

Management of Complex, Heterogeneous, and Hierarchical Memory Resources



David Beckingsale, Maya Gokhale, **Kamil Iskra**, Michael Lang,
Stephen Olivier, Swann Perarnau, Jeffrey Vetter

Exascale Memory Challenge

- Increasing range of memory types (DDR, HBM, NVM)
- Multiple memory types can be present in one system
- Different memory types can be attached to different compute devices
- Complex performance characteristics (even if all memory is byte-addressable)

All initial exascale systems are expected to feature heterogeneous memory hierarchies.

Challenge for ECP projects:

How to use and manage complex, heterogeneous, hierarchical memory resources in operating systems, runtimes, libraries, and applications?

Plan for This Session

1. Project updates (~7 minutes per speaker)

- 2.3.1.16 – SICM: Simplified Interface to Complex Memories (Mike Lang, Maya Gokhale)
- 2.3.1.18 – RAJA/Kokkos (Umpire – David Beckingsale)
- 2.3.1.19 – Argo: Low-level resource management for the OS and runtime (AML – Swann Perarnau; UMap – Maya Gokhale)
- 2.3.2.10 – PROTEAS-TUNE (Jeff Vetter)
- 2.3.6.03 – Sandia National Laboratories ATDM (OpenMP – Stephen Olivier)

2. Panel / Q & A (~45 minutes)

Initial Questions for the Panel

- Hide or expose – what is the right level of abstraction for heterogeneous memory resources?
- Why come up with new interfaces when NUMA API is already here?
- Unified memory – a blessing or a curse?
- Who should be our customers? Share your most encouraging deployment stories.
- Optane DIMMs are (finally) here – were they worth the wait?

Vote on your favorite questions or add your own!

- <https://tinyurl.com/ecpammm>