Trade-off Analysis of Simulation Complexity and Simulation Accuracy using SST/Macro

Zhou Tong
Florida State University
tong@cs.fsu.edu

Scott Pakin and Mike Lang
Los Alamos National Lab
{mlang, pakin}@lanl.org

Xin Yuan
Florida State University
xyuan@cs.fsu.edu

Introduction

Modeling and simulating of HPC applications on contemporary supercomputing platforms is challenging due to the size of the increasingly growing system and application size.

Motivation

Simulation: flow-level vs. packet-level vs. Flit-level

1. Increasing level of simulation complexity yields higher simulation accuracy
2. How much more accurate simulation results finer-grained simulations can deliver and at the cost of higher simulation complexity?

Method

- Investigate the trade-off between simulation complexity and simulation accuracy
- Compare measurement, simulation and modeled results
  - The physical measurement of application execution on three HPC systems.
  - The macroscale component of the Structured Simulation Toolkit (SST)
  - Fast classification Tool

Fast Classification Tool (IPDPS'16)

- Trace-driven performance modeling tool
- One simulation predicts application performance for many network configurations representing various interconnecting network options

SST/Macro

- Simulator for large-scale parallel computer architecture
- Simulate full-scale machine and application
- Trace-driven simulation mode provides performance prediction on a target machine

Table I
APPLICATION TRACES COLLECTED FROM HOST MACHINES WITH THE NUMBER OF MPI RANKS AND THE PERCENTAGE OF COMMUNICATION OBSERVED.

	Machine	Number of Ranks	% Comm.
AMG	Edison	8,27,216,1728	4,21,36,53
AMR	Hopper	64, 1728	16,29
BigFFT	Edison	9, 100, 1024, 10k	1,2,8,4
CR	Hopper	10, 100, 1000	14,45,69
FB	Hopper	125, 1000, 10648	19,22,18
MG	Hopper	125, 1000, 10648	4,4,5
MiniFE	Edison	18, 144, 1152	16,6,7
PARTISN	Mustang	168	54
SNAP	Mustang	168	67

DUMPI traces are available on DOE's Design Forward Web Site
https://portal.nersc.gov/project/CAL/designforward.htm

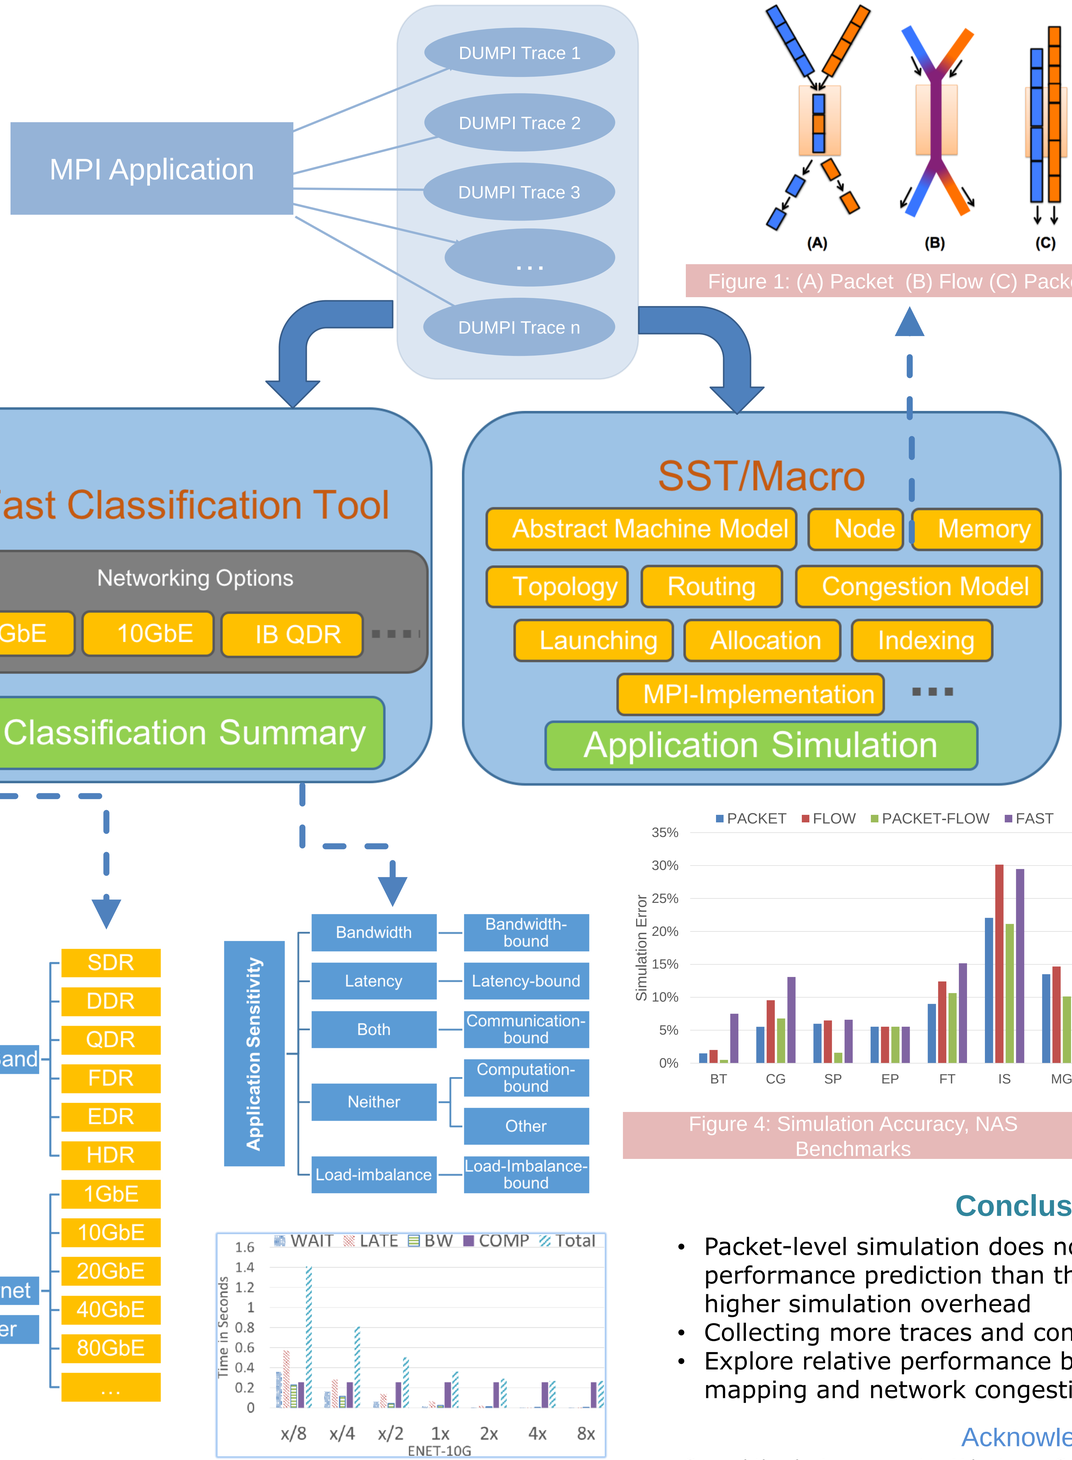

Figure 1: (A) Packet (B) Flow (C) Packet-flow

Preliminary Results

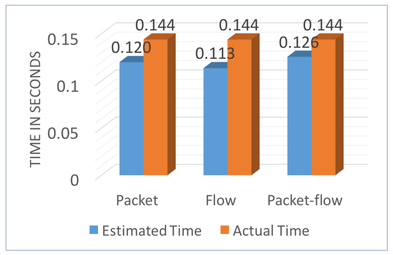

Figure 2. Packet vs. Flow vs. Packet-Flow Trade-off Analysis (CR-Hopper-10)

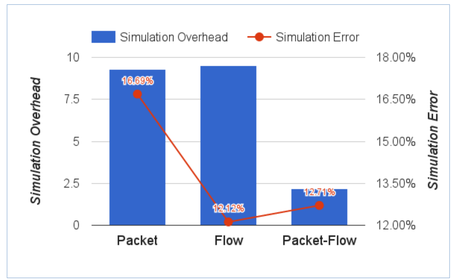

Figure 3: Simulation Overhead vs. Simulation Accuracy (CR-Hopper-10)

Figure 4: Simulation Accuracy, NAS Benchmarks

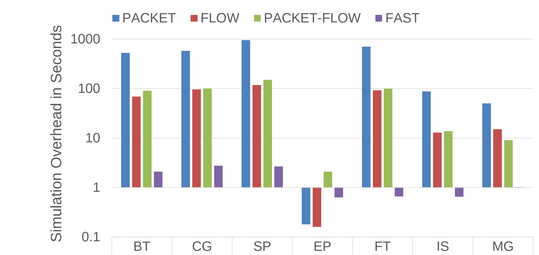

	BT	CG	SP	EP	FT	IS	MG
PACKET	526.15	579.42	952.88	0.18	703.82	87.56	50.29
FLOW	68.97	96.05	117.96	0.16	92.28	12.89	15.11
PACKET-FLOW	90.39	100.39	150.18	2.09	99.64	13.74	9.08
FAST	2.1	2.76	2.67	0.63	0.66	0.65	1.02

Figure 5: Simulation Overhead in Seconds, NAS Benchmarks

Figure 6: CR-64 is sensitive to the speedup and slowdown of network

Conclusion & Future Work

- Packet-level simulation does not guarantee more accurate application performance prediction than the hybrid packet-flow model, but at a much higher simulation overhead
- Collecting more traces and conduct a more comprehensive trade-off analysis
- Explore relative performance benefits of various networking options, task-mapping and network congestion models

Acknowledgements

Special thanks to my mentor Mike Lang, Scott Pakin and my academic advisor Xin Yuan. This work is supported by Ultra-Scale Research Center, New Mexico Consortium, Los Alamos National Lab.