



Award #: 1663887

## Collaborative Research: SI2-SSI : EVOLVE: Enhancing the Open MPI Software for Next Generation Architectures and Applications

PIs: Edgar Gabriel <sup>1</sup>, George Bosilca <sup>2</sup>Co-PIs: Thomas Herault <sup>2</sup>, Aurelien Bouteiller <sup>2</sup><sup>1</sup> University of Houston, <sup>2</sup> The University of Tennessee, Knoxville


### Open MPI

- Widely used implementation of the Message Passing Interface Specification
- Jointly maintained by a consortium of academia, industry and national laboratories
- Over 60,000 downloads per year
- World-wide user community

### Goals of the EVOLVE project

- Extend Open MPI to support new features of the MPI specification
  - Hybrid programming models
  - Fault tolerance
- Support new architectures
- Improve scalability
- Enhance Parallel I/O performance and functionality

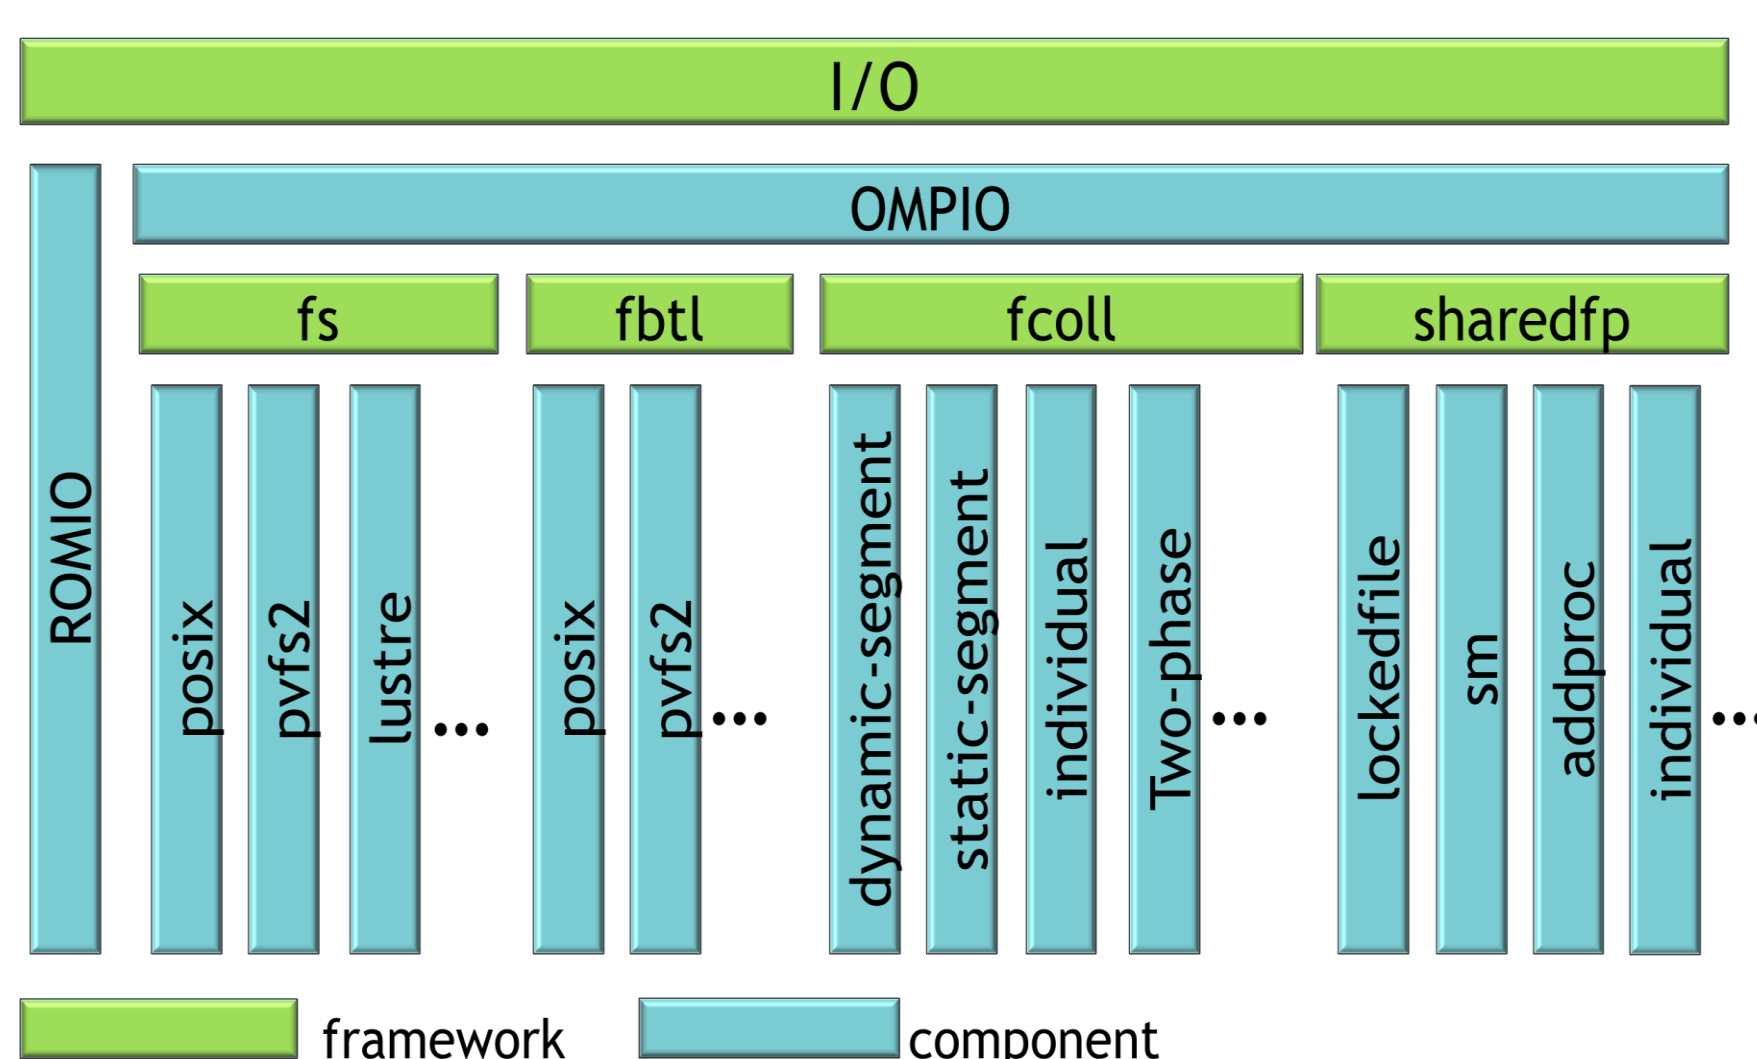
### Open MPI Download statistics



Number of visits to the download page of Open MPI for the v3.0, v3.1, and v4.0 releases in the from 05/14/2018 until 05/14/2019.


### OMPIO

- Parallel I/O components of Open MPI
  - Separates I/O functionality into frameworks
  - Dynamic runtime selection of components
- Integrated with Open MPI derived data type engine
- Integrated with Open MPI progress engine for asynchronous I/O operations



### Overlapping multiple cycles in collective I/O

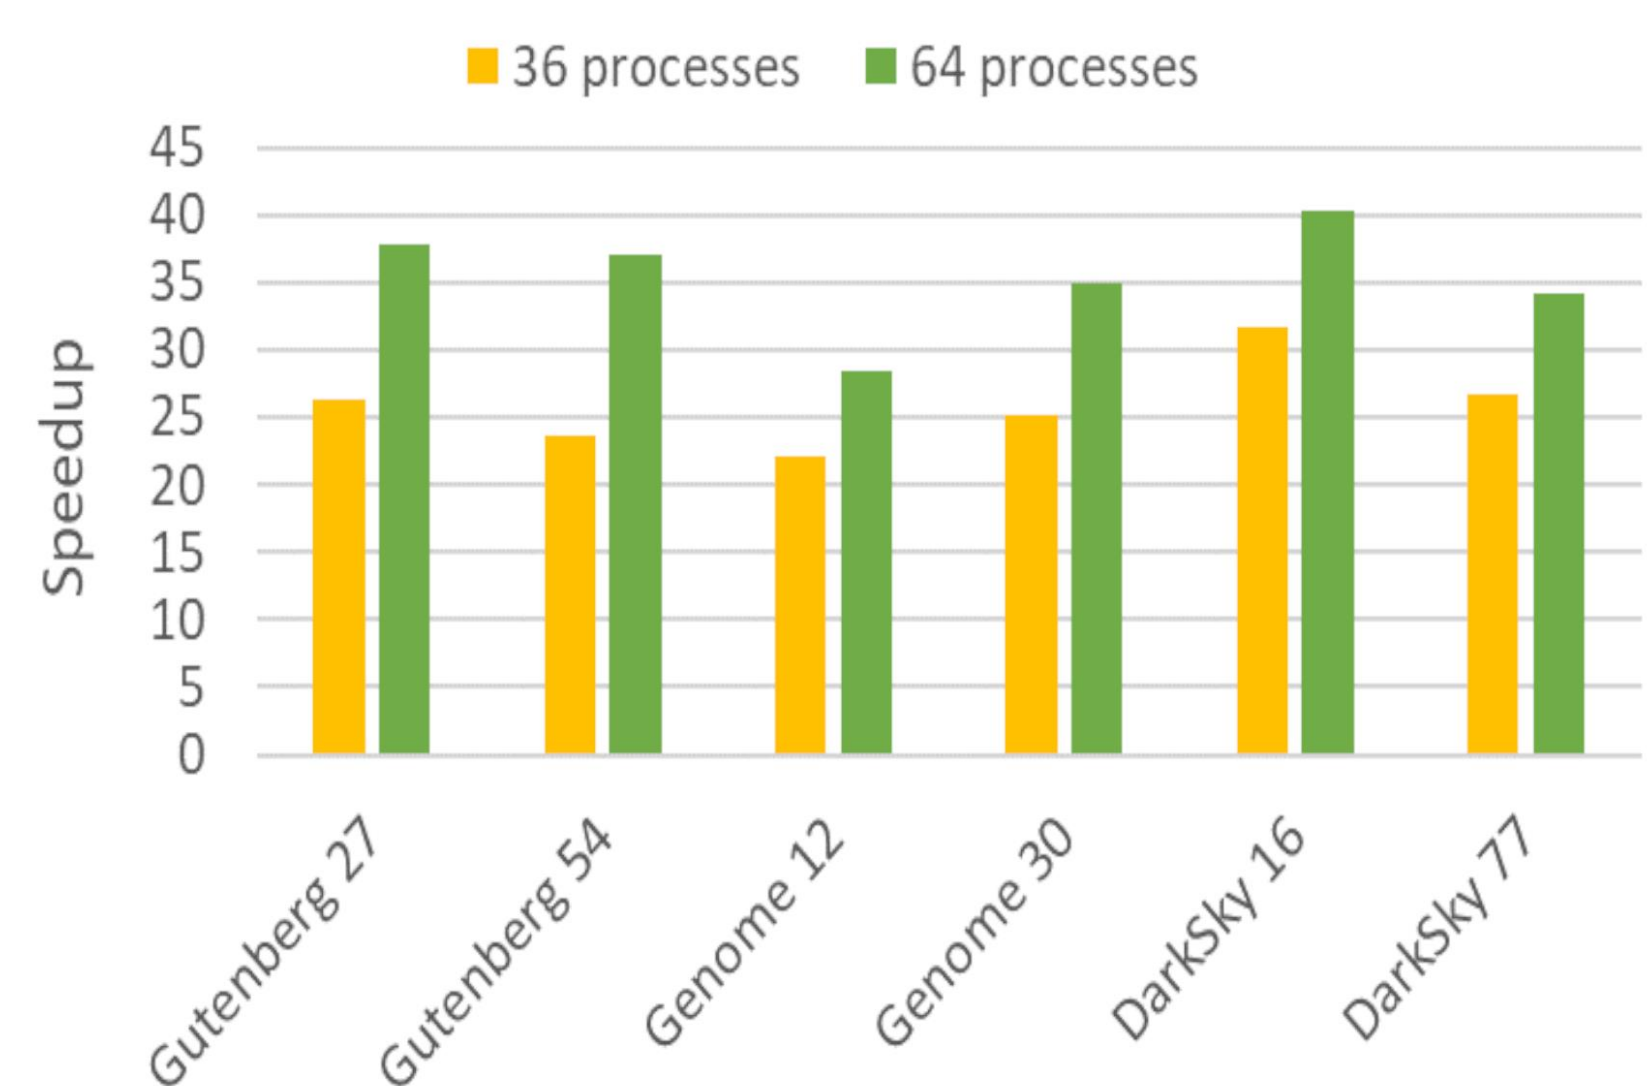
- Performance improvement for collective I/O through internal pipelining
- Multiple algorithms evaluated based on different overlap strategies
- Multiple data transfer primitives (two-sided, one-sided) explored



Average performance improvement of collective write operations for various overlap algorithms

### Compression in Parallel I/O

- Extend collective I/O to support reading and writing compressed data files
- Prototype implementations developed for Snappy compression algorithm



Speedup of a parallel compression tool using collective I/O over a sequential compression tool

Contact: Edgar Gabriel [egabriel@uh.edu](mailto:egabriel@uh.edu)

[1] M. Chaarawi, E. Gabriel, R. Keller, R. Graham, G. Bosilca and J. Dongarra, 'OMPIO: A Modular Software Architecture for MPI I/O, EuroMPI 2011.

[2] Shweta Jha and Edgar Gabriel, 'Performance Models for Communication in Collective I/O Operations', CCGRID Workshops 2017.