Performance of Discrete Event Simulation on a Multiprocessor using Optimistic and Conservative Synchronization

Bruno R. Preiss

Department of Electrical and Computer Engineering
University of Waterloo
Waterloo, Ontario, Canada
N2L 3G1

Yaddes Project

Overview

- empirical performance study
- distributed/parallel DES
- Transputer multiprocessor
- speedup w.r.t. sequential DES
- compart an optimistic and 2 conservative synchronization methods
- closed queueing network benchmarks
- vary topology, lookahead, load

Yaddes System Overview

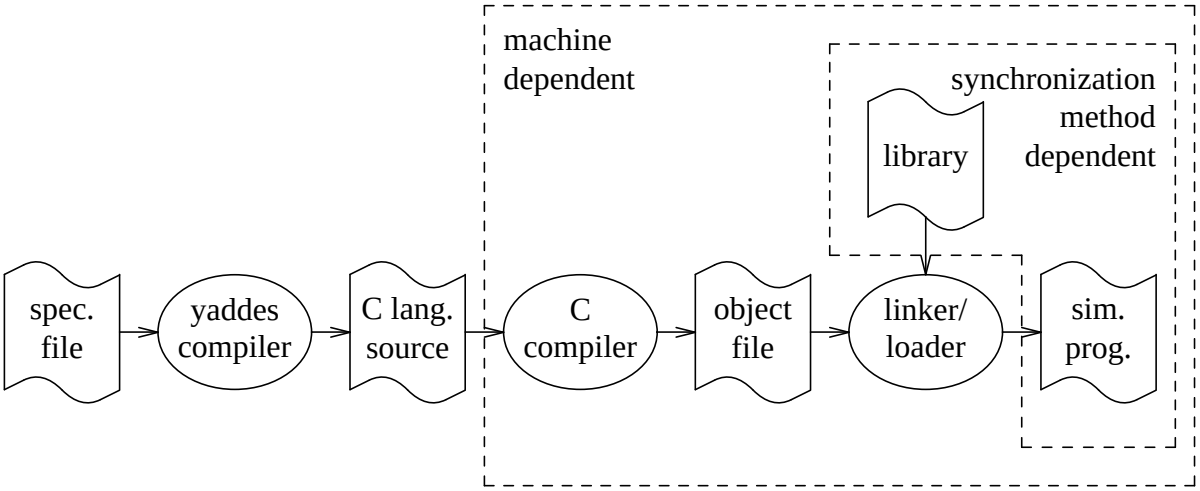

Yaddes Logical Process

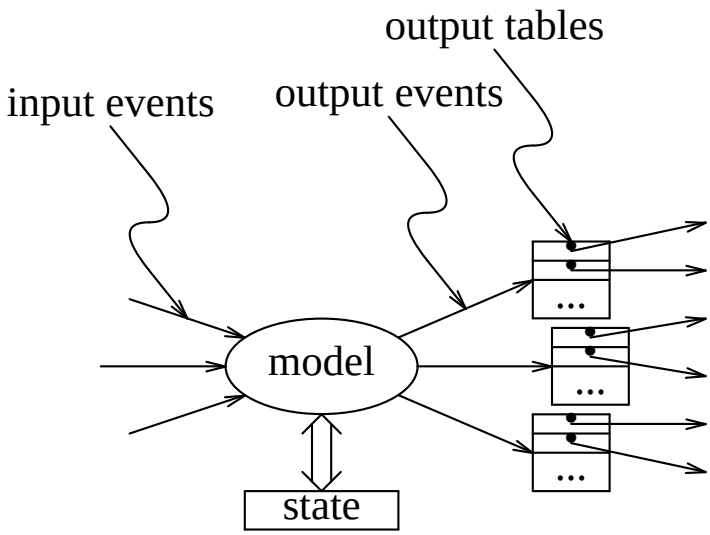

- event combinations
- Yaddes promise: ‘‘If the events on each output arc have monotonically incresing timestamps, then input event combinations will have monotonically increasing timestamps.’’

Synchronization Methods

name	description
EL	Sequential discrete event simulation (event-list-driven).
ML	Distributed discrete event simulation using multiple event lists.
CM	Conservative distributed discrete event simulation (Chandy-Misra-Bryant).
VT	Optimistic distributed discrete event simulation (virtual-time-based).

Queueing Network Benchmarks

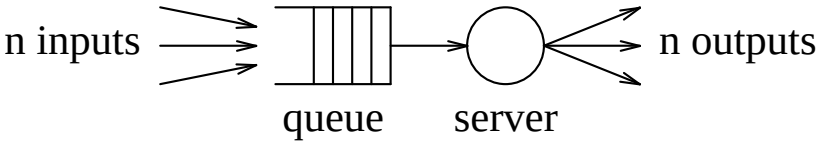

- biased, exponentially distributed service times
- uniformly distributed random routing
- FCFS queueing
- indistinquishable customers

Queueing Network Topologies

network topology	number of LPs	number of inputs per LP	number of outputs per LP	size of smallest loop
Ring	64	2	2	2
Mesh	64	4	4	2
Multistage	64	2	2	8
Hypercube	64	6	6	2

Server Models

- Lookahead: difference between timestamp on an output event and the local time of the LP
- Epsilon lookahead
- Service-time lookahead
- System-time lookahead
- System-time lookahead + Presampling lookahead
  - only affects null messages
  - uses futurelist

Observations

Speedup vs. Processors

- ML rarely does better than 2
- CM can be near linear (high load, large lookahead)
- CM often exhibits slowdown
- VT more uniformly exhibits speedup

Speedup vs. Messages

- best speedup occurs when fewest messages are sent
- in VT, increasing lookahead decreases messages
- increasing lookahead does not affect VT speedup as much as increasing load

Speedup vs. Number of Processors, ML.

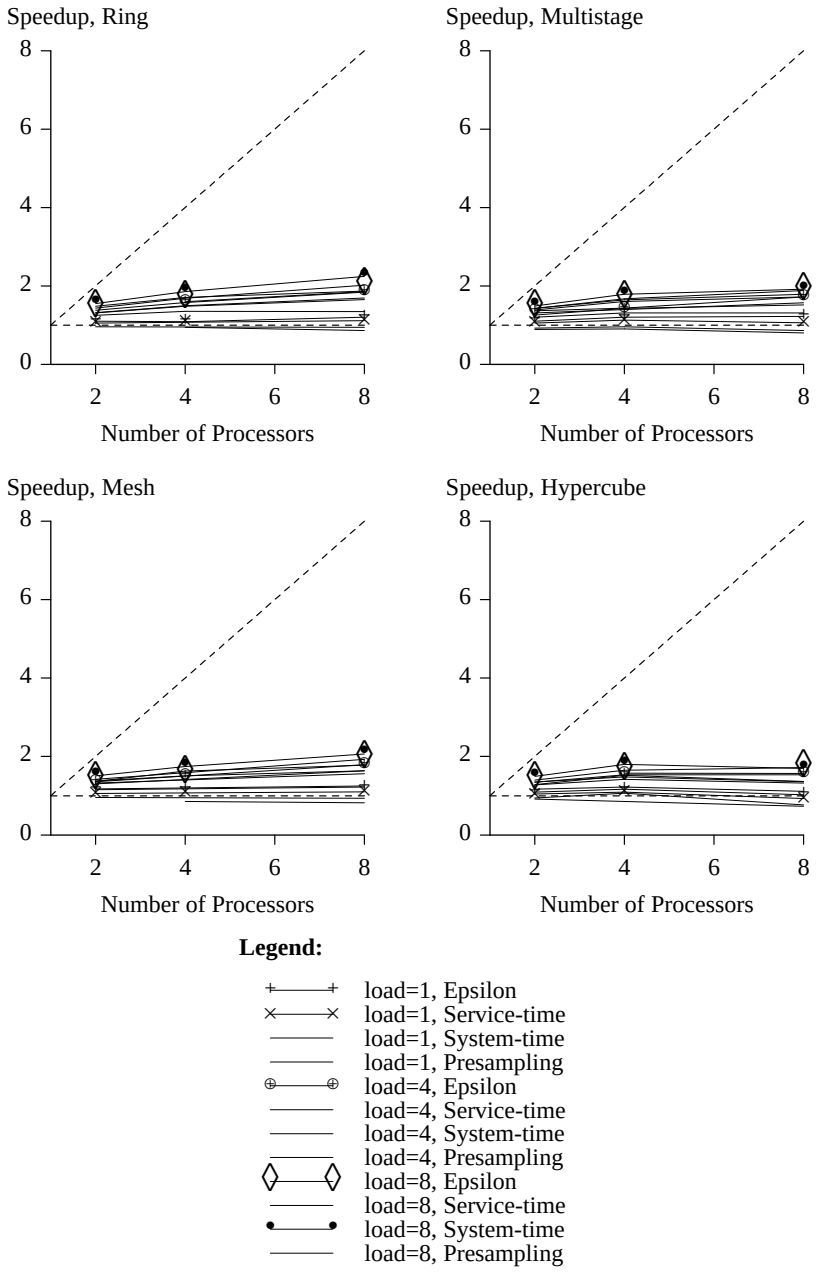

Speedup vs. Number of Processors, CM.

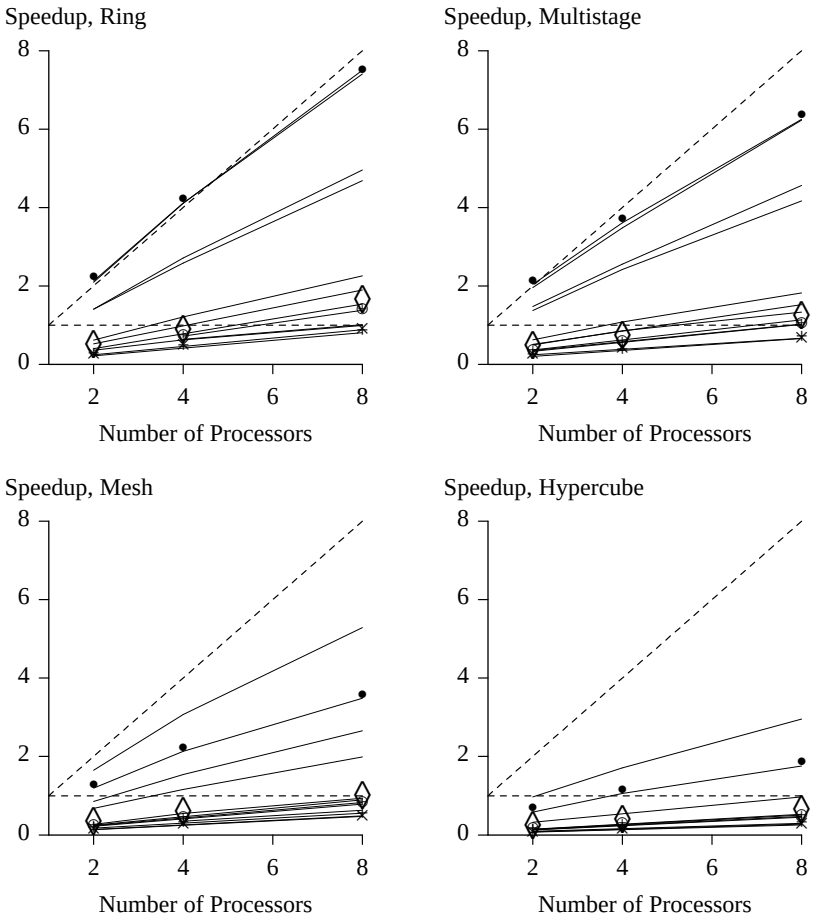

Speedup vs. Number of Processors, VT.

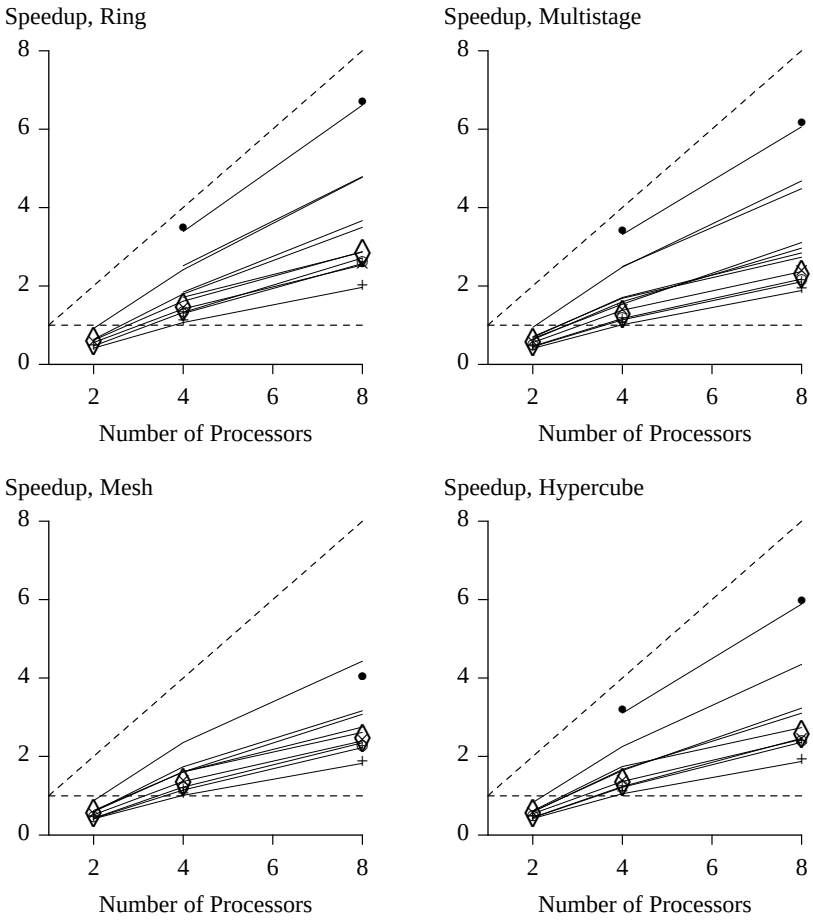

Speedup vs. Number of Messages, CM.

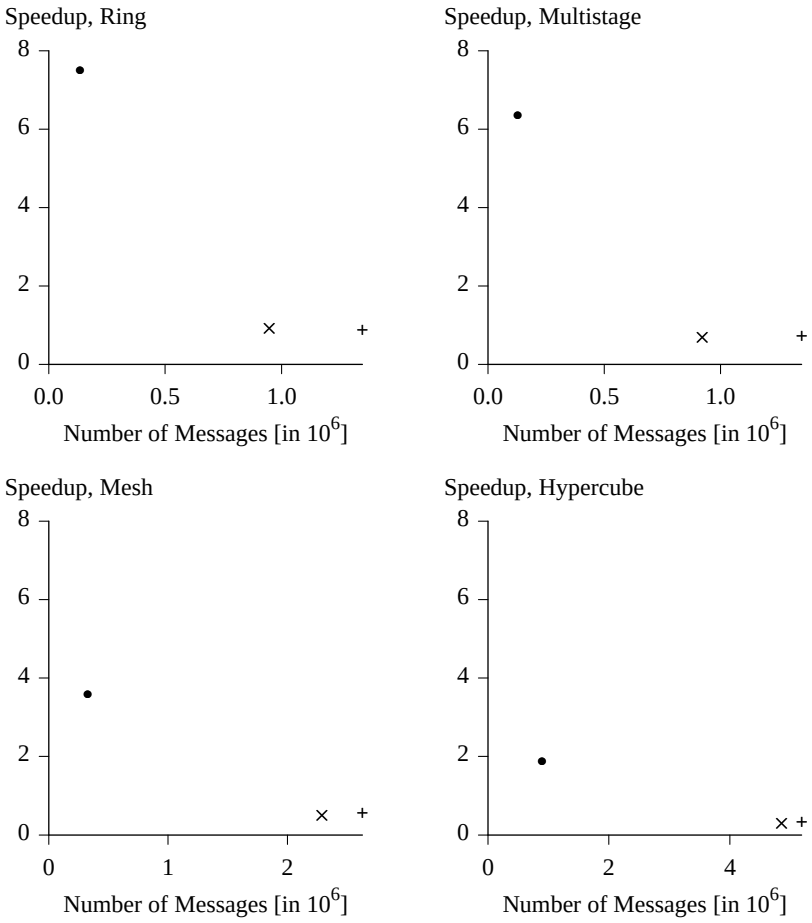

Legend:

+	+	load=1, Epsilon
×	×	load=1, Service-time
		load=1, System-time
		load=1, Presampling
⊕	⊕	load=4, Epsilon
		load=4, Service-time
		load=4, System-time
		load=4, Presampling
◊	◊	load=8, Epsilon
		load=8, Service-time
•	•	load=8, System-time
		load=8, Presampling

Speedup vs. Number of Messages, VT.

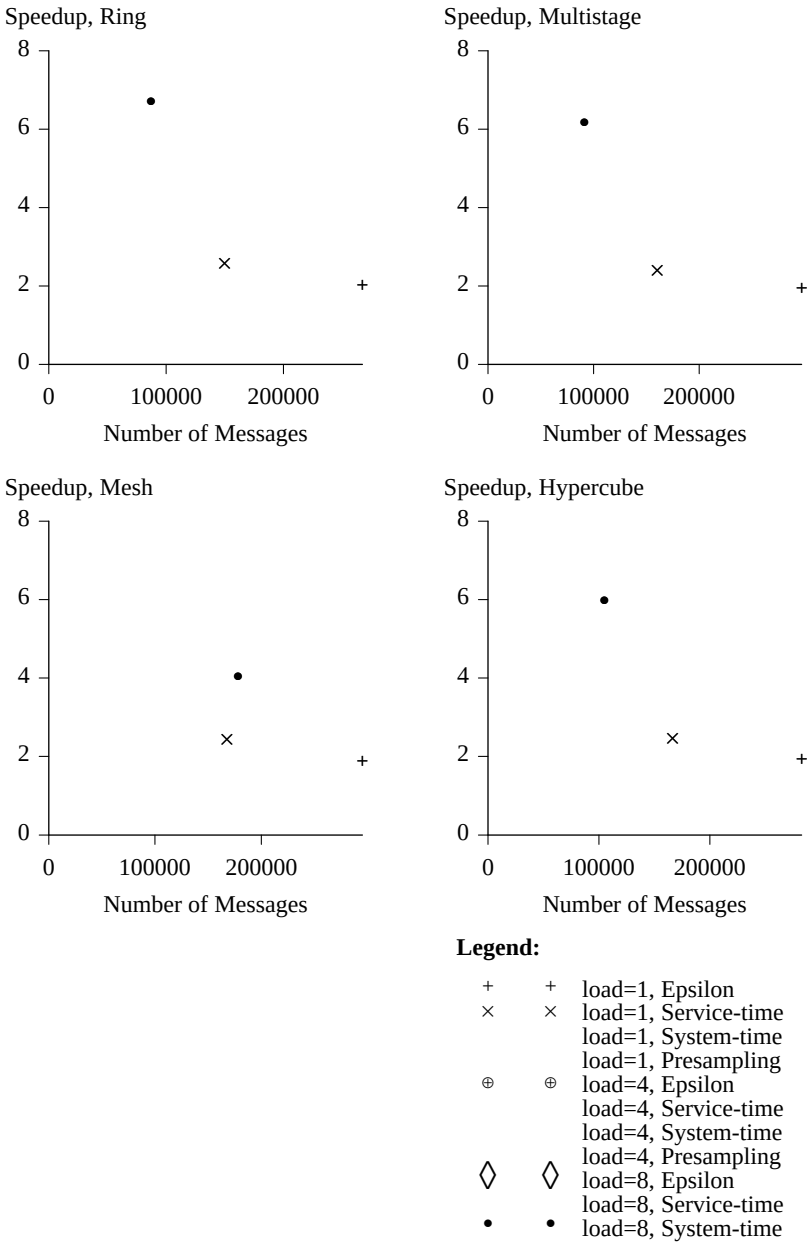

Speedup vs. Number of Model Calls, CM.

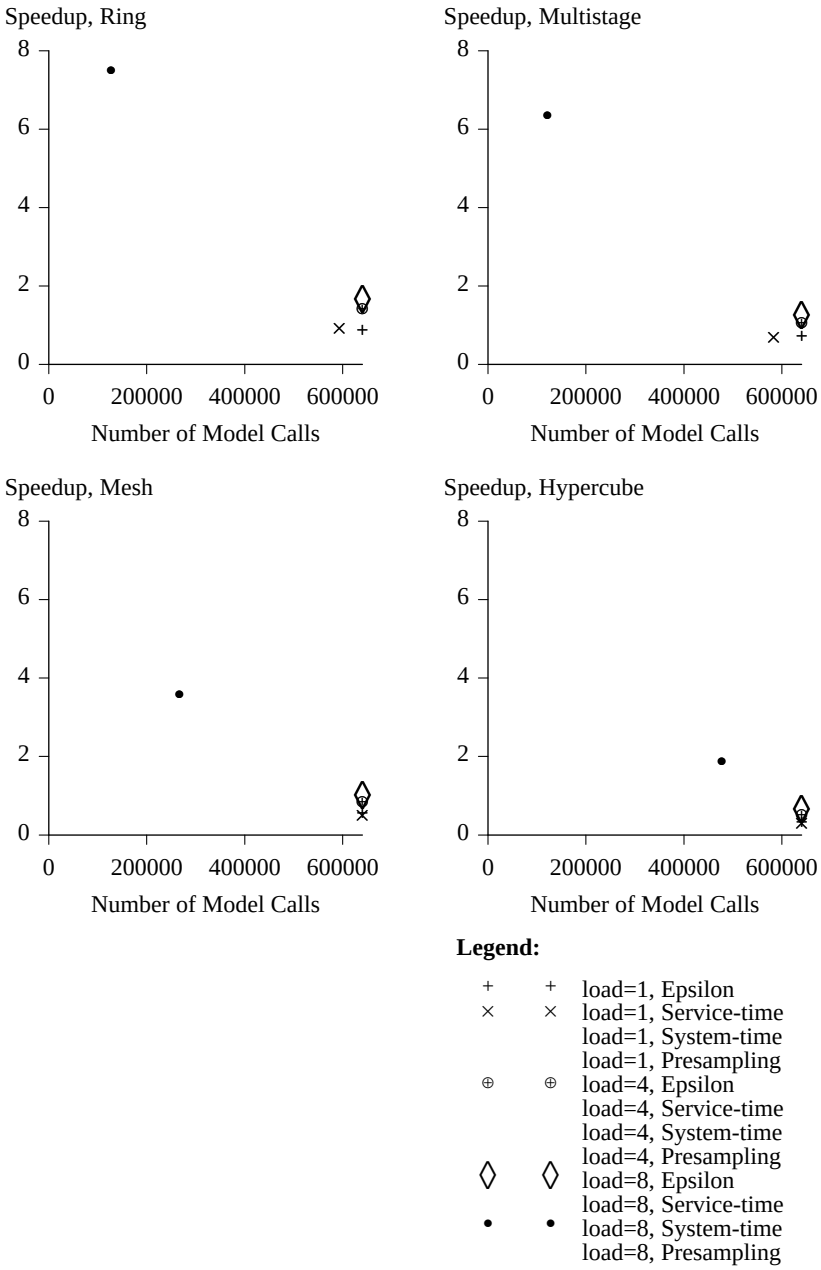

Speedup vs. Number of Model Calls, CM.

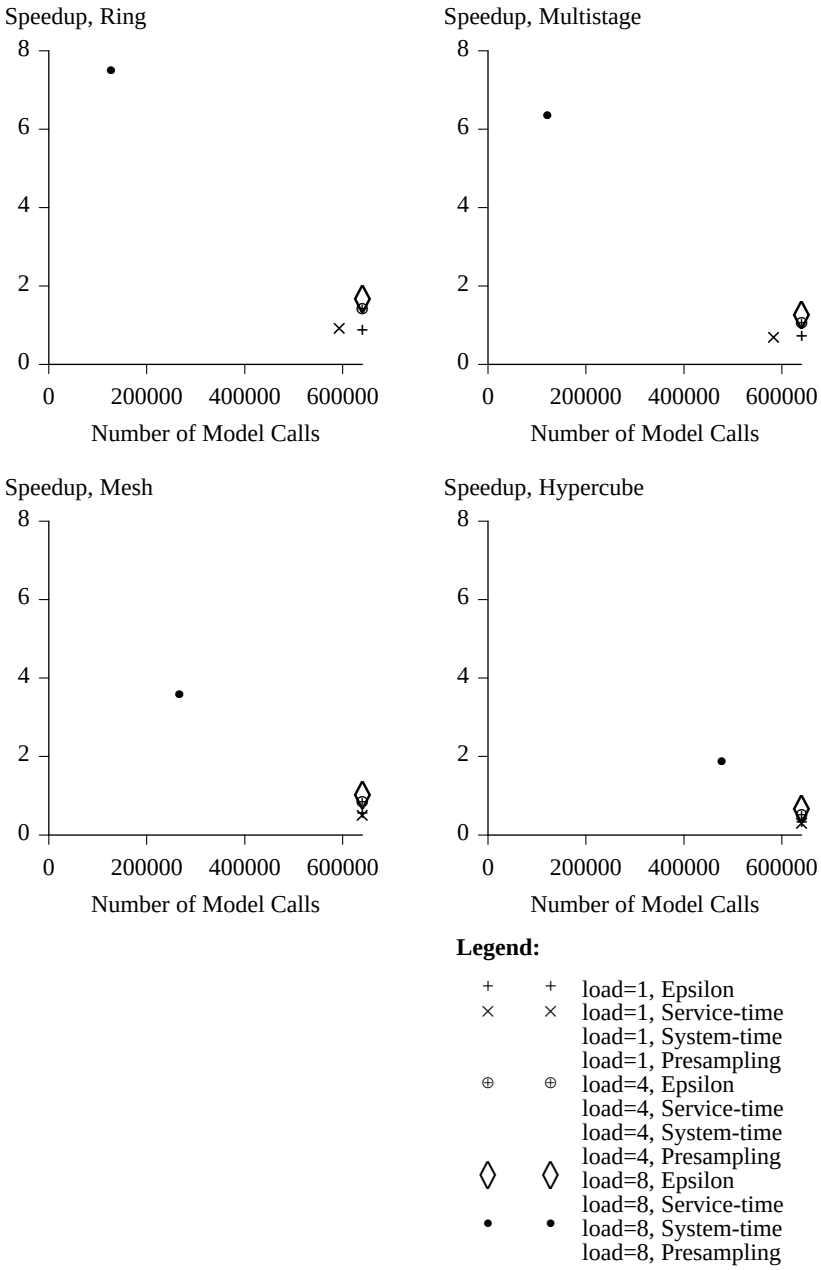

Speedup vs. Number of Model Calls, VT.

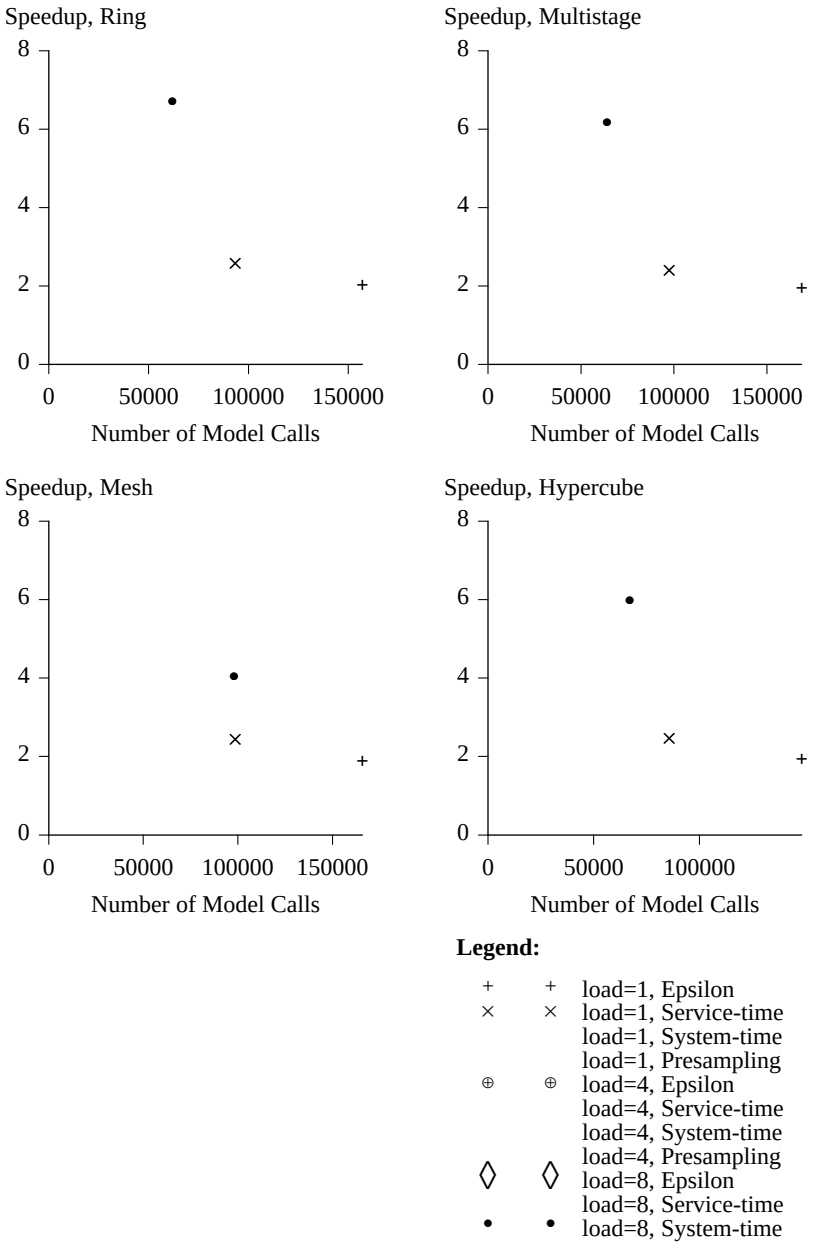

Summary

- empirical study of performance of parallel DES
- closed queueing network benchmarks
- ML is a dead loss
- CM & VT achieve comparable speedup but behave differently w.r.t. topology, lookahead, and load
- best CM: low fanout, large lookahead
- best VT: almost all cases for $n \geq 4$