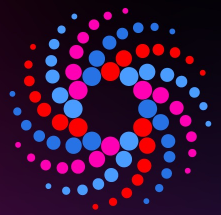

SC22
Dallas, TX | hpc accelerates.

Adopting Heterogeneous Computing Modules: Experiences from a ToUCH Summer Workshop

David P. Bunde, **Kishwar Ahmed**, Sridevi Ayloo, Tisha Brown-Gaines, Joel Fuentes, Vishwesh Jatala, Ruth Kurniawati, Isil Öz, Apan Qasem, Philip J. Schielke, Mary C. Tedeschi, and Thomas Yeh

The Rise of Heterogenous Computing

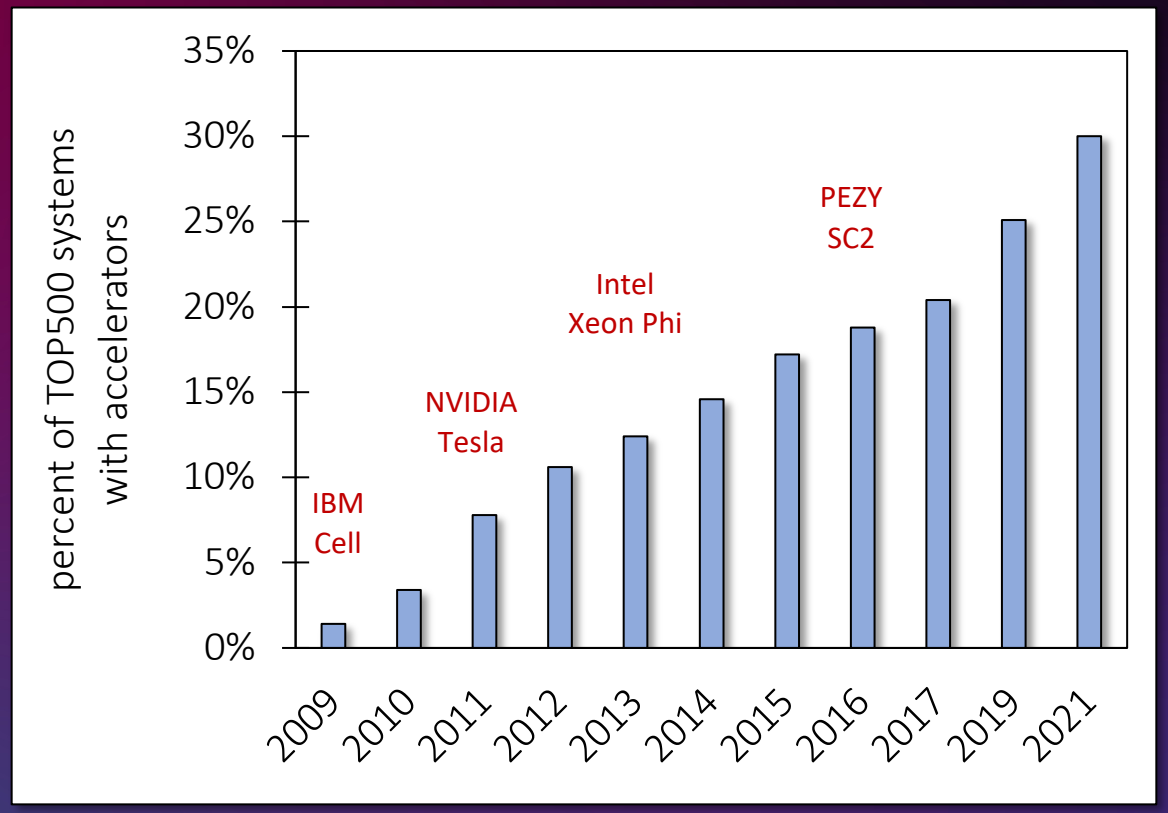

The Rise of Heterogenous Computing

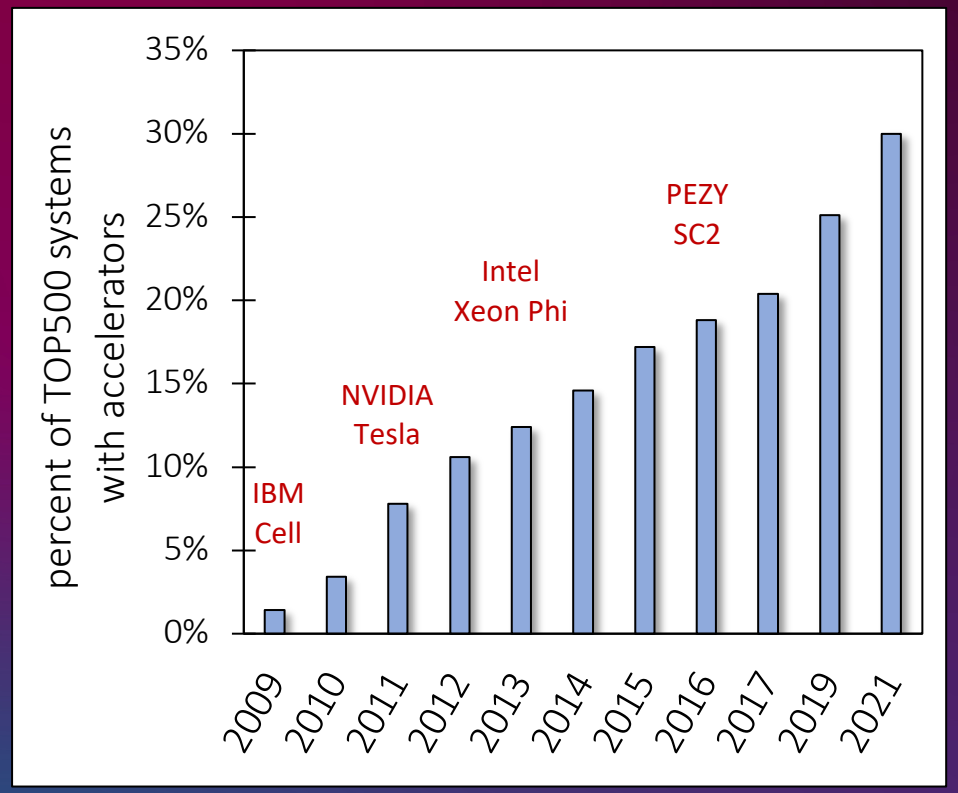

DVFS

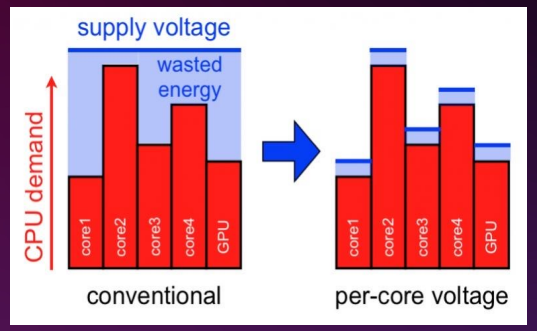

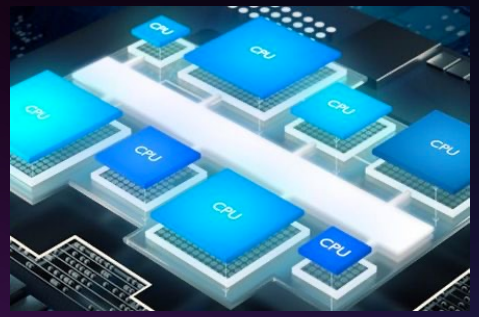

Fat and Thin Cores

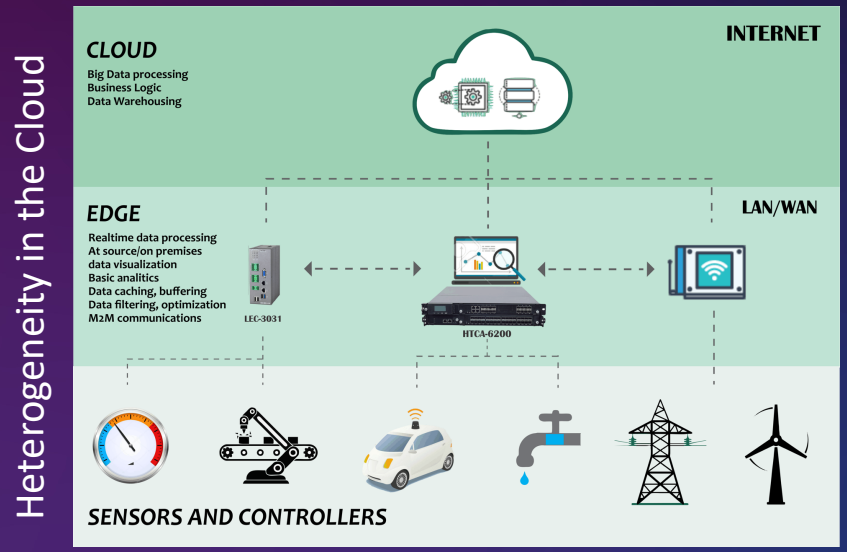

Challenges in Integrating HC into the Curriculum

- Given current industry trends, it is critically important that CS/CE undergraduates are exposed to heterogeneous computing concepts as a required part of the curriculum

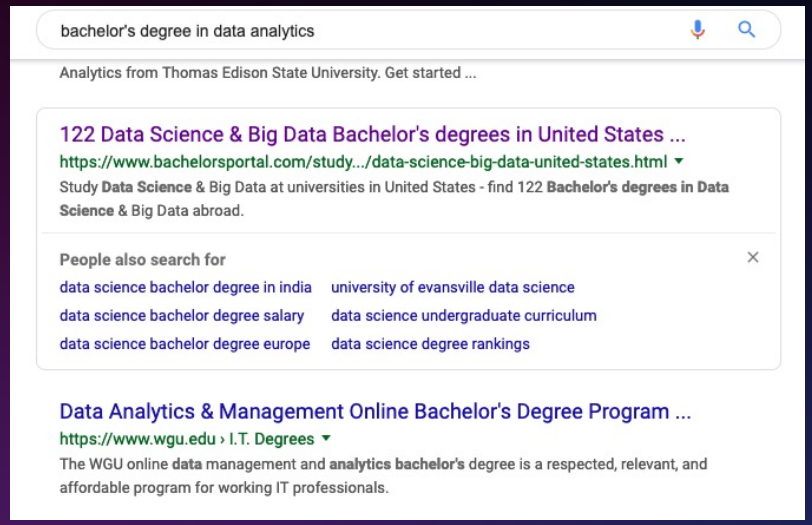

- But, how can it fit?
  - Field is increasingly broad (AI, Cyberscurity, Data Science are also gaining coverage)
  - Not allowed to keep adding courses

The Early-and-Often Approach

- Breaks away from the traditional approach of teaching a topic X as an upper-level elective
- Introduce X early in the curriculum in required courses
- Repeat key concepts in different courses throughout the curriculum
- Tie concepts together in an upper-level capstone course

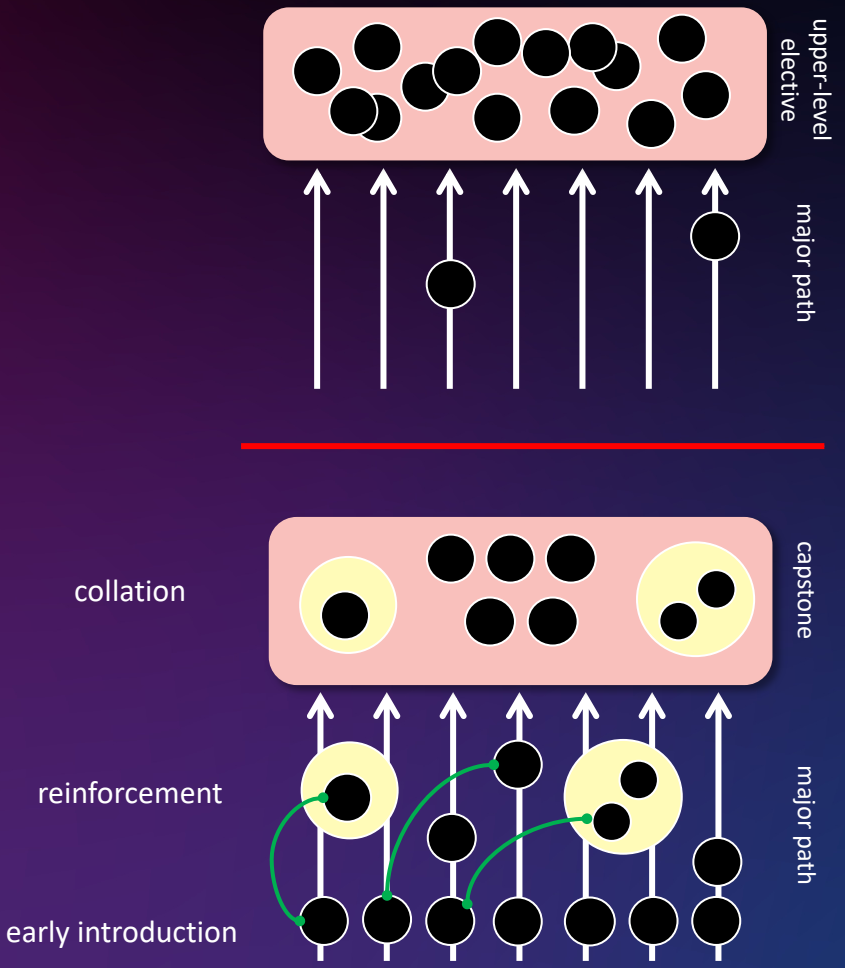

ToUCH project and how we all met

- ToUCH project (Teaching Undergraduates Collaborative and Heterogeneous computing)
  - Implementing modules and training others to adopt them
  - https://touch.cs.txstate.edu/

- One of those trainings: Virtual workshop in Summer 2021
  - Modules implemented in 2021-22 academic year and results accumulated

ToUCH project modules

https://github.com/TeachingUndergradsCHC/modules

- Fundamentals
  - [A1] Heterogeneous Computing: Elementary Notions
  - [A2] Task Mapping on Soft Heterogeneous Systems

  } Previously presented in EduHPC 2019
  - [A3] Pollack's Rule as a Justification for Heterogeneous Computing
- Algorithms
  - [B1] Hybrid Algorithms
- Architecture
  - [C1] Introduction to ARM
  - [C2] GPU Memory Hierarchy
- Programming Models
  - [D1] Introduction to CUDA Programming
  - [D2] Heterogeneous Programming with OpenMP
  - [D3] Introduction to SYCL Programming (based on D1, but developed by an adopter)

Summary of module adoptions

Course	Institution	n	Modules
Intro to Computer Fundamentals	St. John's University	59	A1
Computer System Management and Support	CityTech, City Univ. New York (CUNY)	5	A1
Networking Fundamentals	CityTech, City Univ. New York (CUNY)	5	A1
Intro to Computer Security	CityTech, City Univ. New York (CUNY)	48	A1
Information Systems Development Project	Baruch College	25	A1
Programming II	Belmont University	27	A1
Programming Languages	Belmont University	30	A1
Intro to Computing Systems	Knox College	8	A1, A3, D1
CS 2	Westfield State University	15	A1, A2
Data Structures	Westfield State University	8	A1, A2
Computer Organization	Pomona College	12	A3, D2
Parallelization of Programs	Indian Institute of Technology Bhiliai	9	D1, C2
Computer Architecture	Izmir Institute of Technology	100	C1, C2
Heterogeneous Computing	Universidad del Bío-Bío	21	A1, A3, B1, C1, C2, D1, D3
High Performance Computing	University of South Carolina Beaufort	13	A1
Computer Architecture	University of South Carolina Beaufort	11	C1

Reported in EduPar 2022

Introduction modules (A1 and A2)

- A1: Teach definitions and concepts of parallelism
- A2: Look at the impact of task mapping on system performance
  - Assigning threads with different needs to different kinds of cores

- Already shown to be effective at Texas State. Now tested at other institutions with similar results

Intro students don’t already see parallelism/heterogeneity

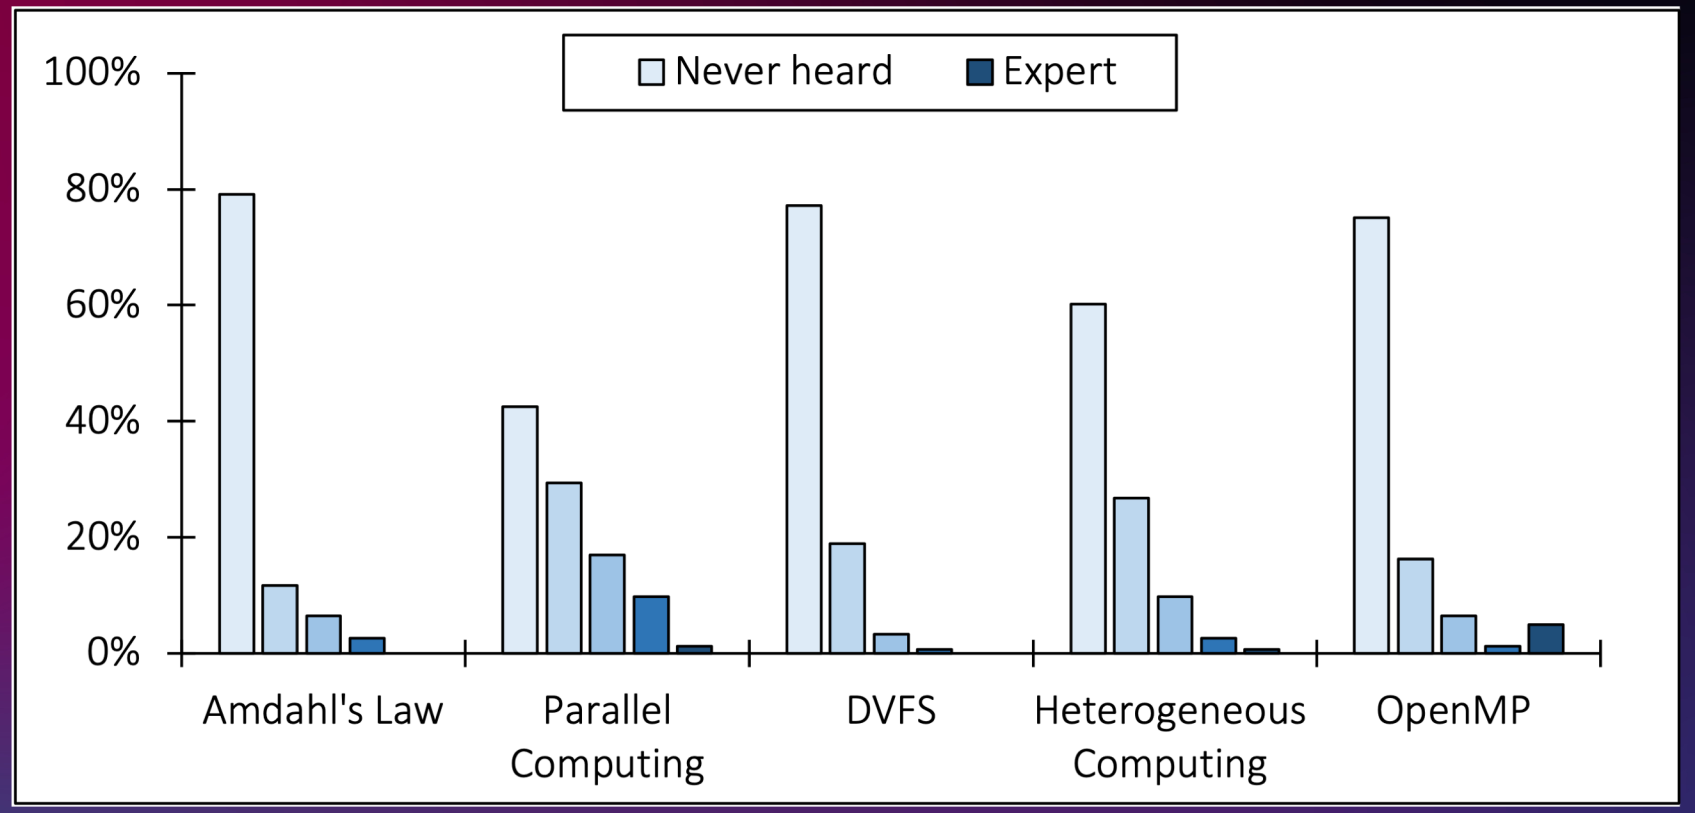

Module A1 combined pre-assessment scores for CUNY, St. John’s Univ,, Baruch College, and Belmont Univ. (n=203)

Pollack's Rule (A3)

- Performance modeling to estimate performance of a core from its size
  - Checking how different configurations of cores affect performance (peak and on kinds of programs)
  - Practice with the concepts of Amdahl's law and motivates use of heterogeneous core sizes
- 6/8 students were able to perform the calculations after the module and another got about half way
- Use by someone other than the creator exposed an implicit assumption made in the module

Architectural modules (C1, D1, C2)

- C1: Introducing ARM assembly
  - Thumb mode gives different tradeoff for power, speed, and code size
  - SIMD instructions
- D1: Introduction to CUDA programming using an image processing application
- C2: Using different kinds of CUDA memory

Qualitative assessment: Harder material, but students seemed to appreciate that understanding architectural features could help them unlock system performance

Again, made improvements to the modules based on issues found through their use

Workshop organization

- Purely virtual
- Split into multiple meetings
  - Orientation meeting
  - Individual/small group meetings with facilitators
  - (week later) Presentations of plans
  - (after next term) Presentation of results
- Stipend split between first set of meetings and presenting results
- Also encouraged adoption with joint paper (this one…)

Workshop organization

- Purely virtual
- Split into multiple meetings
  - Orientation meeting
  - Individual/small group meetings with facilitators
  - (week later) Presentations of plans
  - (after next term) Presentation of results
- Stipend split between first set of meetings and presenting results
- Also encouraged adoption with joint paper (this one…)

Findings

- Doesn’t require travel and cheaper to run
- Greater variety of attendees (including non-US)
  - Created time zone issue and couldn’t fund non-US attendees
- Attendees worked between the meetings
- 10/17 attendees reported on an adoption
- Gathered a lot of data, but challenging to combine it all

General observations

- Modules used in wide variety of ways, far beyond what the creators imagined
  - Before-class preparation for graduate students, in a specialized course, in IT and security programs
- Need to provide assessment instruments with materials
- Slides need more pedagogical information, including more of the text to say and a description of how it fits into the “story” of the lecture
- Hardware is still a barrier to adoption; cloud-based options (Google Colab) were popular for the modules that provided them

Questions?