WASHINGTON STATE UNIVERSITY
School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

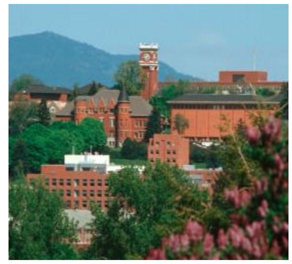

Shervin Hajiamini

A FAST HEURISTIC FOR TASKS ASSIGNMENT IN MANYCORE SYSTEMS WITH VOLTAGE-FREQUENCY ISLANDS

WASHINGTON STATE UNIVERSITY
School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

INTRODUCTION

- Increasing computations in applications has led to faster processing.
  - Use more cores in a chip with more parallelism.
- These applications are typically composed of phases with different computation/memory access characteristics.
- Chip-wide voltage/frequency setting does not address applications with variable computations.

WASHINGTON STATE UNIVERSITY
School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

INTRODUCTION

- Power/Energy-aware methods have to maintain reasonable power budget with increasing the number of cores in a single chip.
- Voltage Frequency Islands (VFI)
  - Multicore system is partitioned into islands.
  - Adjust the islands' voltage/frequency levels to reduce system energy consumption within a desired performance penalty.

September 2018

WASHINGTON STATE UNIVERSITY
School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

ISLANDING

- Fine-grain to coarse-grain
  - Tradeoff aggressive energy saving vs hardware design simplicity and cost

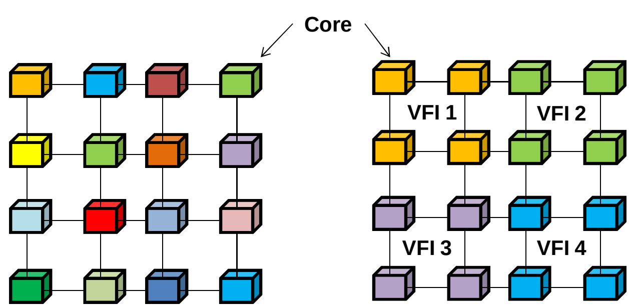

September 2018

WASHINGTON STATE UNIVERSITY

School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

WORKLOAD BALANCE (PER VFI)

- In an island, cores may have different computation workloads.
- Increasing or decreasing the V/F level of such an unbalanced island consumes extra energy or delays the execution time of application.

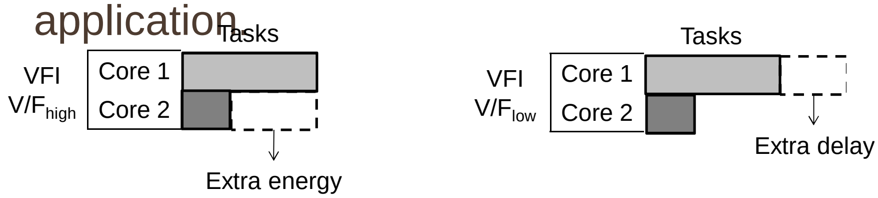

WASHINGTON STATE UNIVERSITY

School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

RESEARCH MOTIVATION

- Address manycore energy efficiency on a symmetric VFI architecture
  - Optimize task-to-island assignments (task partitioning)
  - Optimize V/F level assignments for islands
- V/F levels are dynamically assigned per execution phase of applications
- Optimization objective
  - Minimize application execution time (makespan) without violating energy consumption budget.

RELATED WORK

- Research works that utilize VFIs for energy efficiency fall in one or more of the following categories:
  - Design-time (static) vs. runtime (dynamic) optimization.[ref]
  - Symmetric vs. non-symmetric VFIs.[ref]
  - Solve islanding and V/F level assignment sub-problems simultaneously or individually.[ref]
  - Improve the energy efficiency of VFI-based system w/o considering the similarities of workloads of cores.[ref]
  - Provide exact or (meta) heuristic solutions to solve the VFI's sub-problems.[ref]

WASHINGTON STATE UNIVERSITY | School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

RELATED WORK (CONT.)

- Closely related works
  - K. Duraisamy et al. *Energy Efficient MapReduce with VFI-enabled Multicore Platforms*, DAC 2015.
  - R. Kim et al. *Wireless NoC and Dynamic VFI Codesign: Energy Efficiency Without Performance Penalty*, IEEE TVLSI 2016.
  - S. Pagani et al. *Energy Efficient Task Partitioning based on the Single Frequency Approximation Scheme*, IEEE RTSS 2013.

September 2018

Washington State University | School of Electrical Engineering and Computer Science | Voiland College of Engineering and Architecture

SYSTEM MODEL

- $\sqrt{N} \times \sqrt{N}$ multicore system (N is the number of cores)
- System is composed of VFIs that each associates with a V/F pair

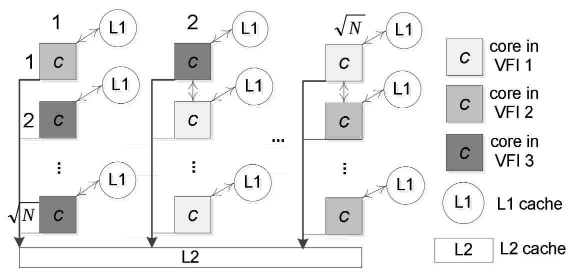

September 2018

Washington State University | School of Electrical Engineering and Computer Science | Voiland College of Engineering and Architecture

TASK EXECUTION MODEL

- Benchmarks
  - Multi-threaded workloads executed on a distributed shared memory platform.
  - The performance of our proposed methodology is optimized using data collected from workload's parallel section (region of interest).
- Execution phase
  - The parallel section consists of a sequence of execution phases.
  - The execution phases are separated by synchronization function calls (locks and

WASHINGTON STATE UNIVERSITY | School of Electrical Engineering and Computer Science, VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

TASK EXECUTION MODEL (CONT.)

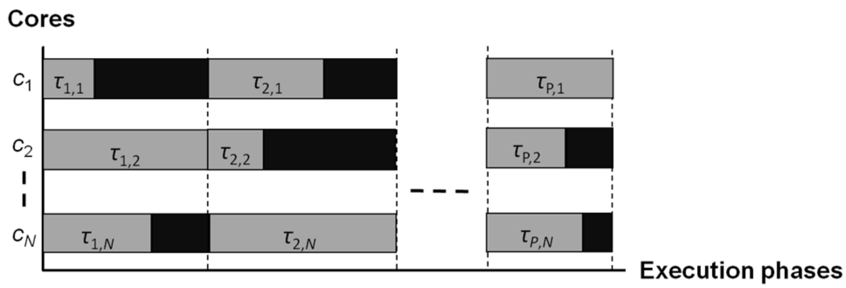

September 2018

WASHINGTON STATE UNIVERSITY | School of Electrical Engineering and Computer Science, VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

TASK EXECUTION PROFILES

- The task set of each benchmark is run multiple times, each time with a fixed V/F level for the entire execution run.
- For each task, three parameters are profiled
  - Execution time
  - Energy consumption
  - Workload/Utilization
- Statically use this profile to optimize the task-to-island and dynamic V/F level assignments.

September 2018

CONTRIBUTION

- Two-step optimization framework for minimizing makespan, given an energy budget, for symmetric VFIs.
  - Formulate the task-to-island assignment problem using mixed integer linear programming .
  - Formulate the dynamic V/F level assignment problem for the VFIs using integer linear programming.
- Propose a fast heuristic that obtains near-optimal solutions for solving the task-to-island assignment problem.

WASHINGTON STATE UNIVERSITY | School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

CONTRIBUTION (CONT.)

- Evaluate the optimization framework on benchmarks with different computational characteristics.
- Compare the energy efficiency of proposed framework to optimal per-core VFIs.
  - Use Energy-Delay Product (EDP) and IPS^2/Watt as well-known performance metrics.

September 2018

WASHINGTON STATE UNIVERSITY
School of
Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

TASK-TO-ISLAND ASSIGNMENT

- Identify and group tasks with similar workloads.
- Define similarity measure based on the difference between the task workload and the maximum workload in a given execution

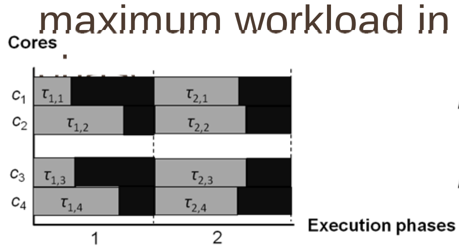

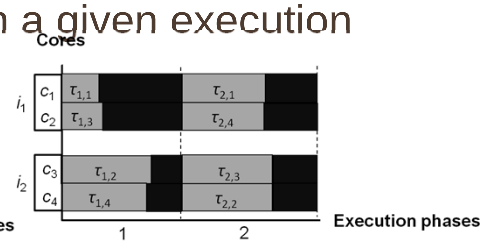

WASHINGTON STATE UNIVERSITY
School of
Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

TASK-TO-ISLAND ASSIGNMENT PROBLEM FORMULATION

	Description	Formulation
Objective	Minimize the wasted workload per phase.	$\underset{y_{i,k},F_k,c_k,z_{k,j},x_{i,k}}{Minimize} \quad Y_j = \sum_{i=1}^{N}\sum_{k=1}^{K} y_{i,k} \quad \forall T_j \in T$
Constraints	Compute the wasted workload for a task.	$y_{i,k} \geq x_{i,k} - t_i \cdot F_k \quad \forall \tau_{j,i} \in T_j, \forall i_k \in I$
	Approximate the maximum workload and its reciprocal.	$c_k = \sum_{j=1}^{r} a_j \cdot z_{k,j} \quad \forall i_k \in I$ $F_k = \sum_{j=1}^{r} \left(\frac{1}{a_j}\right) \cdot z_{k,j} \quad \forall i_k \in I$
	Determine one of the tasks, in the phase, as the one with the maximum workload.	$t_i \cdot x_{i,k} \leq c_k \quad \forall \tau_{j,i} \in T_j, \forall i_k \in I$
	A task is assigned to only one island. All islands have the same number of tasks (symmetry of islands). Compute the probability of approximating the maximum workload with a line segment	$\sum_{k=1}^{K} x_{i,k} = 1 \quad \forall \tau_{j,i} \in T_j$ $\sum_{i=1}^{K} x_{i,k} = Q \quad \forall i_k \in I$ $\sum_{j=1}^{r} z_{k,j} = 1 \quad \forall i_k \in I$

DYNAMIC V/F LEVEL ASSIGNMENT

- An Island may have different amounts of workloads across execution phases.
- Changing the V/F levels of the islands helps increase energy saving and performance.
  - Speed up islands with high workloads.
  - Slow down islands with low workloads.

V/F levels: $s_1 < s_2 < s_3 < s_4$

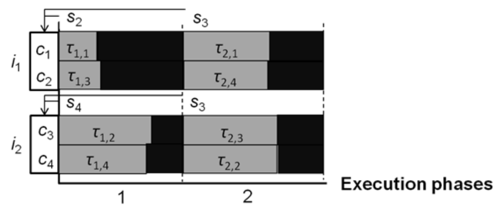

Washington State University | School of Electrical Engineering and Computer Science | Voiland College of Engineering and Architecture

DYNAMIC V/F LEVEL ASSIGNMENT FORMULATION

	Description	Formulation
Objective	Minimize the makespan.	$\underset{\theta_j, a_{k,l,j}}{Minimize} \quad \theta = \sum_{j=1}^{P} \theta_j$
Constraints	Compute the length of an execution phase after assigning the V/F levels of islands.	$\sum_{l=1}^{L} d_{k,l,j} \cdot a_{k,l,j} \le \theta_j \quad \forall i_k \in I, \forall T_j \in T$
	Only one V/F level is assigned to an island in a given execution phase.	$\sum_{l=1}^{L} a_{k,l,j} = 1 \quad \forall i_k \in I, \forall T_j \in T$
	The energy consumption of system is below the pre-defined energy budget.	$\sum_{j=1}^{P}\sum_{k=1}^{K}\sum_{l=1}^{L} e_{k,l,j} \cdot a_{k,l,j} \le EB \quad EB \ge 0$

FAST HEURISTIC FOR TASK-TO-ISLAND ASSIGNMENT

- Tasks are sorted by their workloads ($O(N \log(N))$)
- Sorted tasks are assigned to islands ($O(1)$)
- The heuristic solutions are very close to the MILP-based solutions.

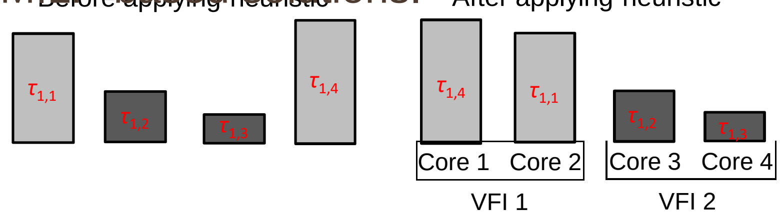

Washington State University
School of Electrical Engineering and Computer Science
Voiland College of Engineering and Architecture

SIMULATION SETUP

- We use GEM5 as a full system simulator to obtain processor-level performance information.
- Processor-level statistics generated by GEM5 simulations are fed to McPAT (Multi-core Power, Area and Timing).
- McPAT generates processor-level power/energy statistics.

Voltage (V)/Frequency (GHz) levels

0.5	1.25
0.667	1.666
0.834	2.083
1.0	2.5

Processor configuration

Processors	64 Alpha cores
L1-cache	64kByte, 4-way associative, 64 Byte
L2-cache	Shared 8 MBytes, 8-way associative, 64 Byte 128KBytes distributed per core
Main memory	512 MBytes

September 2018

BENCHMARKS

- Three applications from SPLASH-2 and PARSEC benchmark suites are considered in our simulations.

Benchmark	Problem size	Application domain
FFT	65536 Data Points	Fast Fourier Transform
LU	512x512 Matrix, 16x16 Blocks	Dense matrix computation
CANNEAL	200,000 Elements	Minimize routing cost in chip with simulated annealing

WASHINGTON STATE UNIVERSITY
School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

VFIS CONFIGURATIONS

- Implemented the heuristic in MATLAB.
- Modeled the problem formulations in AMPL.
  - Used Gurobi to solve the problem.
- Symmetric VFIs: 4 VFIs, 16 cores per VFI.

September 2018

RESULTS

- Compare the proposed Dynamic Coarse-Grain (DCG) VFIs against optimal Fine-Grain (FG) VFIs.
  - Execution time
  - Energy efficiency
- Compare V/F level distributions for VFIs across the benchmarks.

WASHINGTON STATE UNIVERSITY
School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

PERFORMANCE EVALUATION

- Compare the execution times of FG and DCG
  - Energy budget
  - Computation intensiveness of benchmarks
- Three energy budgets (EB) that reduce cores' highest energy usage by
  - a) 7.5% EB(H)
  - b) 22.5% EB(M)
  - c) 37.5% EB(L)

September 2018

WASHINGTON STATE UNIVERSITY

School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

PERFORMANCE EVALUATION (EXECUTION TIME)

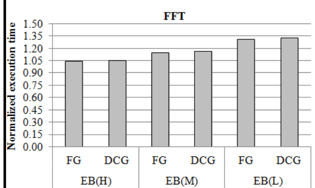

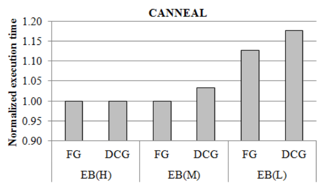

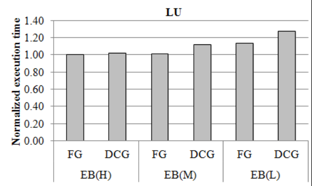

September 2018

WASHINGTON STATE UNIVERSITY

School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

PERFORMANCE EVALUATION (ENERGY EFFICIENCY)

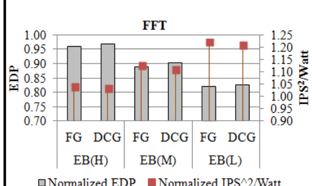

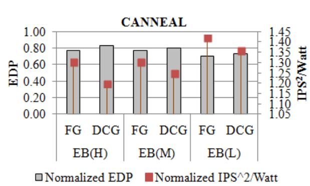

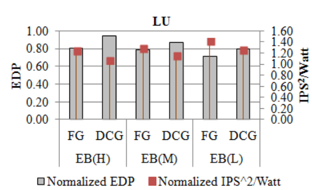

September 2018

PERFORMANCE EVALUATION (V/F LEVEL DISTRIBUTIONS)

Percentage of V/F levels assigned to VFIs

Benchmark	s_1	s_2	s_3	s_4
FFT	0	0	31	69
LU	16	10	12	62
CANNEAL	0	0	65	35

WASHINGTON STATE UNIVERSITY
School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

REAL-WORLD APPLICATION OF THE OPTIMIZATION FRAMEWORK

- Use for kernel applications with moderate workload variations across execution phases.
- Optimize the task and V/F level assignments once at compile-time.
- Store optimization outcomes and use them many times while running the applications.

September 2018

WASHINGTON STATE UNIVERSITY
School of Electrical Engineering and Computer Science
VOILAND COLLEGE OF ENGINEERING AND ARCHITECTURE

CONCLUSION

- Proposed a two-step optimization framework for minimizing makespan while maintaining energy consumption below pre-defined budget.
- Proposed a fast heuristic for balancing the computational workload of each island per execution phase.
- Evaluated the performance of framework against the optimal per-core VFIs, as well as on different compute-intensive benchmarks.