SFIO progress on Swiss-Tx

SCS meeting on scalable distributed file system to Linux October 31, 2000

Emin Gabrielyan

EPFL, Computer Science Dept.
Peripheral Systems Lab.
Emin.Gabrielyan@epfl.ch

- Last (September) optimizations of SFIO read/ write operations
- SFIO on top of MPICH, performance on T1
- SFIO on top of FCI, performance on T1
- Multiple SFIO files access performance on top of FCI of T1. Influence to the performance, stipulated by T1's topology.
- Conclusion
- Future work

Optimisation of SFIO read/write operations for consecusive single block requests

- Control information transfer optimisation for SFIO read and write operations. Control data transmitted in buffered asynchronous mode.

- Optimisation of transmission of data together with control information. Fragmented data together with controlling arrays are grouped into single datatype.

- Asynchronous optimisation of data reception at Compute Node for SFIO read operation.

- Modifications of the SFIO library architecture.

- A graphical demonstration of data flow of optimized SFIO read operation.

Performance improvement by optimisation of control data transmission

Consecutive accesses to a SFIO 10MB/500B/8IO file on Swiss-T0/Hub

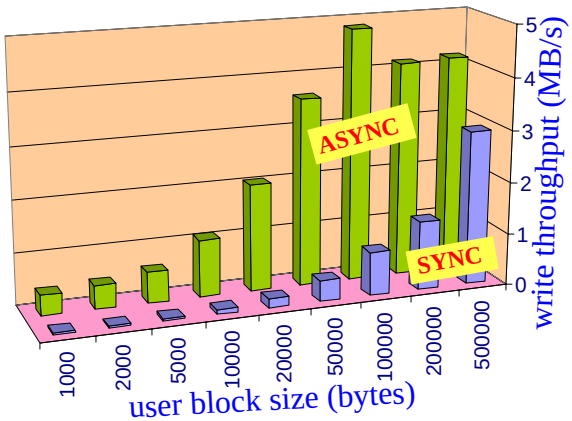

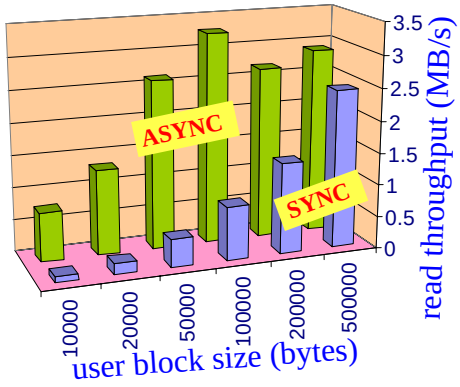

Performance improvement by grouping control data with access data

Consecutive accesses to a SFIO 10MB/500B/8IO file on Swiss-T0/Hub

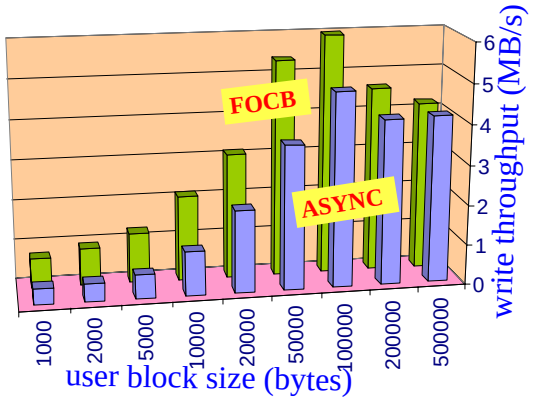

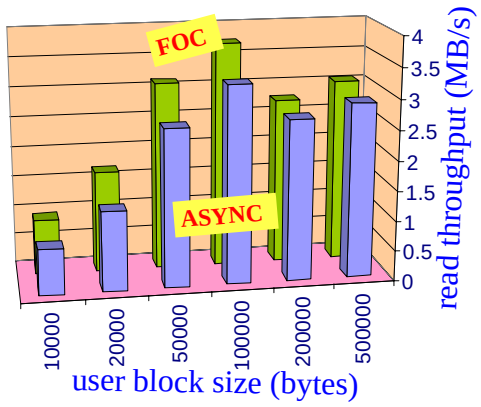

Performance improvement by asynchronous data reception optimization

Consecutive accesses to a SFIO 10MB/500B/6IO file on Swiss-T1Baby/TNet

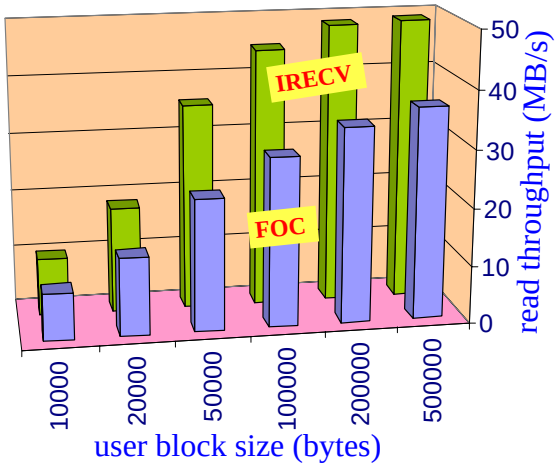

-

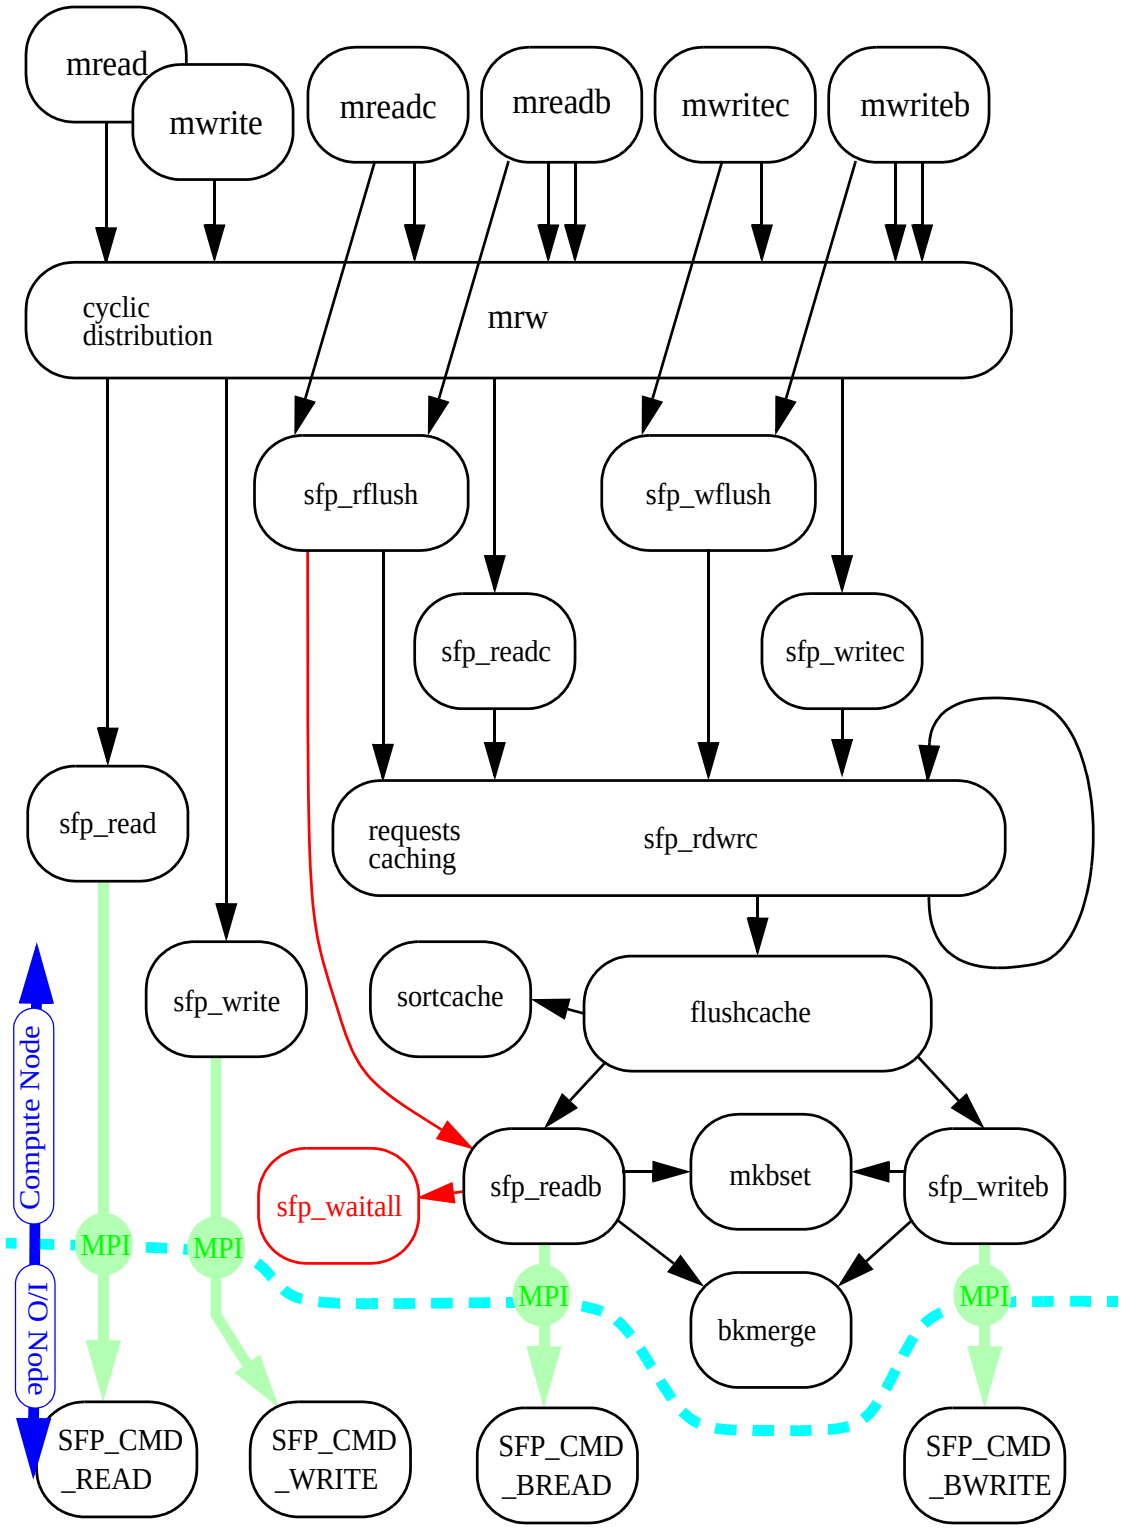
mread
mwrite
mreadc
mreadb
mwritec
mwriteb
cyclic
distribution
mrw
sfp_rflush
sfp_wflush
sfp_readc
sfp_writec
sfp_read
requests
caching
sfp_rdwrc
sfp_write
sortcache
flushcache
Compute Node
I/O Node
sfp_waitall
sfp_readb
mkbset
sfp_writeb
MPI
MPI
MPI
MPI
bkmerge
SFP_CMD
_READ
SFP_CMD
_WRITE
SFP_CMD
_BREAD
SFP_CMD
_BWRITE

-

SFIO Blocking Read, blocking implementation

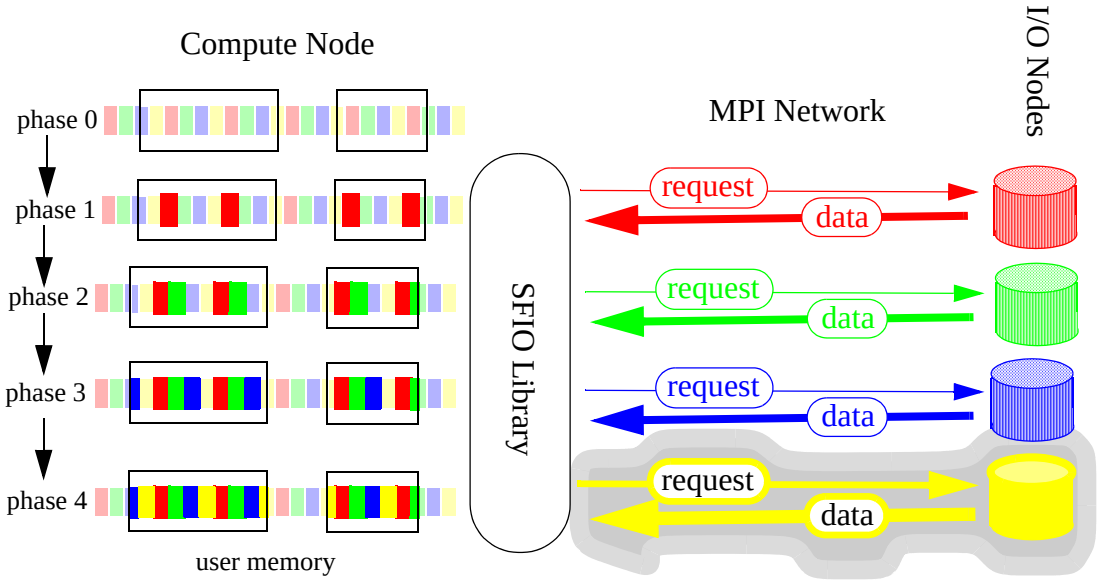

SFIO Blocking Read, non-blocking implementation

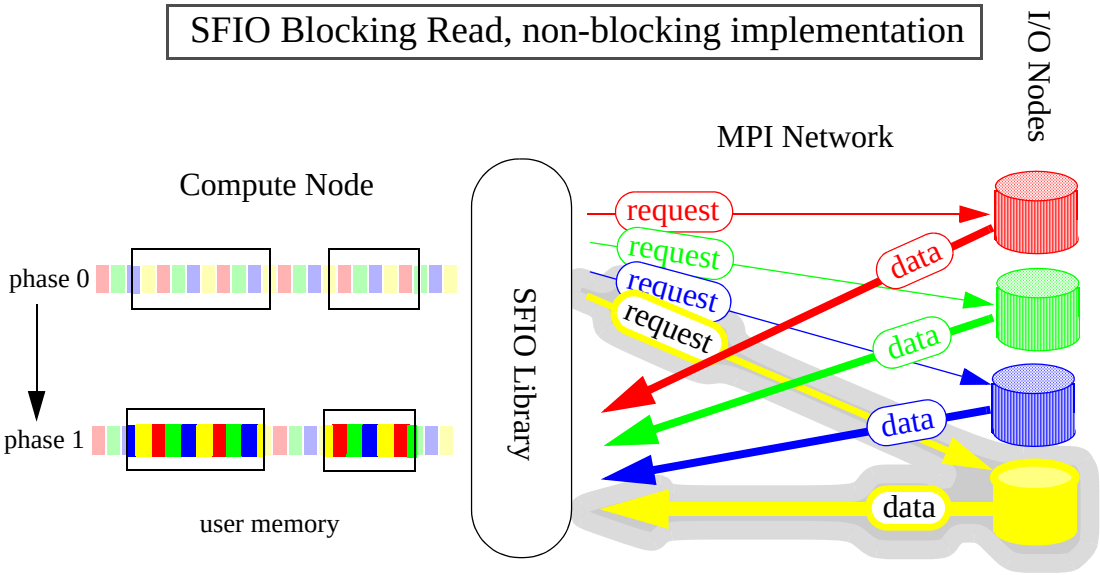

-

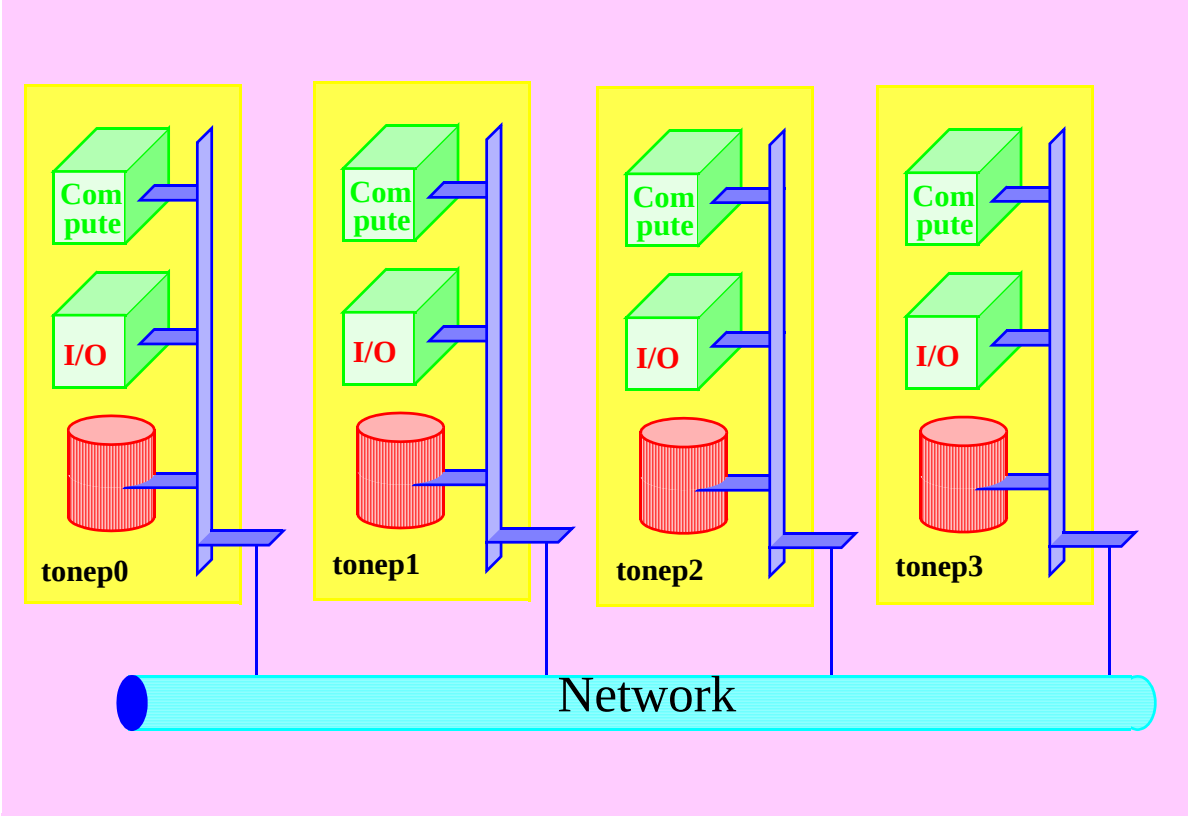

- SFIO All-to-All concurent write access from all compute nodes to all I/O nodes

- Global File size is 2000MByte

- Stripe unit size is 200Byte only

SFIO all-to-all I/O performance on Swiss-T1's Fast Ethernet and Tnet

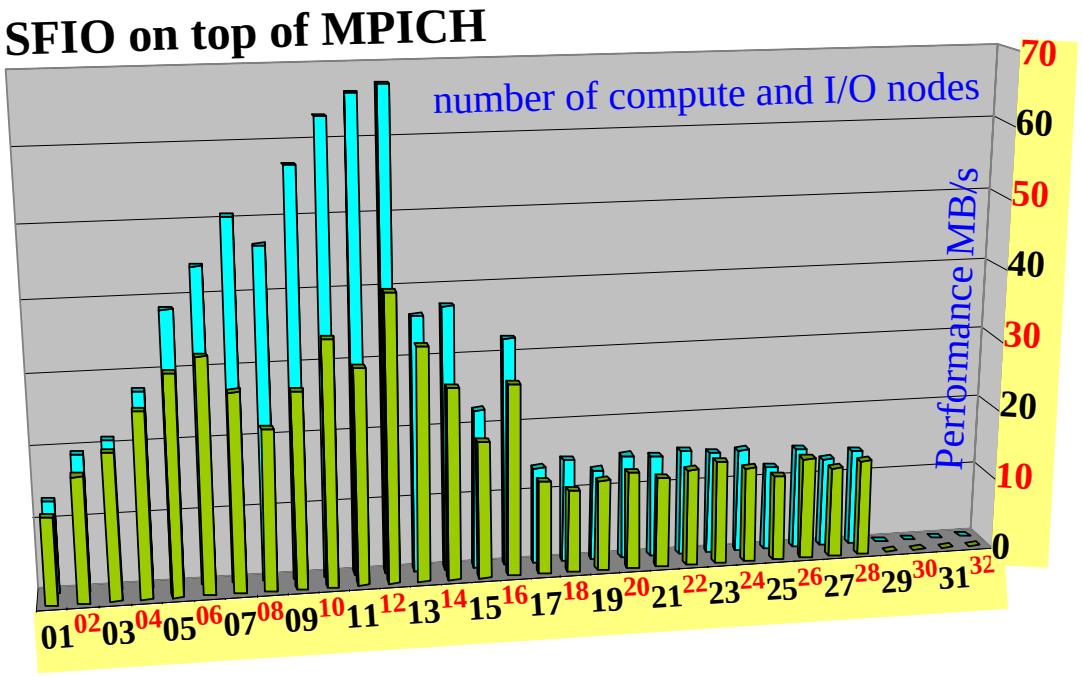

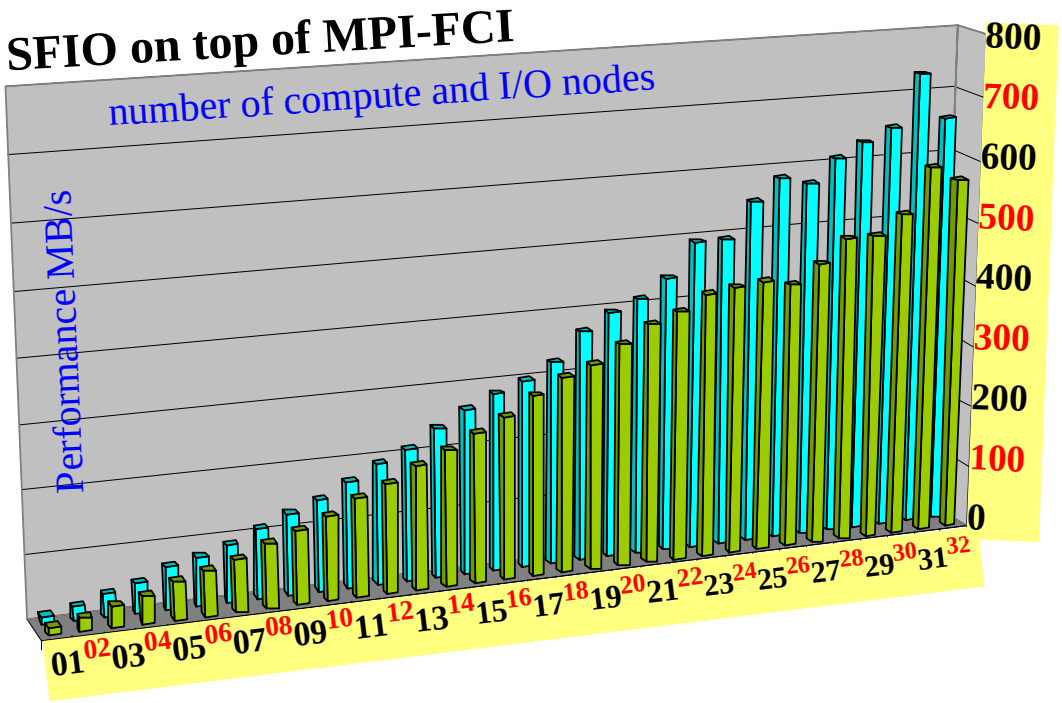

- Superlinear speedup of SFIO/FCI due to augmentation of cache effect when increasing the number of I/O nodes.

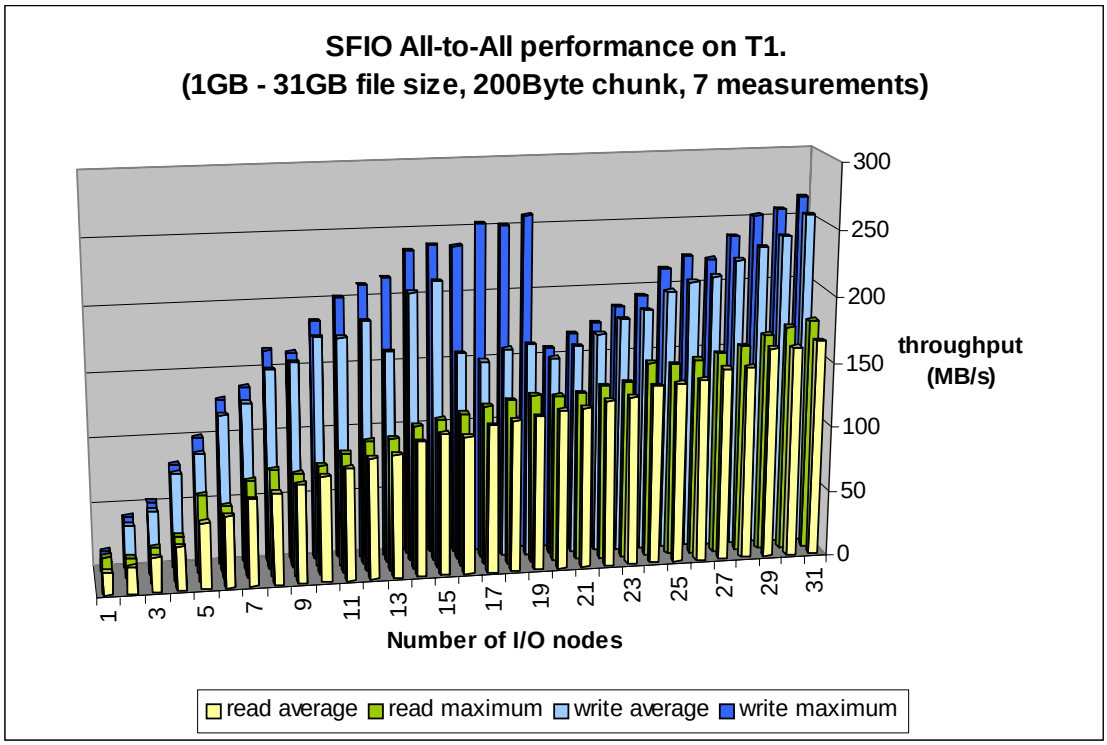

- To avoid the cache effect the total size of SFIO files is increasing when the number of I/O nodes grows.

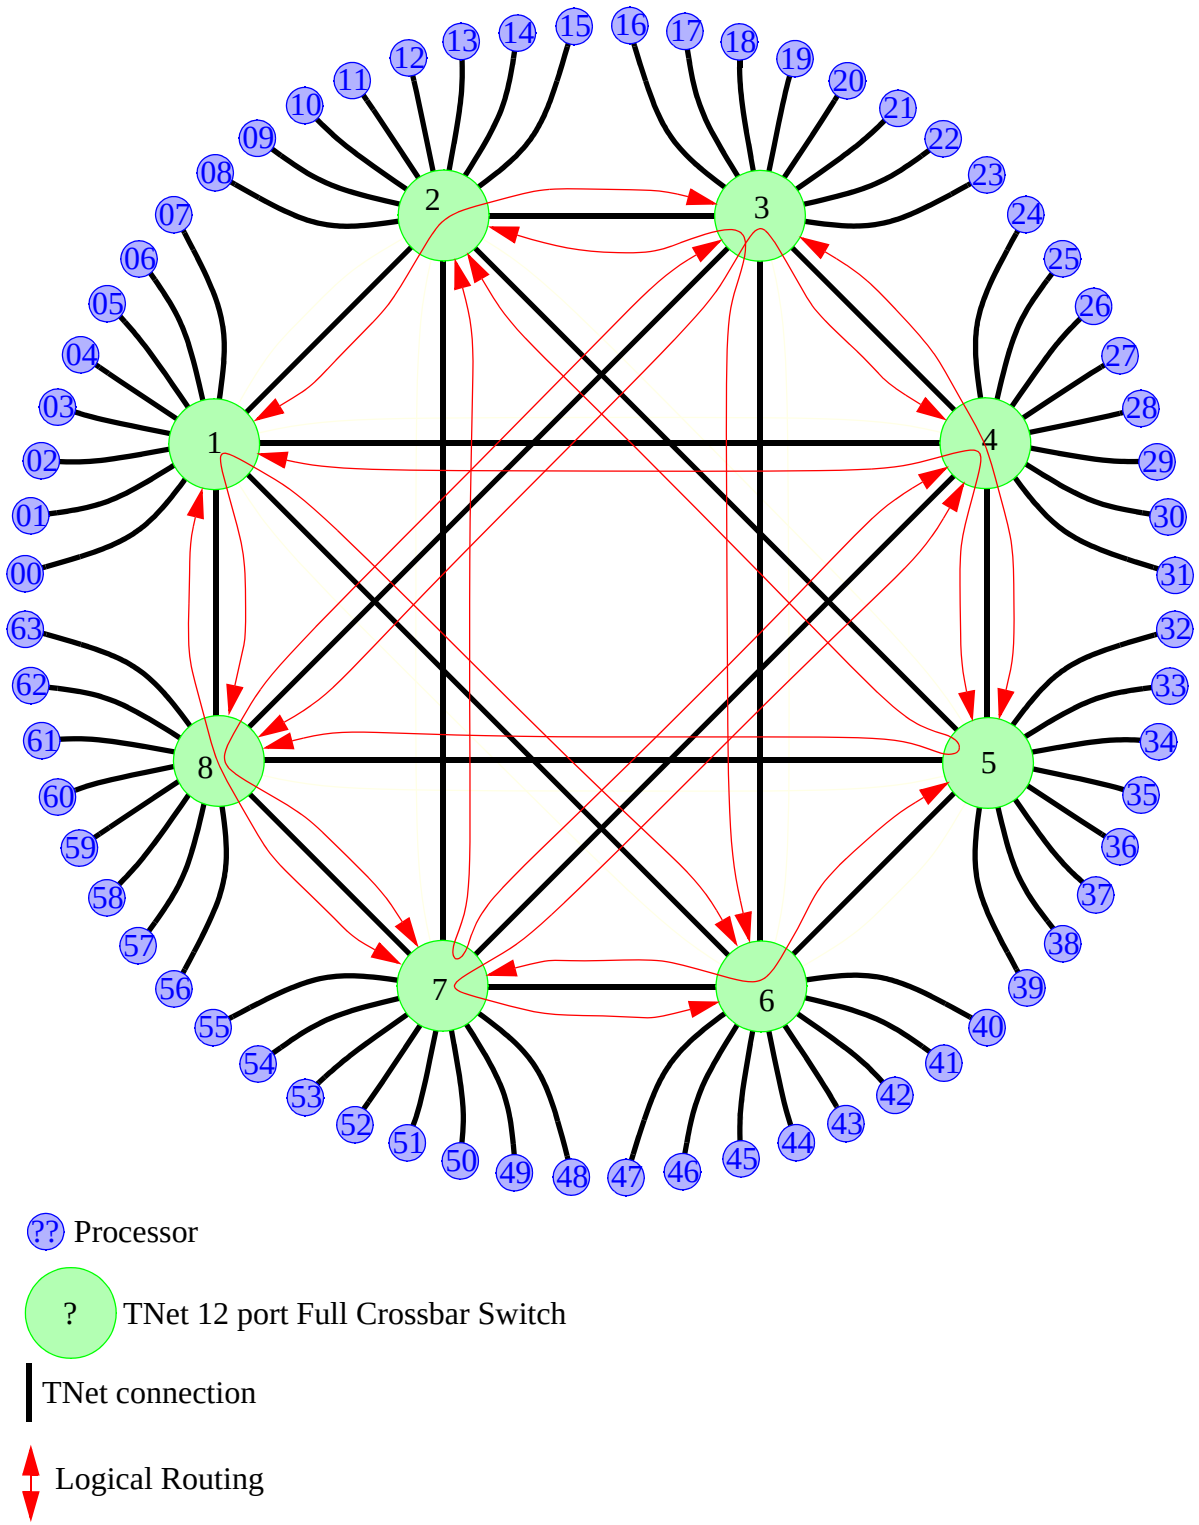

Swiss-T1 TNet interconnection and routing topology

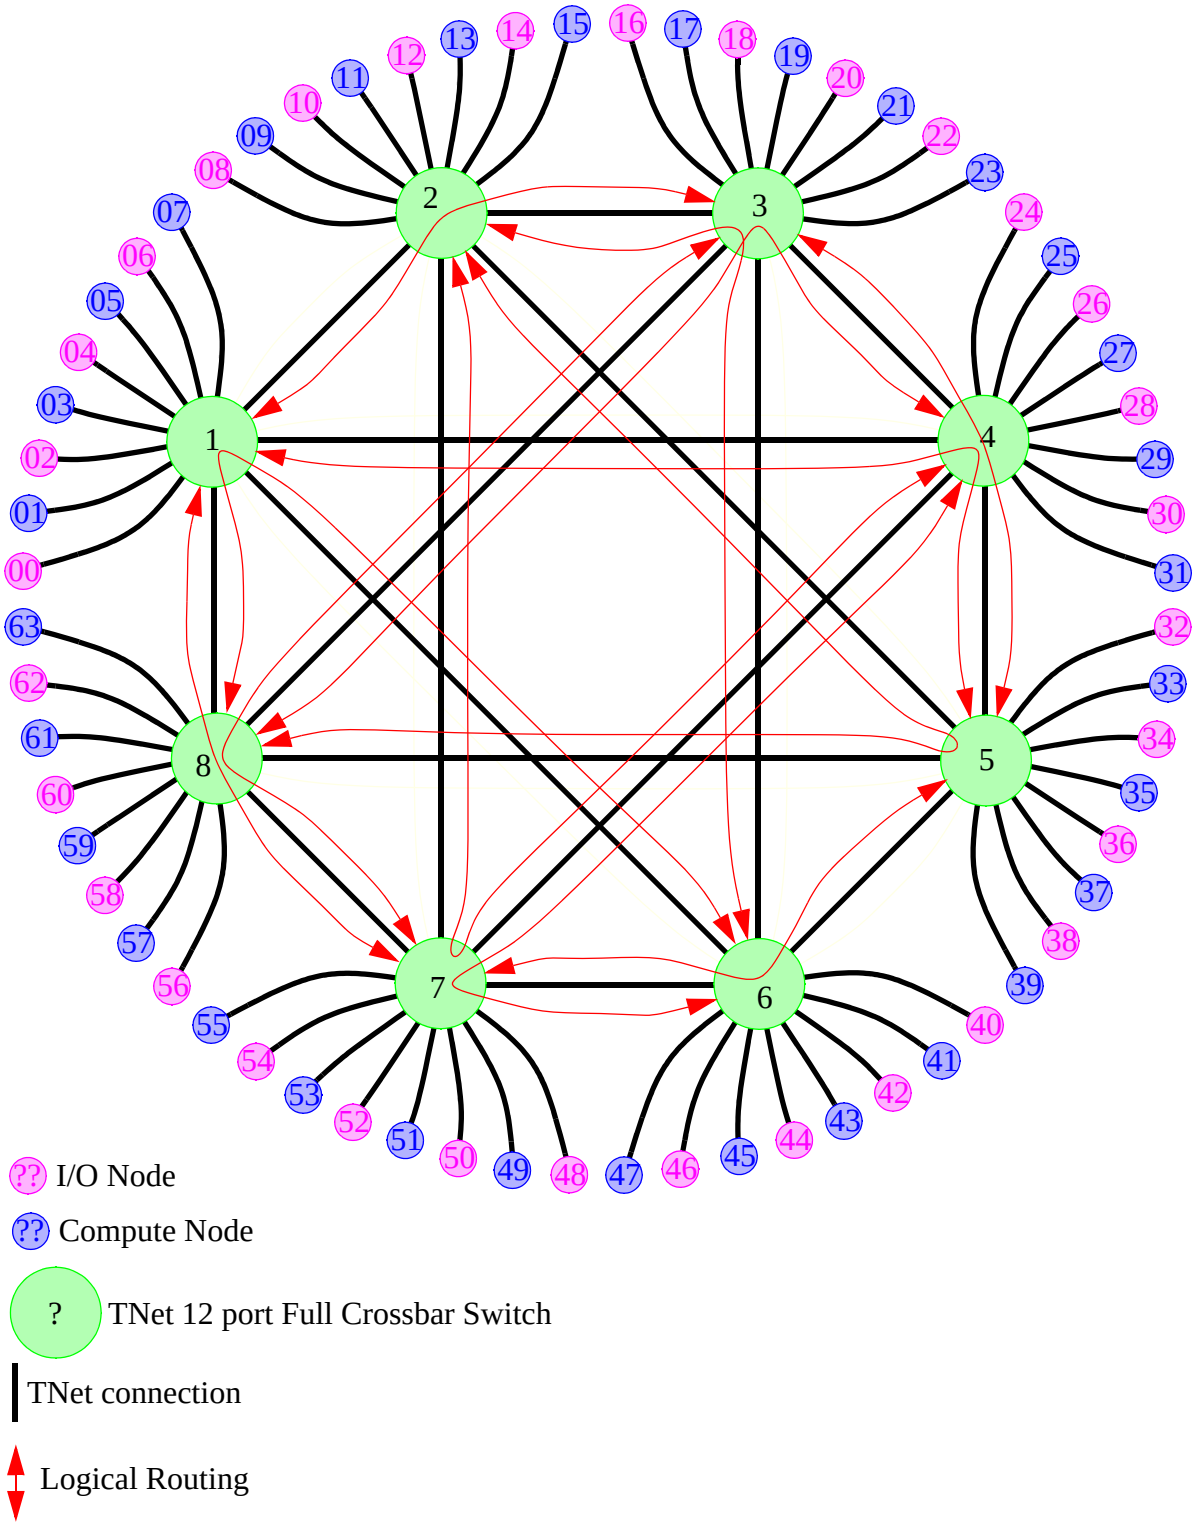

Swiss-T1 SFIO over TNet topology

Conclusion

- SFIO is portable, highly scalable, and ready for the distribution.

Future work

- SFIO performance benchmarking on the large supercomputer of Sandia National Laboratory.
- Adapt from T0 to T1 the modifications of MPICH/ADIO which provide a routing of a sub-set of MPI-I/O operations to the SFIO.
- Performance measurements of MPI-I/O inter-faced to SFIO through MPICH/ADIO.
- Possibly, creation of a portable MPI-I/O interface library to SFIO.
- Asynchronous implementation of blocking write operation. Pipelining on the I/O node.