Towards Characterization of Data Movement in Large-Scale Scientific Applications

Jun Yi, Rajkumar Kettimuthu and Venkatram Vishwanath

Mathematics and Computer Science Division; Computational Institute
Argonne National Laboratory; University of Chicago
<kettimut,venkatv>@mcs.anl.gov

Application Characteristics

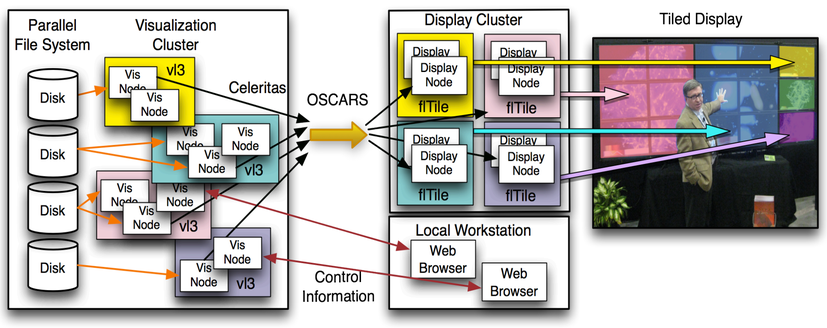

Interactive remote visualization at San Diego of Enzo Cosmology data located at Chicago over ESnet

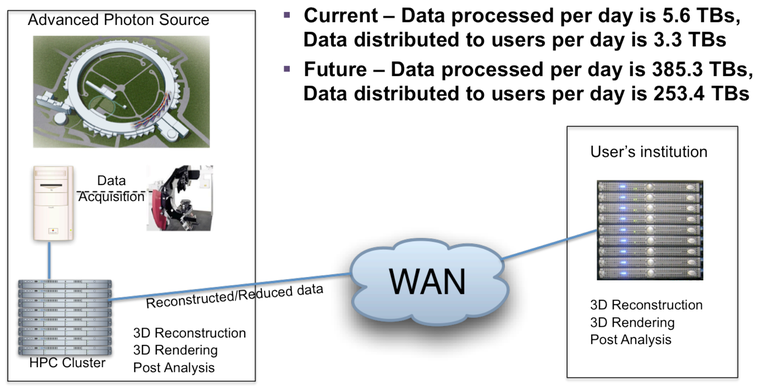

- **Current – Data processed per day is 5.6 TBs, Data distributed to users per day is 3.3 TBs**
- **Future – Data processed per day is 385.3 TBs, Data distributed to users per day is 253.4 TBs**

Experimental-time analysis is critical for enabling interactive changes to experiment parameters

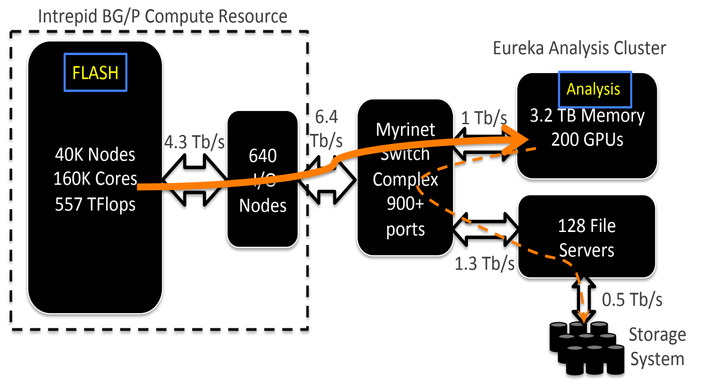

Simulation-time analysis of FLASH Astrophysics simulation data at the Argonne Leadership Computing Facility (ALCF)

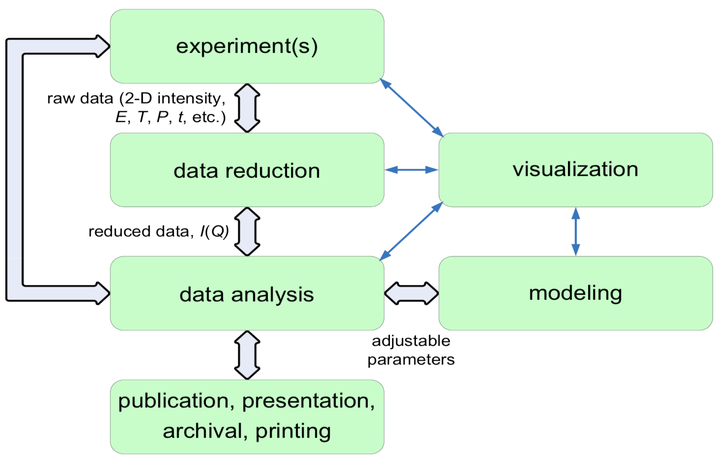

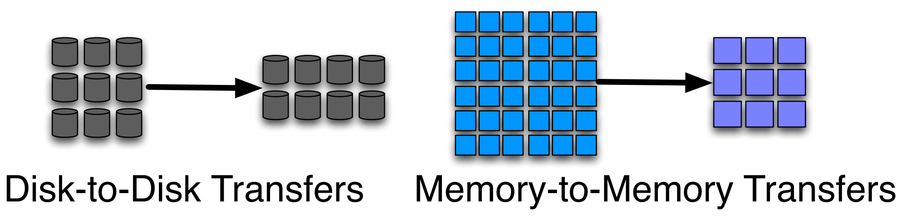

Flow characteristics

App	Type of Flow	# of Flows	BW	Latency	Burstiness	Size	Protocol
Globus Online	Data	n per node	High	N	Y	Large	TCP, UDT
	Control	1 per session	Low	Y	Y	Small	TCP
APS	Data	1 per detector	High	N	Y	Large	TCP
	Control	1 per app	Low	Y	Y	Small	TCP
FLASH Simulation-time Analysis	Data	1 per core	High	N*	Y	Variable	TCP, RDMA
	Control	1 per app	Low	Y	y	Small	TCP, RDMA
ENZO Remote Viz	Data	1 per display	High	Y	N	Large	TCP, UDP
	Control	1 per app	Low	Y	Y	Small	TCP

Application data flows have diverse needs, varying in characteristics including burstiness, latency, reliability, jitter, message size and priority. Parallel m-to-n flows are becoming increasingly common. Bulk data movement is moving beyond file transfers to memory-to-memory, memory-to-disk and disk-to-memory transfers. We aim to develop a framework to capture the requirements and characteristics of various science flows.

Approach

- Design concerted flows API
- Capture Application Requirements
  - Latency, jitter and bandwidth
  - Reliability and error tolerance
  - Deadline, start time, burstiness
  - Contiguous vs non-contiguous
  - Compression technique
  - M-to-N flows
- Create data transfer kernels for representative applications

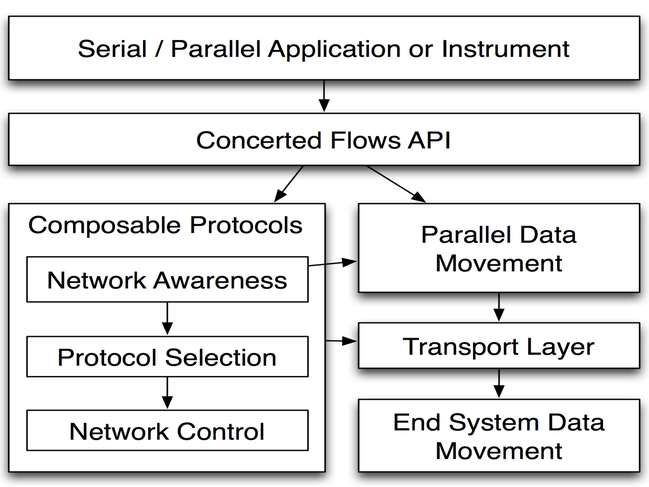

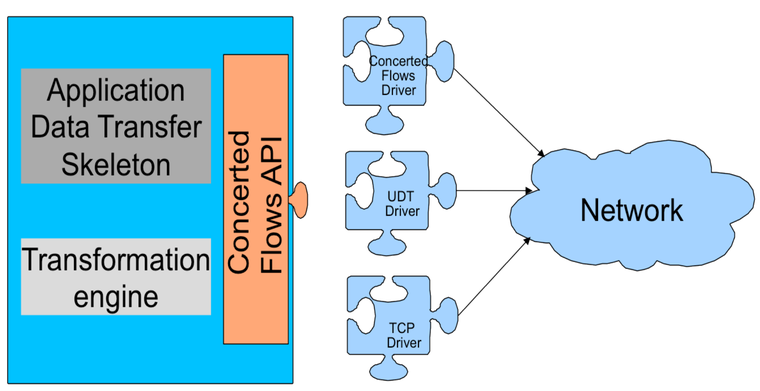

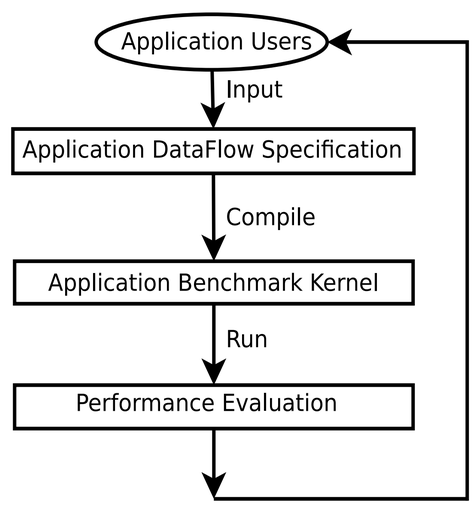

Acknowledgements: Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357,