

# FlexIO: Location Flexible Execution of In Situ Analytics for Large Scale Scientific Applications

Fang Zheng<sup>1</sup>, Hongbo Zou<sup>1</sup>, Greg Eisenhauer<sup>1</sup>, Karsten Schwan<sup>1</sup>, Matthew Wolf<sup>1,2</sup>, Jay Dayal<sup>1</sup>, Tuan-Anh Nguyen<sup>1</sup>, Jianting Cao<sup>1</sup>, Hasan Abbasi<sup>2</sup>, Scott Klasky<sup>2</sup>, Norbert Podhorszki<sup>2</sup>, Hongfeng Yu<sup>3</sup>

<sup>1</sup>CERCS, Georgia Tech <sup>2</sup>Oak Ridge National Laboratory <sup>3</sup>Sandia National Laboratory

## Motivation

### “Big Data” Sciences:

- Scientific Simulations may generate TBs of data per run
- Data is the key to scientific insights
- “Time to Insight” is critical for scientists’ productivity

### I/O Bottleneck:

- I/O capacity not well provisioned on High End machines
- Severe I/O overheads for simulation and analytics at scale


### In Situ Data Analytics: Addressing the I/O Bottleneck

- Analyze data at simulation time
- Online data inspection, manipulation, and processing
- Reduce disk I/O activities
  - Save time and power
- Deliver timely insights
  - Improve productivity

## Challenges

### Need for location flexibility:

- There are multiple options to place in situ analytics
- Placement can impact performance and cost
- No single best placement exists




### Requirements for I/O middleware:

- Support diverse placement options
- Require minimal effort in changing analytics placement
- High performance data movement between simulation and analytics
- Provide performance information for tuning placement

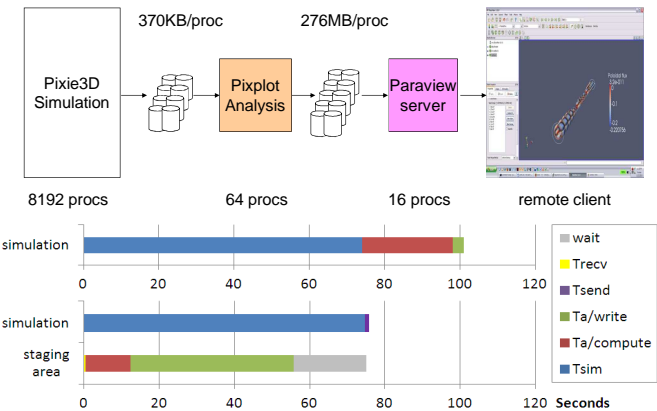
## FlexIO Middleware

### FlexIO middleware for coupling parallel simulations with in-situ analytics:

- Flexibility in ‘where’ and ‘when’ analytics codes are run
- Driven by ADIOS IO interface
- Alter placements without changes to application/analysis codes
- Mobile “Data Conditioning Plug-ins” for dynamic data reduction
- Fast, scalable intra- and inter-node data movement
- Online monitoring of computation and data movement



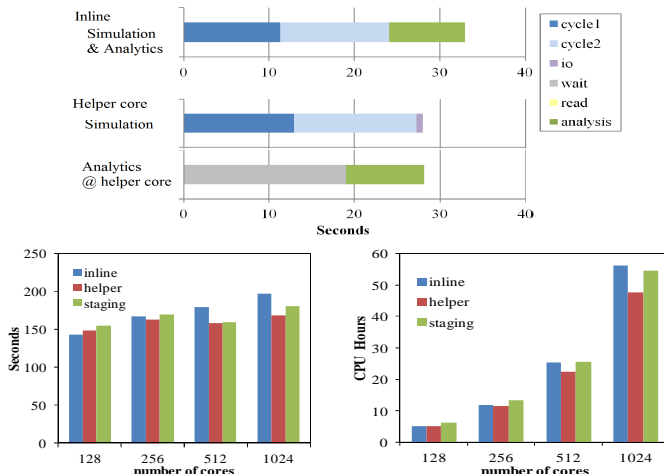
## Case Study I: Pixie3D



Using 0.78% additional nodes offloading Pixplot and I/O to staging area increases performance by 33% in comparison to inline placement

## Case Study II: GTS

GTS fusion simulation code is coupled with a series of analytics (particle distribution function, range query, 2D histogram)



Placing analysis on helper cores in compute nodes achieves the best performance & cost due to spare cycles available there

## Summary

### Conclusions:

- Placement of in situ analytics has significant impact on end-to-end performance and cost
- FlexIO middleware supports diverse placement options for in situ analytics under uniform high level I/O APIs
- Significant performance & cost improvements achieved through flexible placement for real applications

### Ongoing and future work:

- Porting FlexIO middleware to Cray XK6 platform
- Dynamic resource provisioning and automatic placement

## Contact

### Email:

Fang Zheng (fzheng@cc.gatech.edu)  
 Karsten Schwan (schwan@cc.gatech.edu)  
 Matthew Wolf (mwolf@cc.gatech.edu)  
 Scott Klasky (klasky@ornl.gov)

### Web:

<https://research.cc.gatech.edu/sdmatcercs/>  
<http://www.olcf.ornl.gov/center-projects/adios/>

