

CLARISSE: Cross-Layer Abstractions and Run-time for I/O Software Stack of Extreme-Scale Systems

Florin Isaila

ANL & University Carlos III

Jesus Carretero

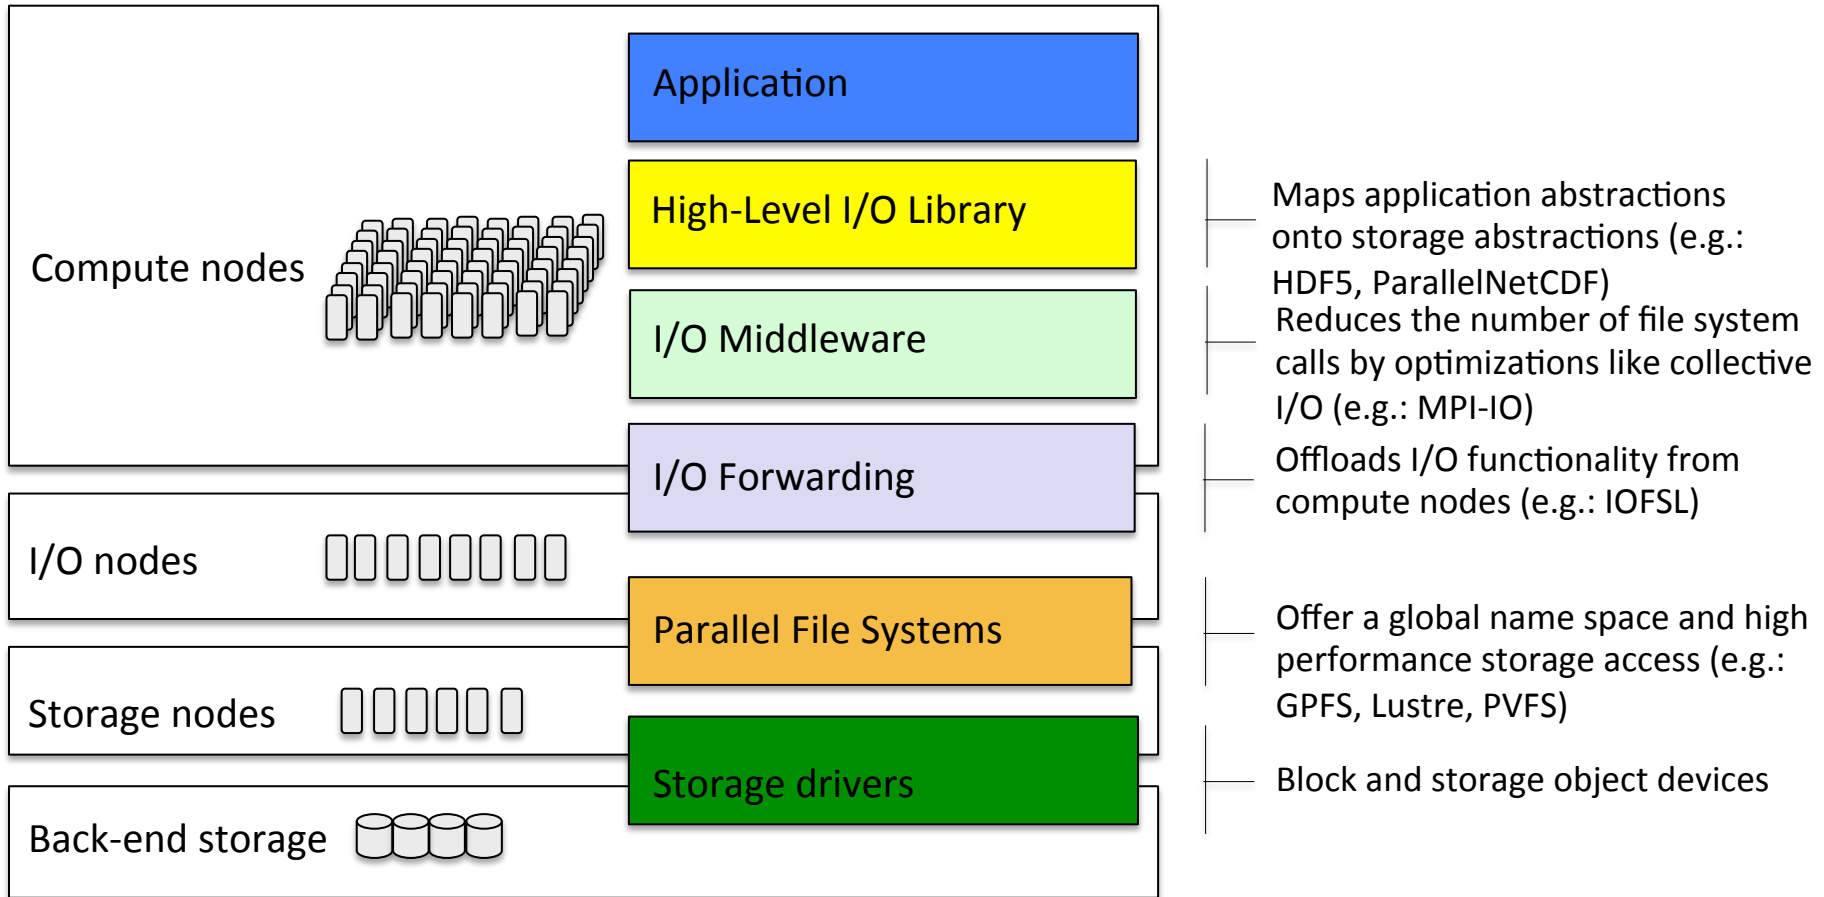
University Carlos III

Rob Ross

ANL

Collaborators: Prasanna Balaprakash (ANL), Phil Carns (ANL), Javier Garcia (UC3M)

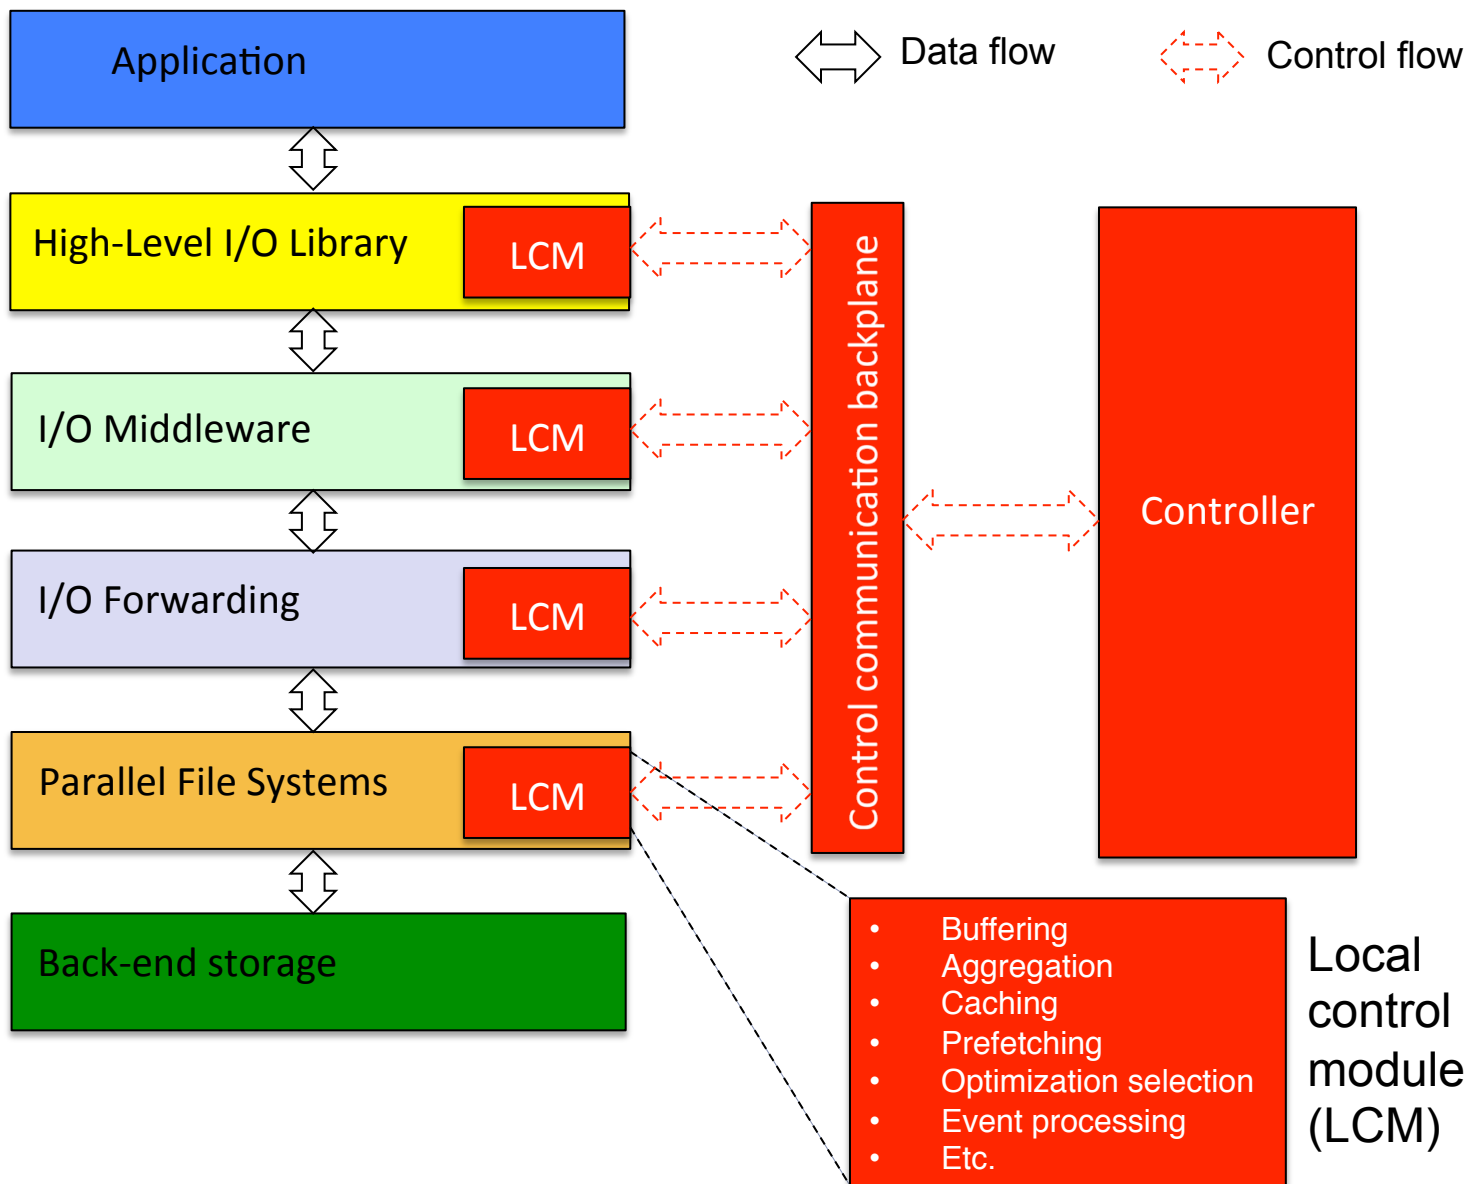
Kevin Harms (ANL), Dries Kimpe (ANL), Rob Latham (ANL), Stefan Wild (ANL)



- ▶ Long path from compute nodes to final storage impacts performance (latency, throughput)
- ▶ Storage I/O optimizations are local: Difficult to perform global optimizations
- ▶ Cross-layer adaptive control mechanisms are not available (e.g for data staging, dynamic load balancing, resilience)



- ▶ Enable global optimizations of the software I/O stack
 - ▶ Evolutionary approach building on many existing optimizations
- ▶ Design novel cross layer control abstractions and mechanisms for supporting data flow optimizations
 - ▶ Collective I/O, data staging, exploit data locality





- ▶ Software Defined Networking (e.g. Open Flow): global control based on separation of control and data flow
- ▶ I/O Flow (Microsoft Research): A Software Defined Storage Architecture for virtualized data centers
- ▶ Fast Forward (Intel et al.): redesign of the storage I/O stack
- ▶ Argo (ANL et al.): whole-system view and optimization based on a OS/R environment
- ▶ Hobbes (Sandia et al.): a lightweight OS/R environment with flexibility to build custom runtimes
- ▶ Cross-layer optimizations for current I/O stack
 - ▶ Parallel I/O autotuning (UIUC & LBNL)
 - ▶ Reduce performance interference (CaLCioM, INRIA & Argonne):
 - ▶ SIOX (University of Hamburg et al.)



- ▶ Mechanisms and abstractions for cross layer control of storage I/O stack
- ▶ Explore the potential of existing storage I/O software stack based on global information
- ▶ Novel cross-layer data flow optimizations: collective I/O, data staging, data locality exploitation

Thank you