Workload Distribution on Heterogeneous Systems based on Metaheuristics and Machine Learning: A Performance and Energy Aware Approach

Suejb Memeti and Sabri Pllana

Linnaeus University, Sweden http://homepage.lnu.se/staff/saplaa/

6th Swedish Workshop on Data Science (SweDS18) Umeå, SE, Nov 20 -- 21, 2018



About Me

- Associate Professor (since 2013)
 - Department of Computer Science, LNU
- Before joining LNU
 - worked 12 years at the Vienna University (est.1365) in Austria
 - coordinated PEPPHER project (EU FP7)
- Selected memberships
 - Senior IEEE Member
 - CHIPSET, European COST action on Big Data
 - ETP4HPC, European platform for HPC
 - HiPEAC, European network on highperformance and embedded systems



PEPPHER presentation at HiPEAC 2012 conference January 23-26, 2012, Paris, France



November 2018 www.top500.org

Rmax is maximal performance achieved in the High Performance Computing LINPACK benchmark

• Pflop/s or petaflop: 10¹⁵ FLOPS

		SPECS	SITE	COUNTRY	CORES	RMAX PFLOP/S	POWER
1	Summit	IBM POWER9 (22C, 3.07GHz), NVIDIA Volta GV100 (80C), Dual-rail Mellanox EDR Infiniband	DOE/SC/ORNL	USA	2,282,544	143.5	11.1
2	Sierra	IBM POWER9 (22C, 3.1GHz), NVIDIA Tesla V100 (80C), Dual-rail Mellanox EDR Infiniband	DOE/NNSA/LLNL	USA	1,572,480	94.6	7.44
3	Sunway TaihuLight	Shenwei SW26010 (2600 1.45 GHz) Custom interconnect	NSCC in Wuxi	China	10,649,600	93.0	15.4
4	Tianhe-2A (Milkyway-2A)	Intel Ivy Bridge (120 2.2 GHz) & TH Express-2, Matrix-2000	NSCC Guangzhou	China	4,981,760	61.4	18.5
5	Piz Daint	Cray XC50, Xeon E5-2690v3 (1202.66Hz), Aries interconnect , NVIDIA Tesla P100	CSCS	Switzerland	319,424	21.2	2.38



Heterogeneous Systems

- Also known as accelerated or asymmetric systems
- Common: GPU-accelerated systems
- Our "Emil" system at LNU uses Intel Xeon Phi as accelerator





Summit at ORNL

- Rank 1 in TOP500 list (Nov 2018)
- 200 Pflop/s (max performance)
- 143.5 Pflop/s (LINPACK performance)
- 4608 computing nodes
- each node comprises
 - two IBM Power9 22-core CPUs
 - six NVIDIA Volta GPUs





Performance of Accelerators and #1 in TOP500





Optimization of Workload Distribution





Genomics is the Largest Data Source



Projected data growth by the year 2025 [Stephens et al. 2015]



Application: DNA Sequence Analysis

- A kind of string matching problem
- Search for a set of patterns in DNA sequences
- Experiment with Human DNA sequence (3GB) from the GenBank sequence database (http://www.ncbi.nlm.nih.gov/genbank)





Experimentation Platform: Emil at LNU





Workload Partitioning of DNA Sequence Analysis





Large Parameter Space

- Number of possible system configurations
 - a product of parameter value ranges
 - each parameter c_i has a value range R_{ci}

$$\prod_{i=1}^{n} R_{c_i} = R_{c_1} \times R_{c_2} \times \ldots \times R_{c_n}$$

Considered parameters on our experimentation platform

Parameters	Parameter values
CPU threads (CPU-T) Accelerator Threads (ACC-T) CPU thread affinity (CPU-A) Accelerator thread affinity (ACC-A) CPU Workload Fraction (CPU-W) Accelerator Workload Fraction (ACC-W)	<pre>{12, 24, 36, 48} {60, 120, 180, 240} {none, scatter, compact} {balanced, scatter, compact} {1100} {100 - CPU-W}</pre>



Optimization using Metaheuristics and Machine Learning

- Challenges
 - large discrete parameter space
 - evaluation of all possibilities is impractical
- Metaheuristics: search for a near-optimal solution
 - avoid evaluating all feasible solutions
- Simulated Annealing: a metaheuristic that is inspired by thermodynamics
 - slowly cooled liquid forms a crystal (minimum energy state)
- Machine Learning: learn from historical performance data
 - build a performance model of system under study
 - use Linear Regression, Boosted Decision Tree Regression,...





Comparison of Linear Regression (LR) and Boosted Decision Tree Regression (BDTR)

 Relative error of BDTR is smaller than of LR for larger number of training experiments





Enumeration and Measurements (EM)

- Parameter space exploration using Enumeration
- Solution evaluation using Measurements





Enumeration and Machine Learning (EML)

- Parameter space exploration using Enumeration
- Solution evaluation using Machine Learning





Simulated Annealing and Measurements (SAM)

- Parameter space exploration using Simulated Annealing
- Solution evaluation using Measurements





Simulated Annealing and Machine Learning (SAML)

- Parameter space exploration using Simulated Annealing
- Solution evaluation using Machine Learning





Results: Power Efficiency [MB/W]

- SAML and SAM power efficiency optimization results are close to EM
 - SAML and SAM use only 5% of experiments required by EM





Summary

Features of studied approaches for workload distribution

Method	Space Exploration	Sys. Conf. Evaluation	Effort	Accuracy	Prediction
EM	Enumeration	Measurements	High	Optimal	No
EML	Enumeration	Machine Learning	High	Near-optimal	Yes
SAM	Simulated Annealing	Measurements	Medium	Near-optimal	No
SAML	Simulated Annealing	Machine Learning	Medium	Near-optimal	Yes

Key papers

- Memeti, S., & Pllana, S. (2017). Combinatorial optimization of DNA sequence analysis on heterogeneous systems. Concurrency and Computation: Practice and Experience, 29(7), e4037.
- Memeti, S., & Pllana, S. (2018). A machine learning approach for accelerating DNA sequence analysis. The International Journal of High Performance Computing Applications, 32(3), 363-379.



