Tangram: Integrated Control of Heterogeneous Computers

Raghavendra Pradyumna Pothukuchi*, Joseph L. Greathouse†, Karthik Rao†, Christopher Erb†, Leonardo Piga†, Petros Voulgaris*, Josep Torrellas*

*ILLINOIS †AMD

1 Heterogeneity and Modularity

Heterogeneous systems are integrated from independent subsystems

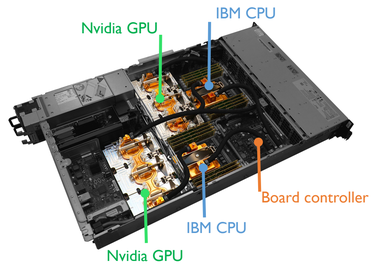

A node of the Summit supercomputer (discrete integration)

Image courtesy: NVIDIA

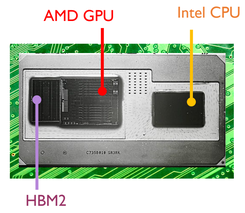

Intel Kaby Lake G (On-package integration)

Challenge: achieve full system efficiency by distributing resources (e.g., power, storage) and goals (e.g., performance) across the subsystems

2 State-of-the-Art Resource Controllers

Limitation 1: many piecemeal heuristics

- Conflicting and ineffective
- Fragile and difficult to reuse

Limitation 2: centralized control architectures

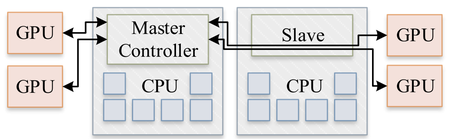

- Slow
- Ineffective because subsystems do not expose internal details
- Difficult to reuse

Why is the state-of-the-art so limited?

Making distributed control work correctly is hard!

Heuristics (falsely) appear intuitive

3 Designing a Modular Controller

Resource control goals		
	Safety	Immediate response
	Enhancement	Search and optimal control
	Preconfigured	Immediate optimal control

Enhancement with Robust Control

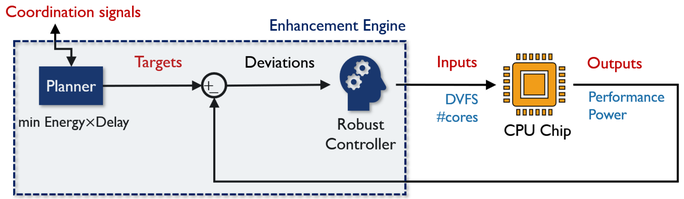

Robust controller:

Matrices **[A, B, C, D]**, and *state*

$$\text{inputs} = \mathbf{C} \times \text{state} + \mathbf{D} \times \text{deviations}$$

$$\text{state}_{new} = \mathbf{A} \times \text{state} + \mathbf{B} \times \text{deviations}$$

Feature: deviations are small even when runtime conditions differ from design conditions

Proposed Controller

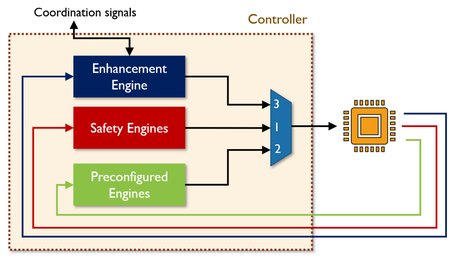

Enhancement is robust to interference from other engines and subsystems

Proposed Interface

Optimize locally and propagate constraints hierarchically

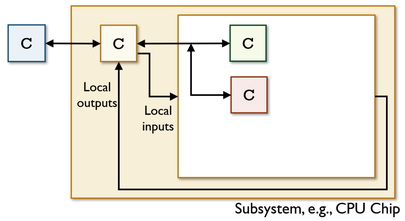

Parent to Child

max power,
min performance, etc.

Child to Parent

power used,
performance achieved, etc.

4 Tangram: A Fast, Modular and Coordinated Framework

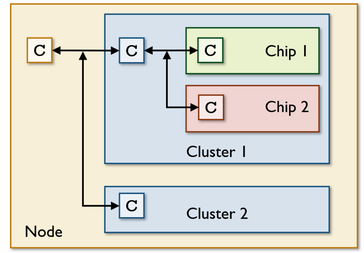

Naturally fits the underlying system organization

5 Prototype and Evaluation Highlights

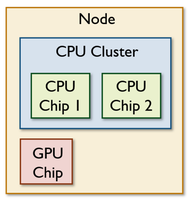

Minimize Energy Delay Product with power constraints,
Current and temperature safety

CPUs: AMD Ryzen™ 1800x , GPU: AMD Radeon™ RX 580

Responsiveness

Tangram: **common case: 15 ms (local response)**
worst case: 515 ms (node response)

Centralized: **200 ms always**

Evaluation with Chai Applications

(Fine-grained concurrent CPU-GPU execution)

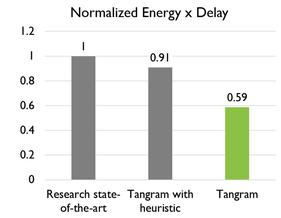

Example: *rsct*

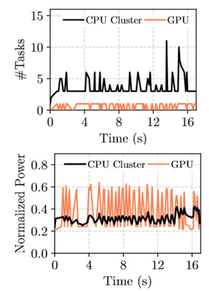

Tangram's 41% better EDP = 32% faster with 13% lower energy

Tangram: Modular Efficient Systematic