Proxy Apps Mysteries Revealed

PI: David Richards (LLNL)
Team: Jeanine Cook (SNL), Hal Finkel (ANL), Christoph Junghans (LANL), Peter McCorquodale (LBNL) Shirley Moore (ORNL), Omar Aaziz (SNL), Tanner Juedeman (SNL), Courtenay Vaughan (SNL), Brian Homerding (ANL), Thomas Uram (ANL), Abhinav Bhatele (LLNL), Xavier Andrade (LLNL), Robert Pavel (LANL), Vinay Ramakrishnaiah (LANL), Tiffany Mintz (ORNL), Greg Watson (ORNL)

Mission

To curate a suite of proxy applications that are representative of the intended characteristics of their respective parent applications and are easy to obtain and use. Characteristics include hardware bottlenecks (e.g., memory, computation, communication) and programming models.

Proxy App Suite v2.0

AMG	miniVite
CANDLE Benchmarks	NEKbone
Ember	PICSARlite
ExaMiniMD	SW4lite
Laghos	SWFFT
MACSio	thornado-mini
miniAMR	XSBench
miniQMC	

- 53 proxy applications now in our catalog. Always looking for more
- Each proxy app can be installed with a simple Spack command
  - Dependencies are easily identified and provided
  - See https://proxyapps.exascaleproject.org/downloads
- Improved metadata to track provenance and sponsors of proxies
- The Proxy App Team worked with ECP Projects to develop or enhance new proxy apps
  - PICSARlite: Reduced code size by 80%
  - miniVite: Consulted on design of proxy and selection of algorithm
  - thornado-mini: Helped with documentation and testing

Proxy Catalog Performance Characterization

Latency

Bandwidth

Compute

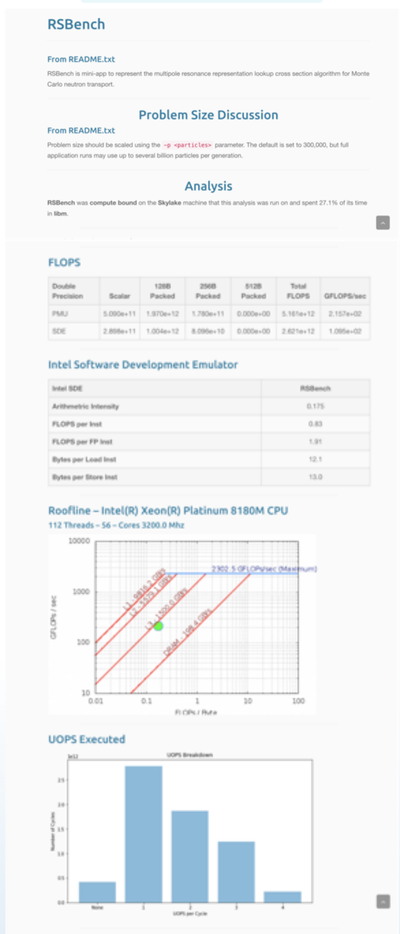

Added additional information to the proxy app catalog

- Performance bounds and data
- Problem size and run information
- Important code segments with cache and memory bandwidth data

Quantitative Assessment

Goal: Understand how well proxies represent parent applications quantitatively at hardware level → better proxies in future

- Representative problems/sizes
- Detailed profiling
- Quantitative characterization
- Statistical proxy/parent app comparison

Quantitative Comparison

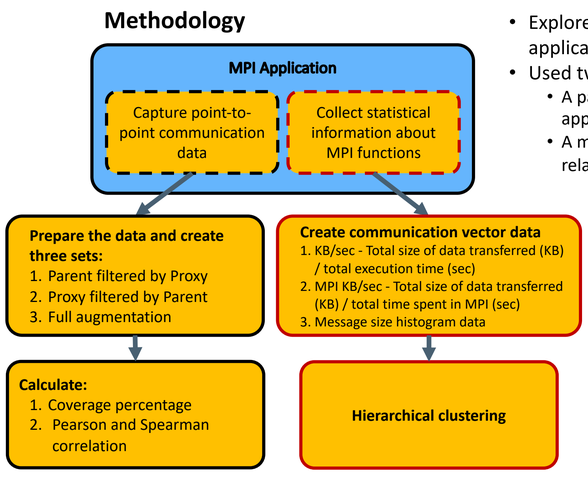

- Explored how the communication of a proxy application relates to its parent application
- Used two quantification methods:
  - A pairwise communication quantification method that captures how much one application matches the other
  - A message characteristics quantification method that produces a clustering-based relatedness measure of the parent applications and their proxies

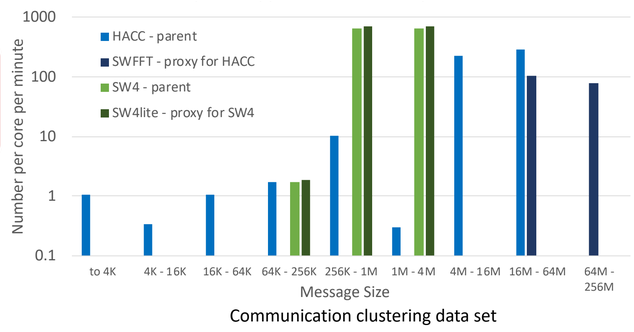

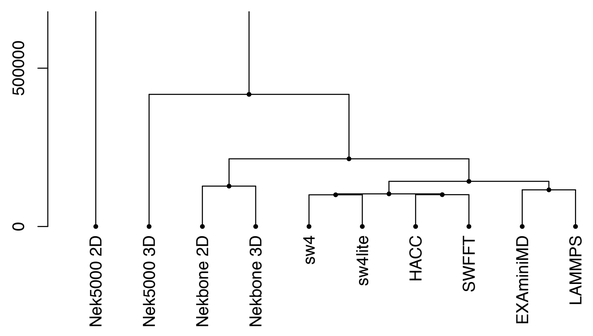

Communication Similarity, Clustering Message Size, and Frequency

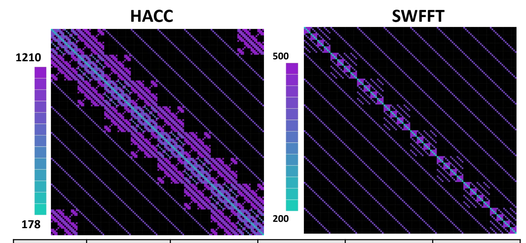

Parent/Proxy	Parent in Proxy		Proxy in Parent		Full Set		Parent in Proxy		Proxy in Parent	
	#msg	#pair	#msg	#pair	PCorr	SCorr	PCorr	SCorr	PCorr	SCorr
HACC/ SWFFT	68.7	41.1	100	100	0.97	0.99	0.92	0.84	0.97	0.99

Point-to-Point Communication Analysis

GPU Characterization of Proxy Apps

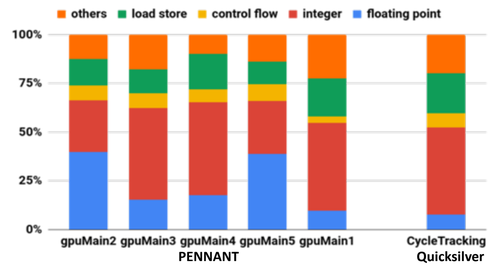

Instruction mix for PENNANT (sedovbig) and Quicksilver (Coral2)

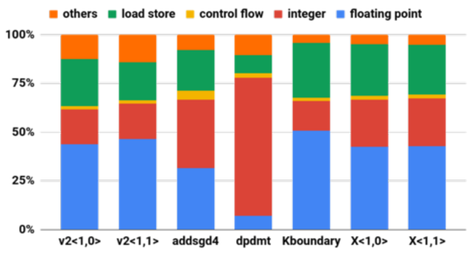

Instruction mix for SW4lite (LOH1-h100)

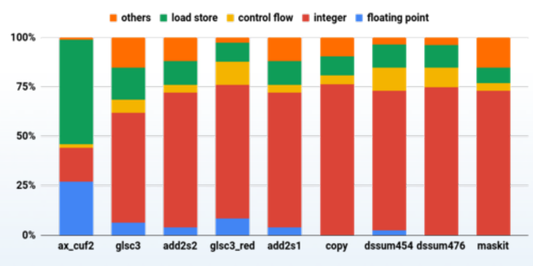

Instruction mix for NEKbone

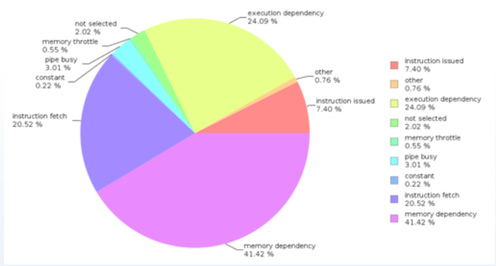

Stall cycle breakdown for PENNANT gpumain2 kernel on V100

Observations:

- All four proxy apps are limited by memory performance on P100 and V100 GPUs.
- Applications that are latency-bound on P100 remain latency-bound to same degree on V100, in spite of improvements to cache/memory hierarchy.
  - SW4lite improved from 14% to 16% of peak
  - PENNANT remained at 8%
- Above instruction mixes are shown for P100. Integer instruction counts, which include address calculations, are reduced and can be performed simultaneously with floating point operations on V100 and thus have less impact on V100.

U.S. DEPARTMENT OF ENERGY
National Nuclear Security Administration

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.
SAND No. _____

EXASCALE COMPUTING PROJECT
OAK RIDGE National Laboratory
Argonne NATIONAL LABORATORY
Los Alamos NATIONAL LABORATORY EST. 1943
BERKELEY LAB
Lawrence Livermore National Laboratory
Sandia National Laboratories