

## CompE Seminar Series

### Assistant Professor Matt Sinclair

Computer Sciences Department at the University of Wisconsin-Madison

Tuesday, November 11, 2025

Presentation: 12:00 Noon – 1:30 P.M.

MSEE 112

<https://purdue-edu.zoom.us/j/91378862058>

## Rethinking the Control Plane for Chiplet-Based Heterogeneous Systems

**Abstract** In recent years, system designers have increasingly been turning to heterogeneous systems to improve performance and energy efficiency. Specialized accelerators are frequently used to improve the efficiency of computations that run inefficiently on conventional, general-purpose processors. As a result, systems ranging from smartphones to datacenters, hyperscalers, and supercomputers are increasingly using large numbers of accelerators (including GPUs) while providing better efficiency than CPU-based solutions. In particular, GPUs are widely used in these systems due to their combination of programmability and efficiency. Traditionally, GPUs are throughput-oriented, focused on data parallelism, and assume synchronization happens at a coarse granularity. However, programmers have begun using these systems for a wider variety of applications which exhibit different characteristics, including latency-sensitivity, mixes of both task and data parallelism, and fine-grained synchronization. Thus, future heterogeneous systems must evolve and make deadline-aware scheduling, more intelligent data movement, efficient fine-grained synchronization, and effective power management first-order design constraints. In the first part of this talk, I will discuss our efforts to apply hardware-software co-design to help future heterogeneous systems overcome these challenges and improve performance, energy efficiency, and scalability. Then, in the second part I will discuss how the on-going transition to chiplet-based heterogeneous systems exacerbates these challenges and how we address these challenges in chiplet-based heterogeneous systems by rethinking the control plane.

**Bio** Matt Sinclair is an Assistant Professor in the Computer Sciences Department at the University of Wisconsin-Madison. He is also an Affiliate Faculty in the ECE Department and Teaching Academy at UW-Madison. His research primarily focuses on how to design, program, and optimize future heterogeneous systems. He also designs the tools for future heterogeneous systems, including serving on the gem5 Project Management Committee and the MLCommons Power, HPC, and Science Working Groups. He is a recipient of the DOE Early Career and NSF CAREER awards, and his work has been funded by the DOE, Google, NSF, and SRC. Matt's research has also been recognized several times, including an ACM Doctoral Dissertation Award nomination, a Qualcomm Innovation Fellowship, the David J. Kuck Outstanding PhD Thesis Award, and an ACM SIGARCH - IEEE Computer Society TCCA Outstanding Dissertation Award Honorable Mention. He is also the current steward for the ISCA Hall of Fame.