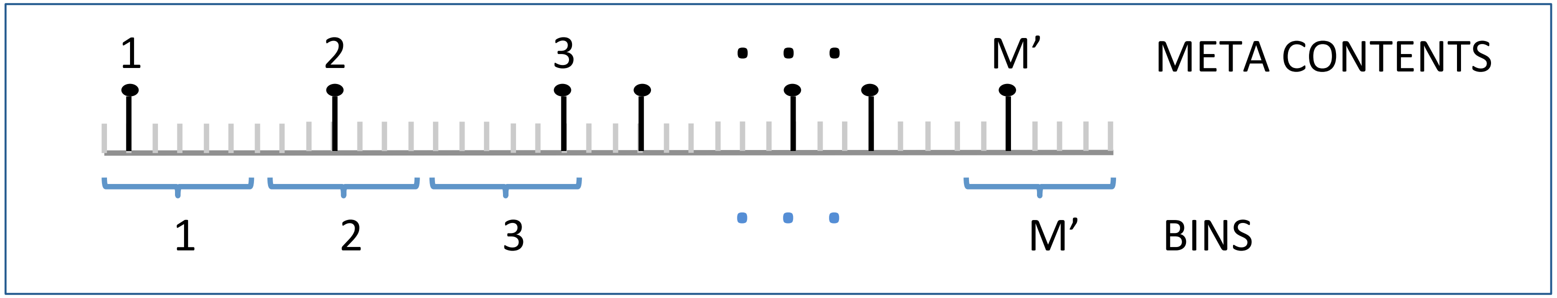

Requested with probability $p \propto \lambda_t / \sum_{j \in \mathcal{M}'_n} \lambda_j$

Meta content n-th bin → $\mathcal{M}'_n$ Δ contents of the n-th bin

$t \in [(n-1)\Delta+1; n\Delta]$

→ NOT SCALABLE!
M' • O(Δ) init cost
M' • O(Δ) space cost
O(logΔ) sampling time

Rejection Inversion Sampling (VERY SCALABLE!)

Extract Zipf's distributed random numbers between [1, Δ]
No Memory and O(1) runtime complexity

STABILITY CHECK

Pivotal role to simulate the R'[1] requests at steady state

Dynamic transient period (routing, meta-caching, topology,…) → **Adaptive Stability Monitor**

- Coefficient of Variation (**CV**) of the mean hit probability (CV = std(p_{hit}) / E(p_{hit}))
- Batch mean of W samples (new sample iif active cache and state change)
- Check stability (i.e., CV < $5 \cdot 10^{-3}$) for the first **N' = Y•N** nodes, where Y ∈]0,1].

(1)

ED-Sim: end = R / (Λ*Cl), with $\Lambda = \sum_{i \in M} \lambda_i$

MC-TTL: end' = R' / (Λ'*Cl), with $\Lambda' = \sum_{j \in M'} \lambda_i \approx \Lambda / \Delta$

Since end = end' → R' = R / Δ

T_C CORRECTION & CONSISTENCY CHECK

Hp: each TTL cache will store, on average, $\mathbf{C_T = C/\Delta}$ contents at steady state if its *eviction time* corresponds to the *characteristic time* T_C of its equivalent LRU non-scaled cache.

Controlled Variable = Measured Cache Size $\mathbf{C_M}$

$$C_{M\,i}(k+1) = \frac{C_{M\,i}(k)\ t(k) + B_i(k+1)\ [t(k+1) - t(k)]}{t(k+1)}$$

$C_{Mi}(k+1)$ = online avg of the cache size of the i-th node @ k-th measurement time

$B_i(k+1)$ = actual # contents stored inside the TTL cache

Samples are taken every miss event with probability p=0.1

Consistency Check (after R' events)

$$\frac{1}{N'} \sum_{i \in N'} \frac{|C_T - C_{M\,i}|}{C_T} < 0.1$$

Same N' nodes
(coherent with stability check)

YES → End MC-TTL simulation

NO → Correct Tc and start new MC-TTL cycle

$$T_{C\,i}(z+1) = T_{C\,i}(z)\ \left(\frac{C_T}{C_{M\,i}(z)}\right)$$

C_M distance from C_T connected to input T_C distance from real T_C (the higher T_C the bigger C_M)

T_C SENSITIVITY - I

4-level Binary Tree

|Cache| - C = 1e3
|Catalog| - M = 1e6 Downscaling Factor - Δ = 1e2
Requests - R = 1e7

	Tc Values [s]			
Level	**LCE**	**FIX0.1**	**LCD**	**2-LRU (Name/Main)**
0 (Root)	11	115	13	14 / 654
1	22	218	1090	27 / 1040
2	43	400	1250	51 / 1420
3 (Leaves)	75	570	815	75 / 1255
P_hit @ Stab.	22.8	25.9	28.5	29.2
P_hit @ End.	22.8	26.3	28.1	29.4
Stab. Time	8.4	10.3	14.4	23.7
End Time	482.7	390	386	456.3

T_C SENSITIVITY - II

4-level Binary Tree

|Cache|	-	C = 1e5			
|Catalog|	-	M = 1e8	Downscaling Factor	-	Δ = 1e4
# Requests	-	R = 1e8			

LCE

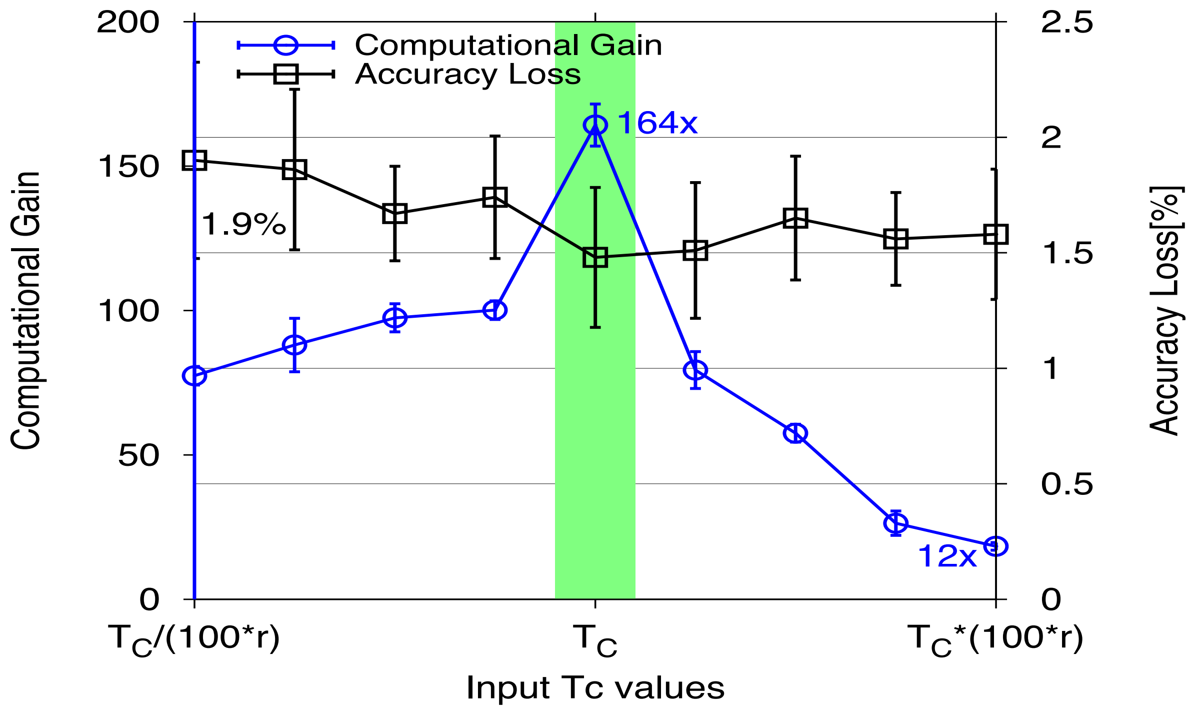

Simulation: Phit = 30.6% - CPU = 4089 s - Mem = 673 MB
Model: Phit = 29.8% - CPU = 740 s - Mem = 24240 MB
MC-TTL: - Mem ≈ 38 MB (18x)

YOTTA SENSITIVITY

NDN Testbed

|Cache| - C = 1e3
|Catalog| - M = 1e6
Requests - R = 1e7

Downscaling Factor - Δ = 1e3

Sim. Type	Yotta	P_hit @ Stab.	P_hit @ End	Stab. Time [s]	End Time [s]
ED-Sim	1	33.8	33.9	99.7	458.0
	0.95	33.7	33.9	42.4	432.5
	0.9	33.7	33.9	35.0	421.9
	0.75	33.6	33.9	29.3	408.13
	0.5	32.8	33.9	15.4	401.33