

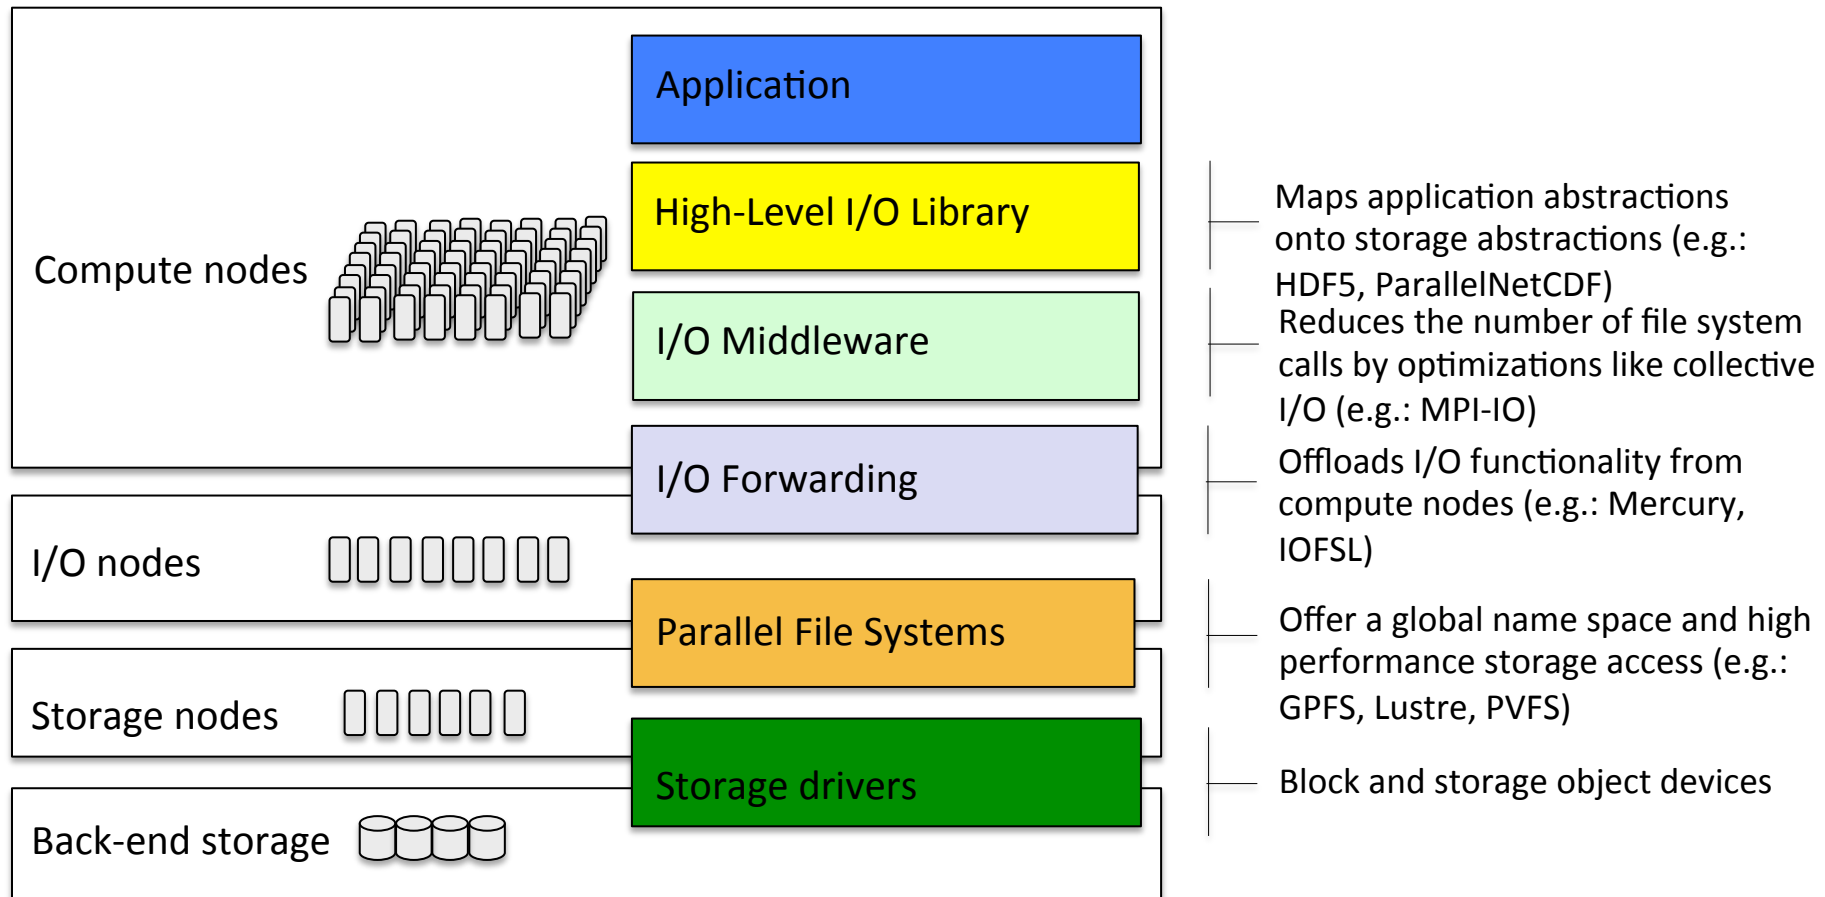
Optimizing data staging based on autotuning, coordination, and locality exploitation on large scale supercomputers

Florin Isaila

ANL & University Carlos III

Collaborators: Prasanna Balaprakash (ANL), Phil Carns (ANL), Jesus Carretero (UC3M), Francisco Duro (UC3M), Javier Garcia (UC3M), Kevin Harms (ANL), Paul Hoveland (ANL), Dries Kimpe (ANL), Rob Latham (ANL), Tom Peterka (ANL), Rob Ross (ANL), Stefan Wild (ANL)

Current problems of storage I/O software stack

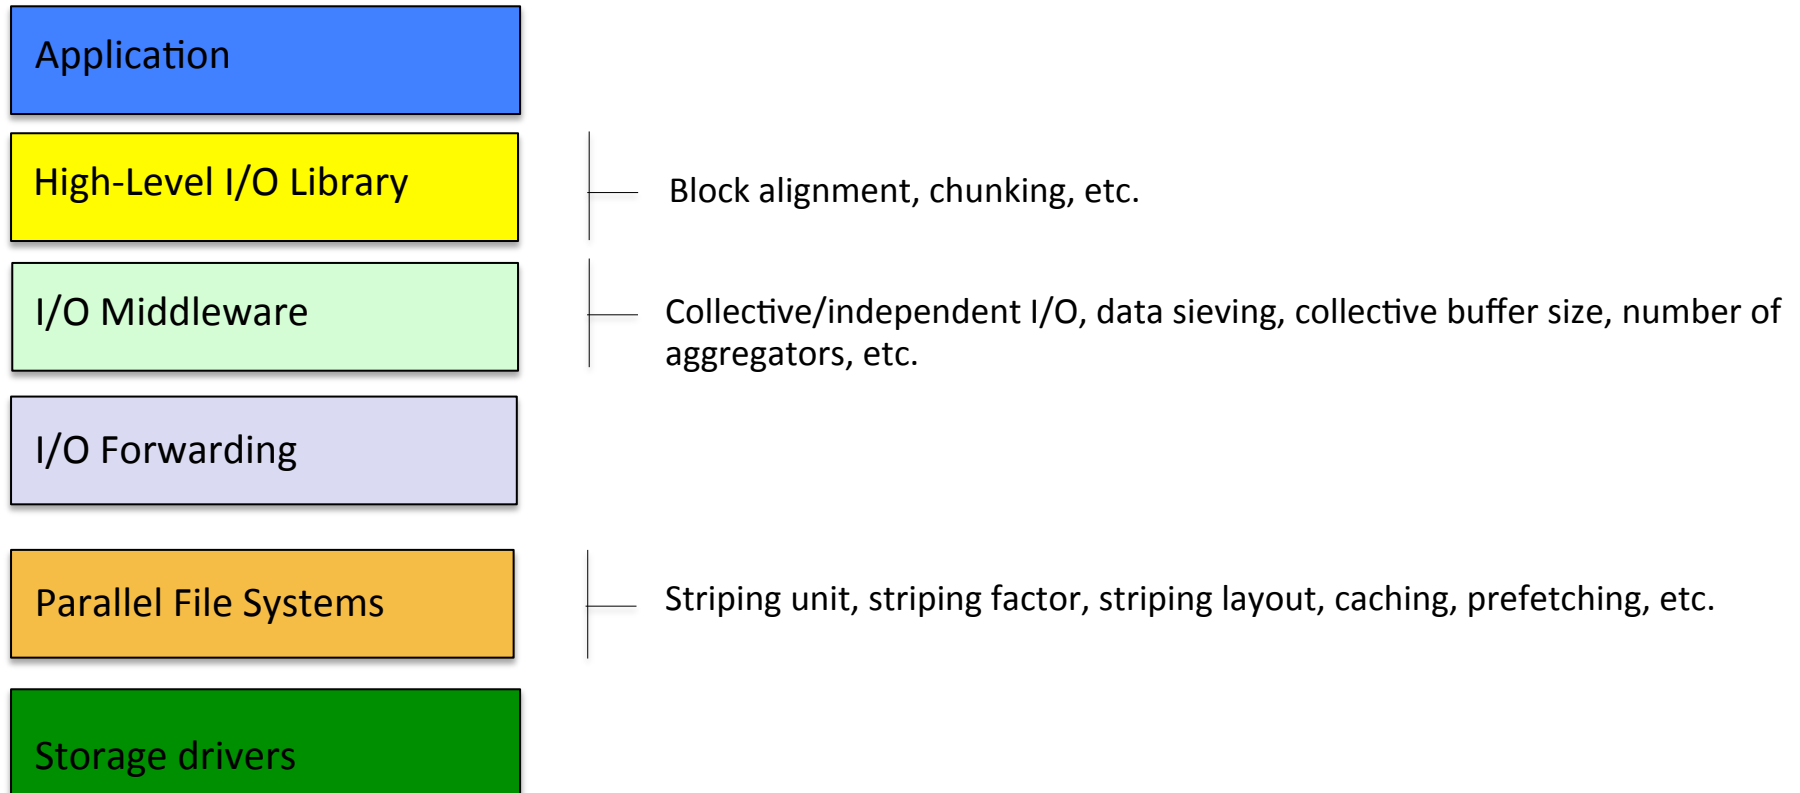


- ▶ Optimization: complex stack, deep distributed storage hierarchy
- ▶ Coordination: poor state of programmable control mechanisms are not available (e.g., for data staging, dynamic load balancing, resilience)
- ▶ Exploit data locality




- ▶ ***Optimization:*** Model-based autotuning of collective I/O
- ▶ *Coordination:* Data staging coordination
- ▶ *Exploit data locality:* Improving the scalability and performance of the Swift workflow language by leveraging data locality through Hercules

- ▶ Huge parameter space of the storage I/O software stack
- ▶ Domain knowledge is increasingly harder: software and hardware complexity



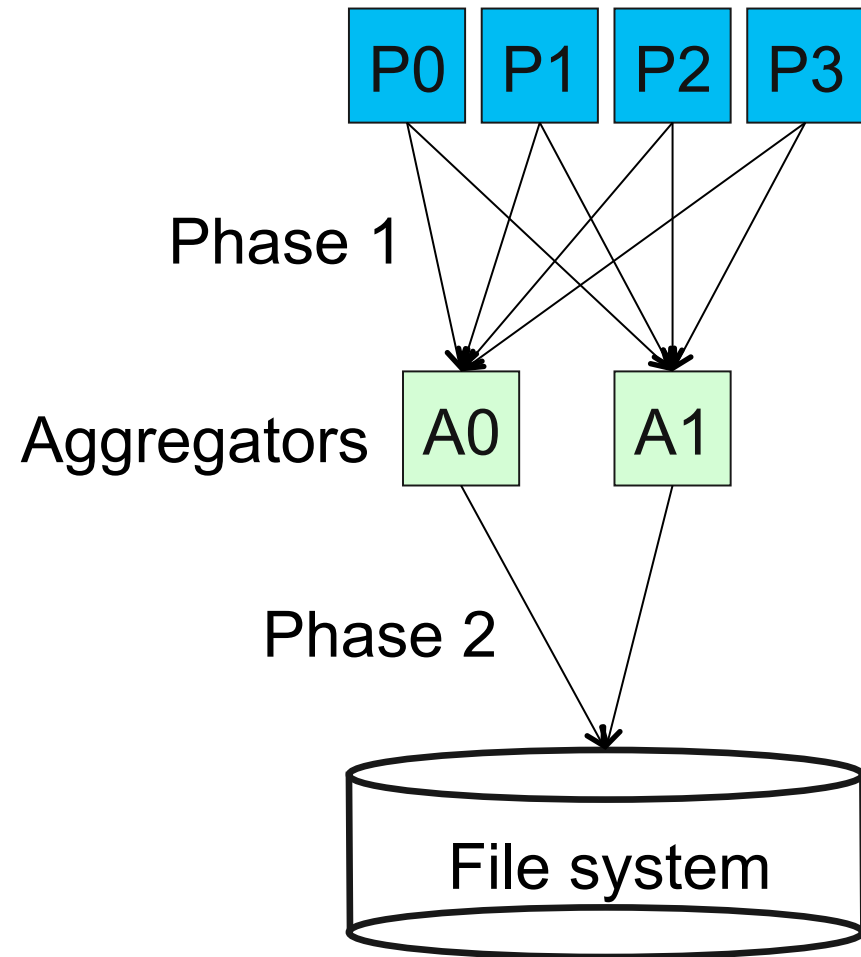
- ▶ **Model based tuning**
 - ▶ Analytical: Extensive domain knowledge required: software stack, architectural characteristics
 - ▶ Machine learning [Kumar2013, Yu2012]
- ▶ **Search-based tuning**
 - ▶ Genetic algorithms [Bezhad2013]
 - ▶ Simulated annealing [Chen2000]
- ▶ **Hybrid [Bezhad2014]**



This work

- Model-based tuning of two-phase-I/O, the most popular collective I/O implementation from ROMIO
- Combination of analytical and machine learning models
- IEEE Cluster 2015 paper

- ▶ N processes collectively write or read to a file
- ▶ Two-phase I/O write
 - ▶ Computation and communication for mapping writes to the file domain
 - ▶ Communication for sending data to aggregators
 - ▶ Storage I/O for storing the data to the file system



4 parameters

n : number of nodes

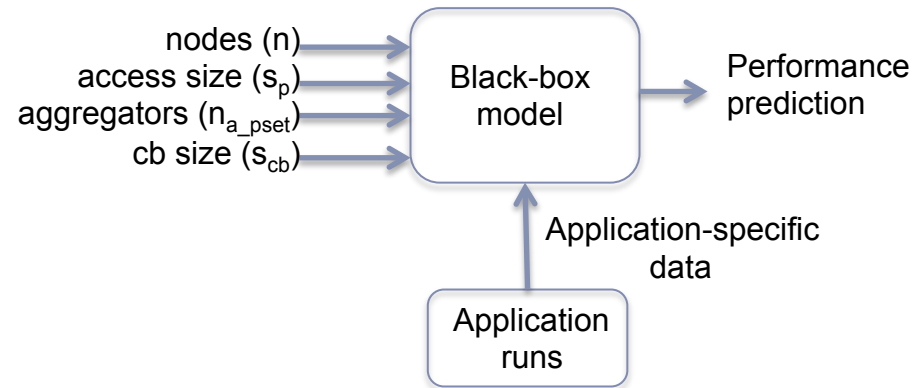
s : access size

n_a : number of aggregators

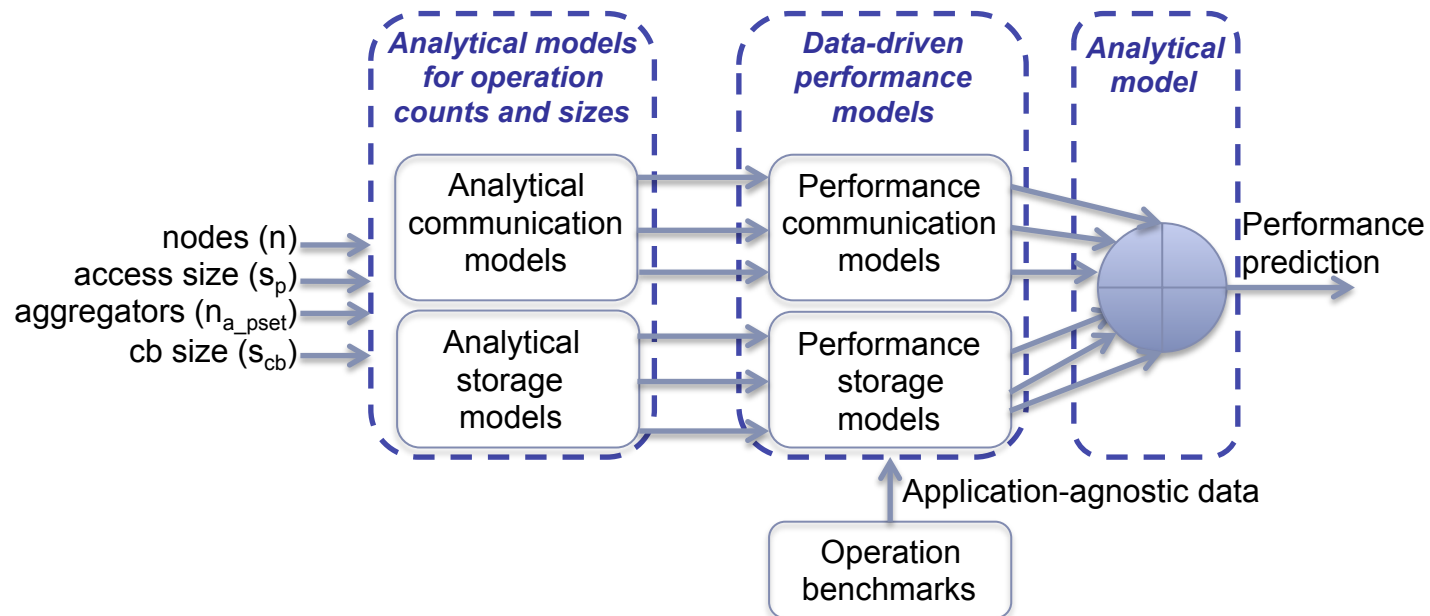
s_{cb} : collective buffer size

Goal: tune n_a and s_{cb}

Black box model



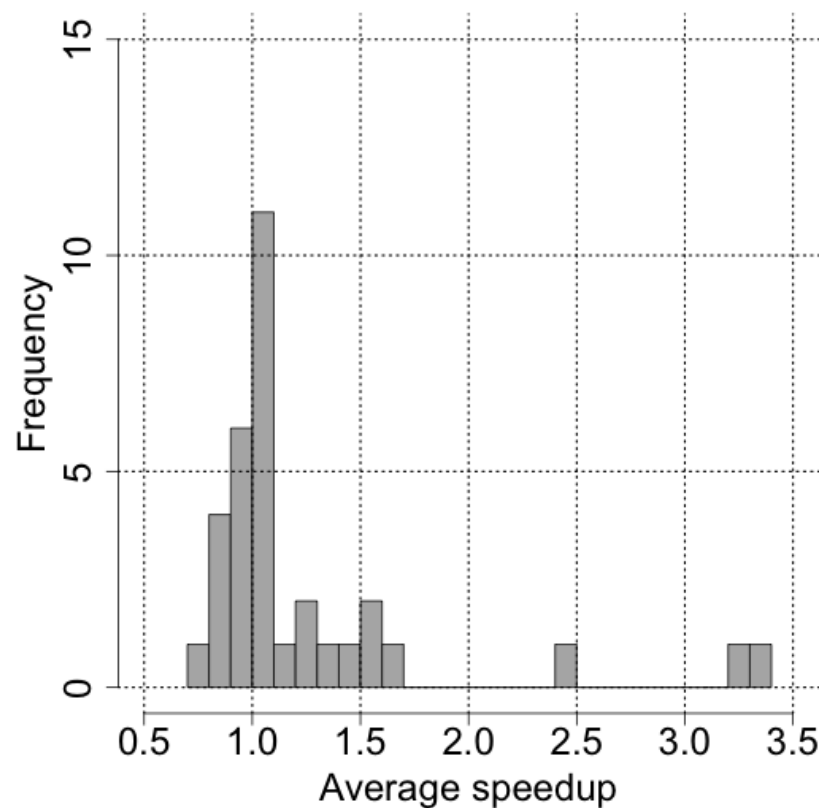
Hybrid model




- ▶ IOR benchmark: N processes concurrently write and non-overlapping region to the file system through MPI-IO
- ▶ MPICH 3.1
- ▶ Vesta Blue Gene/Q
 - ▶ 2048 compute nodes 1.6 GHz 16 cores 16GB RAM
 - ▶ 5D torus network interconnecting compute nodes
 - ▶ 1 I/O node per 32 compute nodes
 - ▶ Client (I/O node) cache: 4GB
 - ▶ GPFS 3.5: Block size: 8MB, 40 NSD SATA data drives (Max throughput: 250 MB/s)
- ▶ Benchmark for performance models: ALCF MPI benchmark

- ▶ Black box models and performance models
 - ▶ linear regression, neural networks, support vector machines, random forests, and cubist
 - ▶ Selected the model with best RMSE and R^2
- ▶ Data set:
 - ▶ Black box model: 297 points
 - ▶ Processes: 2048 (128 nodes on 16 cores), 4196, and 8392
 - ▶ Transfer sizes/core (MB): 1, 2, 4, 8, 16, 32, 64, 96, 128, 192, 256
 - ▶ Collective buffer size (MB): 8, 16 (default), 32
 - ▶ Number of aggregators per 128 nodes: 40, 136, 520 (default)
 - ▶ Performance models
 - ▶ Alltoall: : 51 points for 2,048, 4,096, and 8,192 ranks and for message sizes between 1 byte and 256KB.
 - ▶ Alltoallv: 1,044 points for distributing message sizes between 1 byte and 64 MB (in powers of 2) for subsets of 2,048, 4,096, and 8,192 ranks.
 - ▶ Allreduce: 57 points for 2,048, 4,096, and 8,192 ranks and for message sizes between 4 bytes and 1 MB.
 - ▶ POSIX: 567 points for various sizes and various subsets of 2,048, 4,096, and 8,192 ranks.

Speedup of ml; training set size=99




Speedup of hybrid over default




Comparison among various approaches

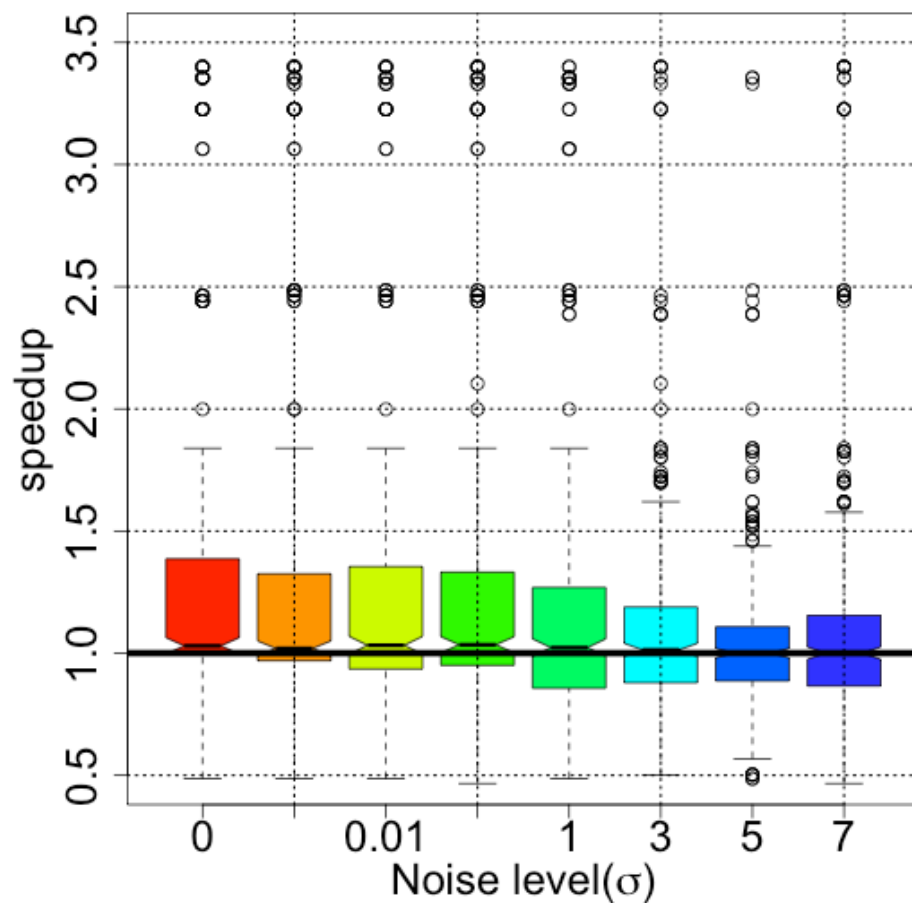
Speedup of various approaches over default



Speedup of hybrid over default



Speedup of ml over default



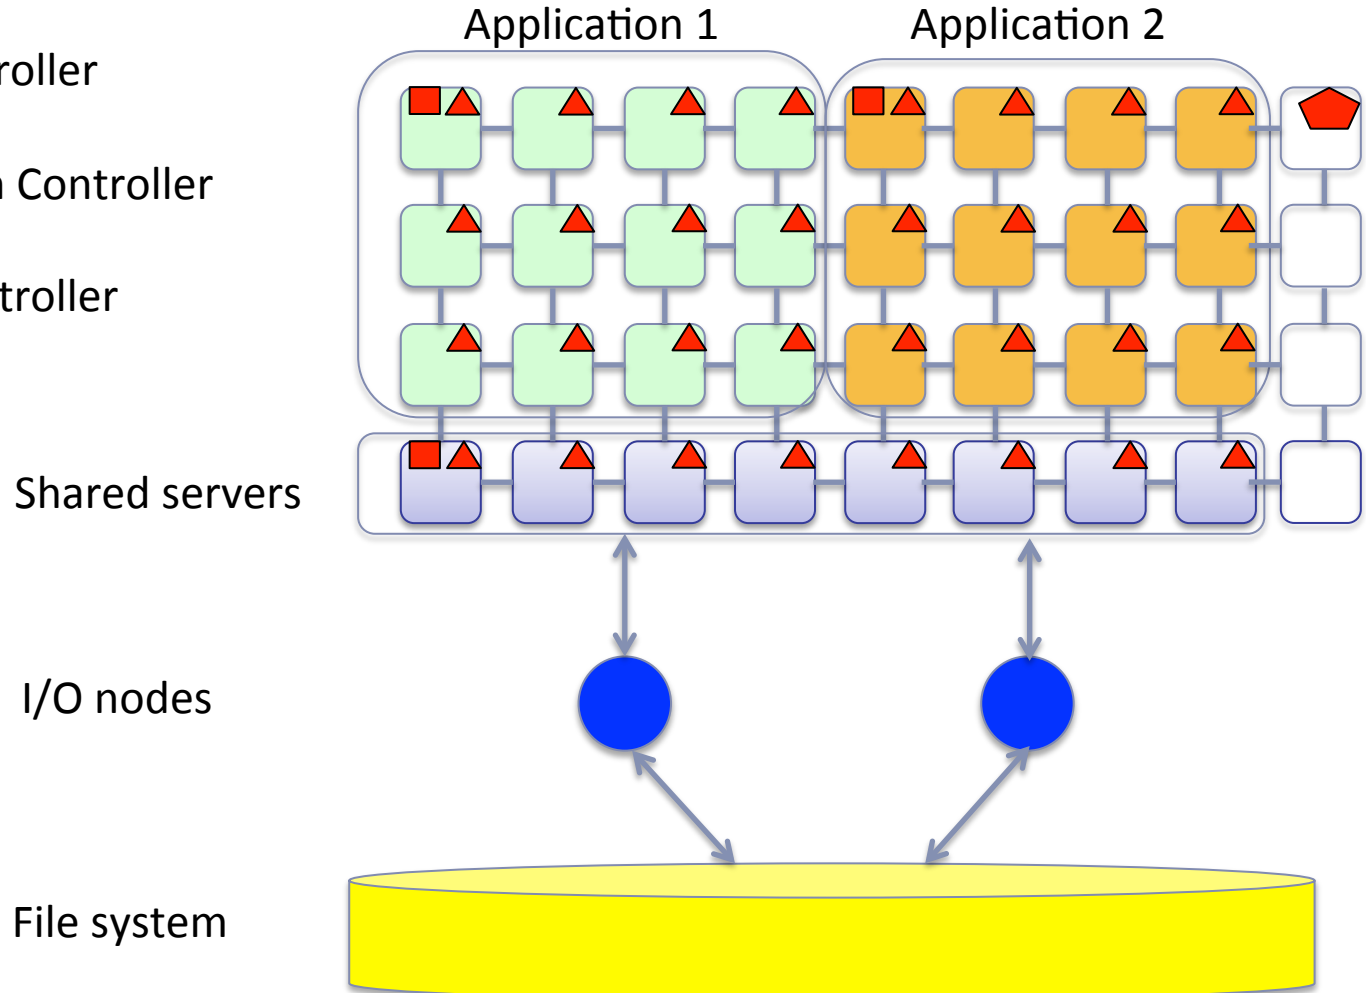
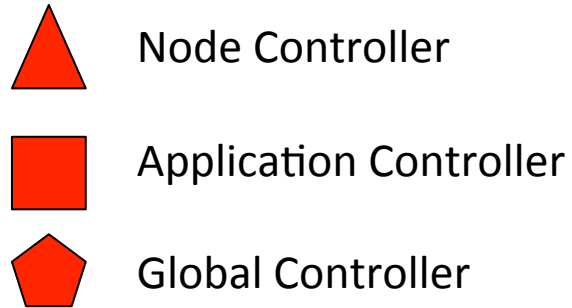


- ▶ *Optimization*: Model-based autotuning of collective I/O
- ▶ ***Coordination*: Data staging coordination**
- ▶ *Exploit data locality*: Improving the scalability and performance of the Swift workflow language by leveraging data locality through Hercules



- ▶ Concurrent parallel data flows
 - ▶ Lack of data staging coordination
 - Among applications
 - Between applications and the system
- ▶ Increasing storage hierarchy
- ▶ Lack of standards for dynamic monitoring of large scale infrastructures (e.g. load, faults)
- ▶ Coupled control and data mechanisms
- ▶ Goal: offer novel mechanisms for data staging coordination to improve
 - ▶ Load balance
 - ▶ Resilience
 - ▶ Parallel I/O scheduling

- Decouple the data and control paths
- Data-path: abstractions used to implement data access operations
 - Collective I/O
 - 2 implementation: view based I/O, list-IO (can be used as both server-based I/O and client-based I/O)
- Control path: Based on a publish/subscribe substrate (e.g. Beacon)
 - Processes can subscribe to events having certain properties
 - Associate call-back
 - Wait for an event
 - Check for the arrival of an event
- Hierarchical control
 - Global controller
 - Application controller
 - Node controller
 - All nodes participate in control





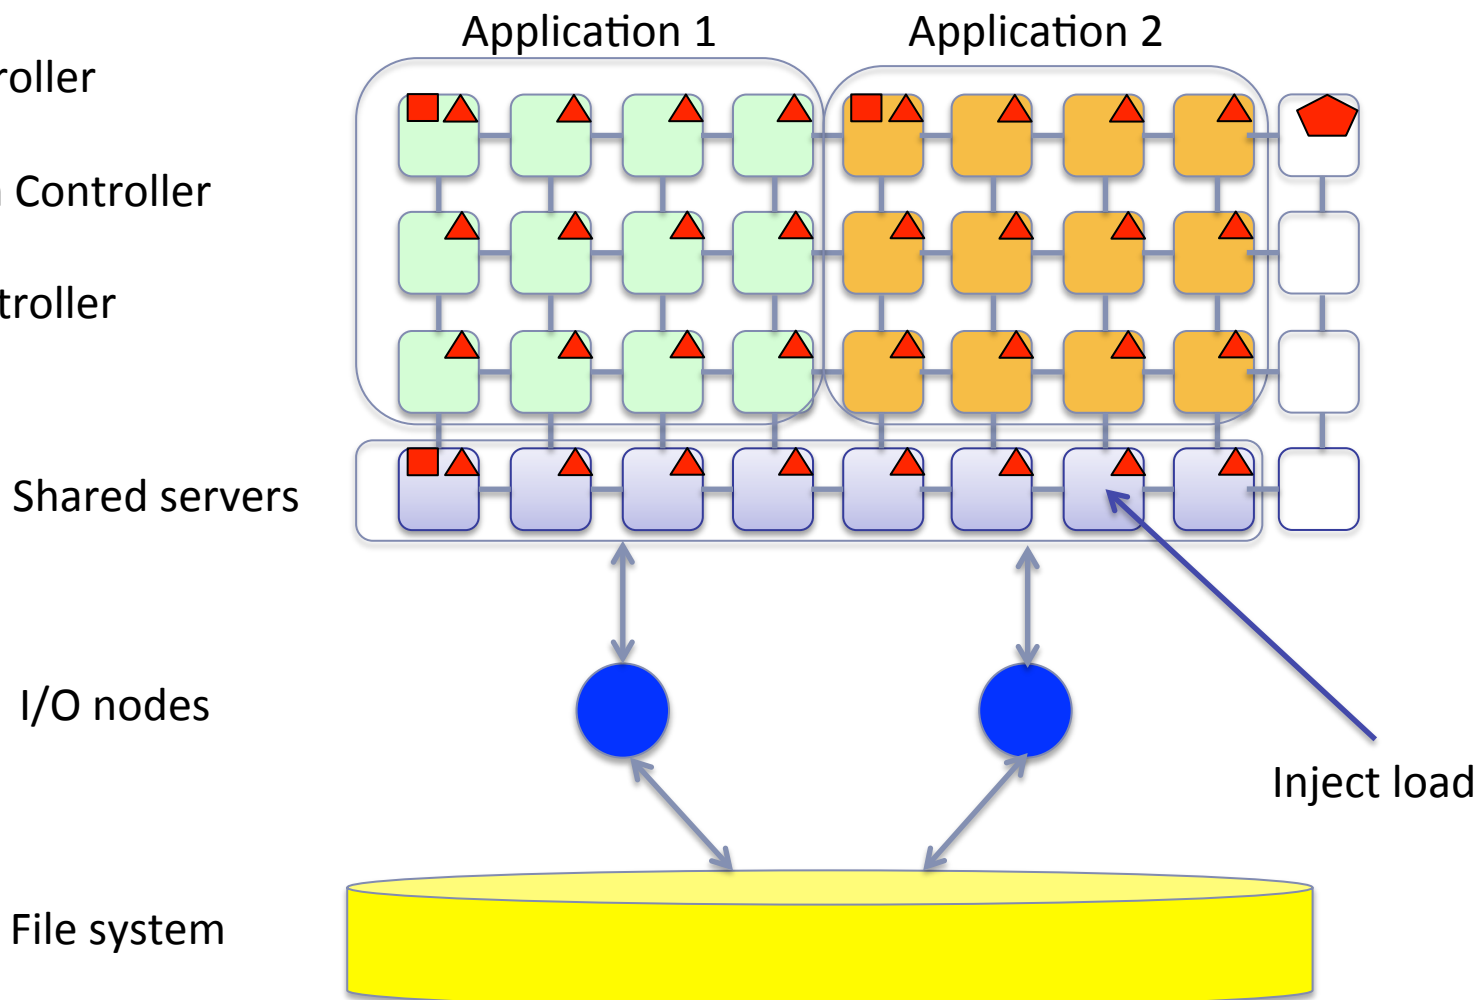
Node Controller



Application Controller

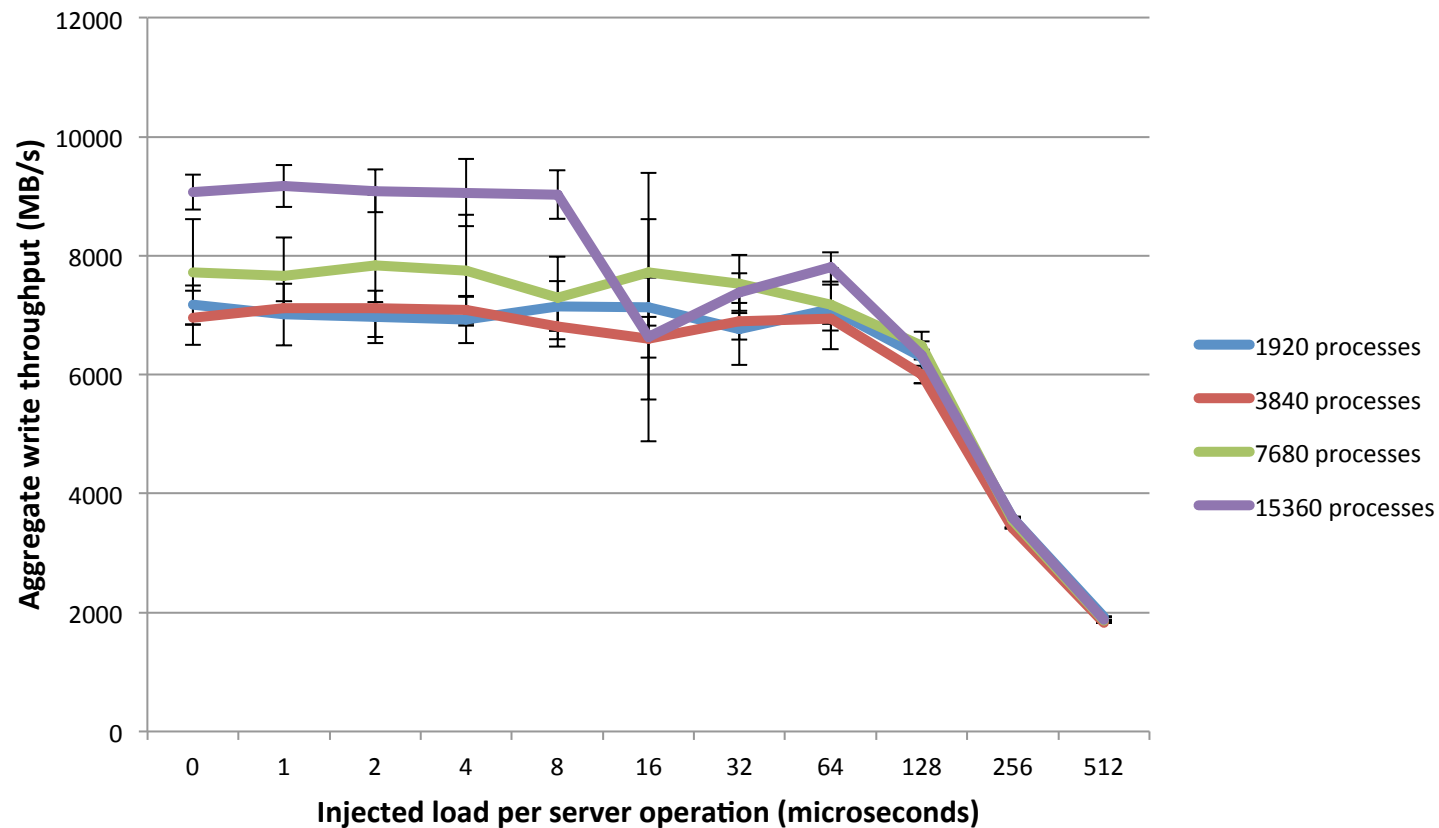


Global Controller




Load injection at 1 server (1 application)

Aggregate write throughput for injecting load on one server (one operation of 30 Gbytes)




Write time (10 operations, 3840 processes, 256/255 servers)



Detect loaded server
Reconstruct server map

New epoch with fewer servers


Write time (10 operations, 15360 processes, 1024/1023 servers)






- ▶ Several applications share
- ▶ The application controller notifies the global controller
- ▶ The global controller schedules the next application to be run
- ▶ Several policies possible
 - ▶ FCFS evaluation

Write timeline for two parallel clients with 3840 processes each - No scheduling



Write timeline for two parallel clients with 3840 processes each - FCFS scheduling



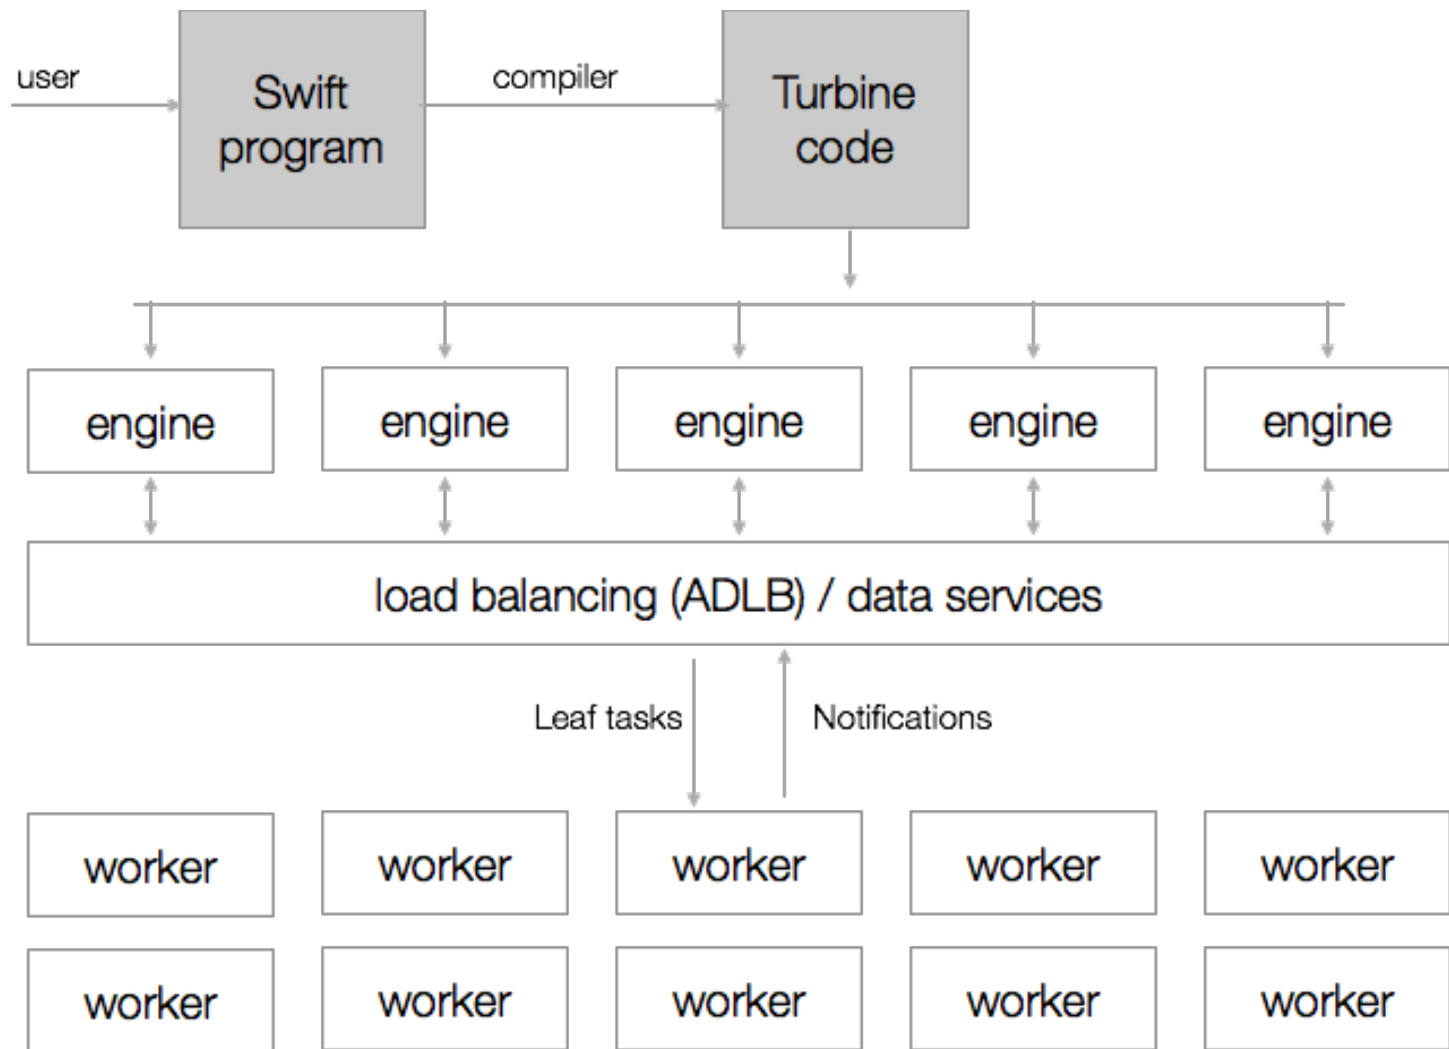


- ▶ *Optimization*: Model-based autotuning of collective I/O
- ▶ *Coordination*: Data staging coordination
- ▶ ***Exploit data locality***: Improving the scalability and performance of the Swift workflow language by leveraging data locality through Hercules

- ▶ Swift/T: Language and runtime for dataflow applications

```
(int r) myproc (int i, int j)
{
    int f = F(i);
    int g = G(j);
    r = f + g;
}
```

- ▶ F() and G() implemented in native code or external programs
- ▶ F() and G() run concurrently in different processes
- ▶ r is computed when they are both done






- ▶ Load balancer is not locality-aware
- ▶ Tasks communicate through the parallel file system (bottleneck)
- ▶ Objectives:
 - ▶ Improve the performance of inter-task communication
 - ▶ Data locality
 - ▶ Investigate the tradeoffs between data locality and load-balance in workflow execution
 - ▶ Ideal load balance, but poor locality
 - ▶ Ideal data locality, but poor load balance (not all nodes used)

- ▶ **Hercules**
 - ▶ persistent key value store based on Memcached
 - ▶ On-demand deployment of servers on application nodes
- ▶ **Data placement over the servers**
 - ▶ Consistent hashing (original Memcached)
 - ▶ Locality-aware (implemented)
 - ▶ Load-aware (under implementation)
 - ▶ Capacity aware
- ▶ **New Swift language constructs**
 - ▶ Soft location: best effort task placement
 - ▶ Hard location: enforce data locality

File-copy Strong Scalability - Aggregated Throughput*


1024 files x 256 MBytes (R+W)



Fusion Linux cluster

- 320 nodes, 2 x quad core, 36 GB RAM
- Infiniband QDR (4 GB/s) and gigabit ethernet
- GPFS: up to 2500 MB/s

MapReduce-like WC application 256 files x 256 MB - 64 GB Total execution time




- ▶ **Model-based autotuning I/O**
 - ▶ Performance predictability
 - ▶ Improve individual models
 - ▶ Noise
 - ▶ Load and noise modeling for load detection in data staging (multiple servers)
- ▶ **Data staging coordination**
 - ▶ Topology-aware server/aggregator placement - JL Colaboration with E. Jeannot, F. Tessier (INRIA), V. Vishnavath (ANL)
 - ▶ Multiple stage coordination (aggregation – burst buffer – file system)
 - ▶ Load prediction based on Omnisc'IO (Mathieu Dorrier – ANL)
 - ▶ Adaptive buffering in parallel applications workflows (Decaf project)
 - ▶ Adopt Global Information Bus from Argo and Hobbes (Beacon, Exposé)
 - ▶ Need for sub-second monitoring and notification
- ▶ **Exploit locality in workflows**
 - ▶ Load-aware placement
 - ▶ Tradeoff locality – load balance
- ▶ **New applications? New architectures? New coordination scenarios?**

Thank you


Conclusions - autotuning

- ▶ Automatic parameter configuration
 - ▶ Machine learning and hybrid models approaches outperform the default values in most cases
 - ▶ Hybrid models higher robustness to noise than pure machine learning
 - ▶ Hybrid model do not require application reruns
- ▶ Factors that limit efficiency of the I/O stack optimization
 - ▶ POSIX consistency semantics: File locking
 - ▶ File system noise
 - ▶ The lack of information about the state of storage hierarchy (e.g. cached versus non-cached)
 - ▶ Performance predictability needs to improve


Aggregate write throughput for 1920 processes



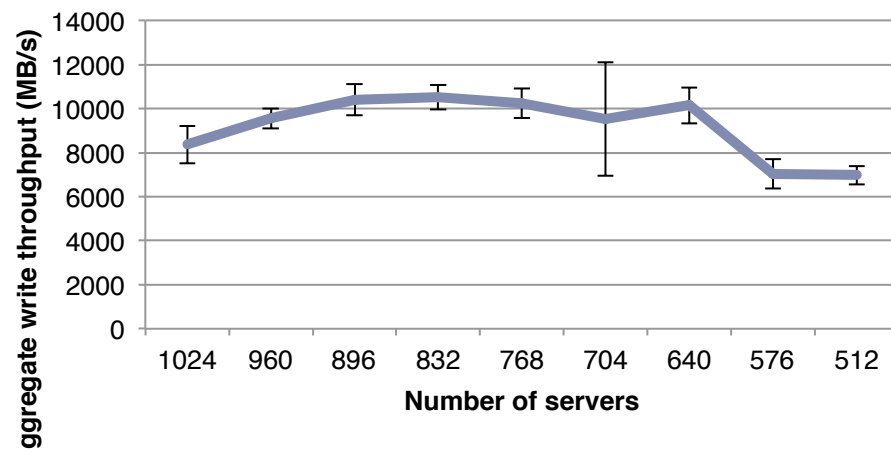
Aggregate write throughput for 3840 processes



Aggregate write throughput for 7980 processes



Aggregate write throughput for 15360 processes





- ▶ Assumes the availability of a load detection mechanism
- ▶ One application process detects a loaded server
- ▶ Notifies the application controller
- ▶ Application controller informs all node controllers and ask them to prepare to start a new epoch with less servers
- ▶ Node controller
 - ▶ Decides the last operations to be executed from the current epoch
 - ▶ Suspends all operation from the future epoch
 - ▶ Updates the server map
 - ▶ Notifies the application controller
 - ▶ Application controller ask all nodes to start a new epoch
- ▶ Each node controller resumes the suspended operations if any



- ▶ Data staging coordination
- ▶ Separation of data and control
- ▶ Hierarchical controlling
- ▶ Significant benefits
 - ▶ Load/Fault aware sever-scale down
 - ▶ Parallel I/O scheduling
- ▶ Scalable load and fault monitoring is required



- ▶ Integration Swift/T - Hercules
- ▶ Substantially improves the throughput over shared file systems
- ▶ I/O performance scales up with the number of application nodes
- ▶ Exploit data locality in workflows
- ▶ Less sensitive to file system noise and contention

