pallas

Competence in High Performance Computing

Portable MPI Tools at Work - Cracking Performance Problems

Werner Krotz-Vogel, Pallas GmbH

<krotz@pallas.com>

42nd CUG Conference, May 02000

Pallas GmbH
Hermülheimer Straße 10
D-50321 Brühl, Germany

info@pallas.com
www.pallas.com

Pallas Tools Software

State-of-the-art program development tools ...

... in detail:

- Vampir-2.5 (online-Demo), Vampirtrace-2.0, Dimemas

... briefly:

- Etnus TotalView 4.0 Multi-process Debugger
- PGI 3.1 x86 Compilers, Cluster Development Kit (CDK)
- KAP/Pro Toolset 3.8, OpenMP
- KAI C++ 3.4, ISO standard
- FORESYS - Fortran Restructuring Tool

... free open source:

- PMB - Pallas MPI Benchmark Suite (incl. “effective Bandwith”)

PGI Cluster Development Kit (CDK)

Compilers & Tools ...

- PGI 3.1 x86 compilers , C, C++, F77, F90, HPF, pgrof, pgdbg
- SMP/OpenMP support for C, C++, F77, F90

... plus convenient add-on´s:

- parallel ScaLAPACK
- optimized BLAS, LAPACK
- MPI/mpich
- PVM
- PBS - Portable Batch System
- Tutorial, examples
- Cluster management utilities

KAP/Pro Toolset - Assure Example

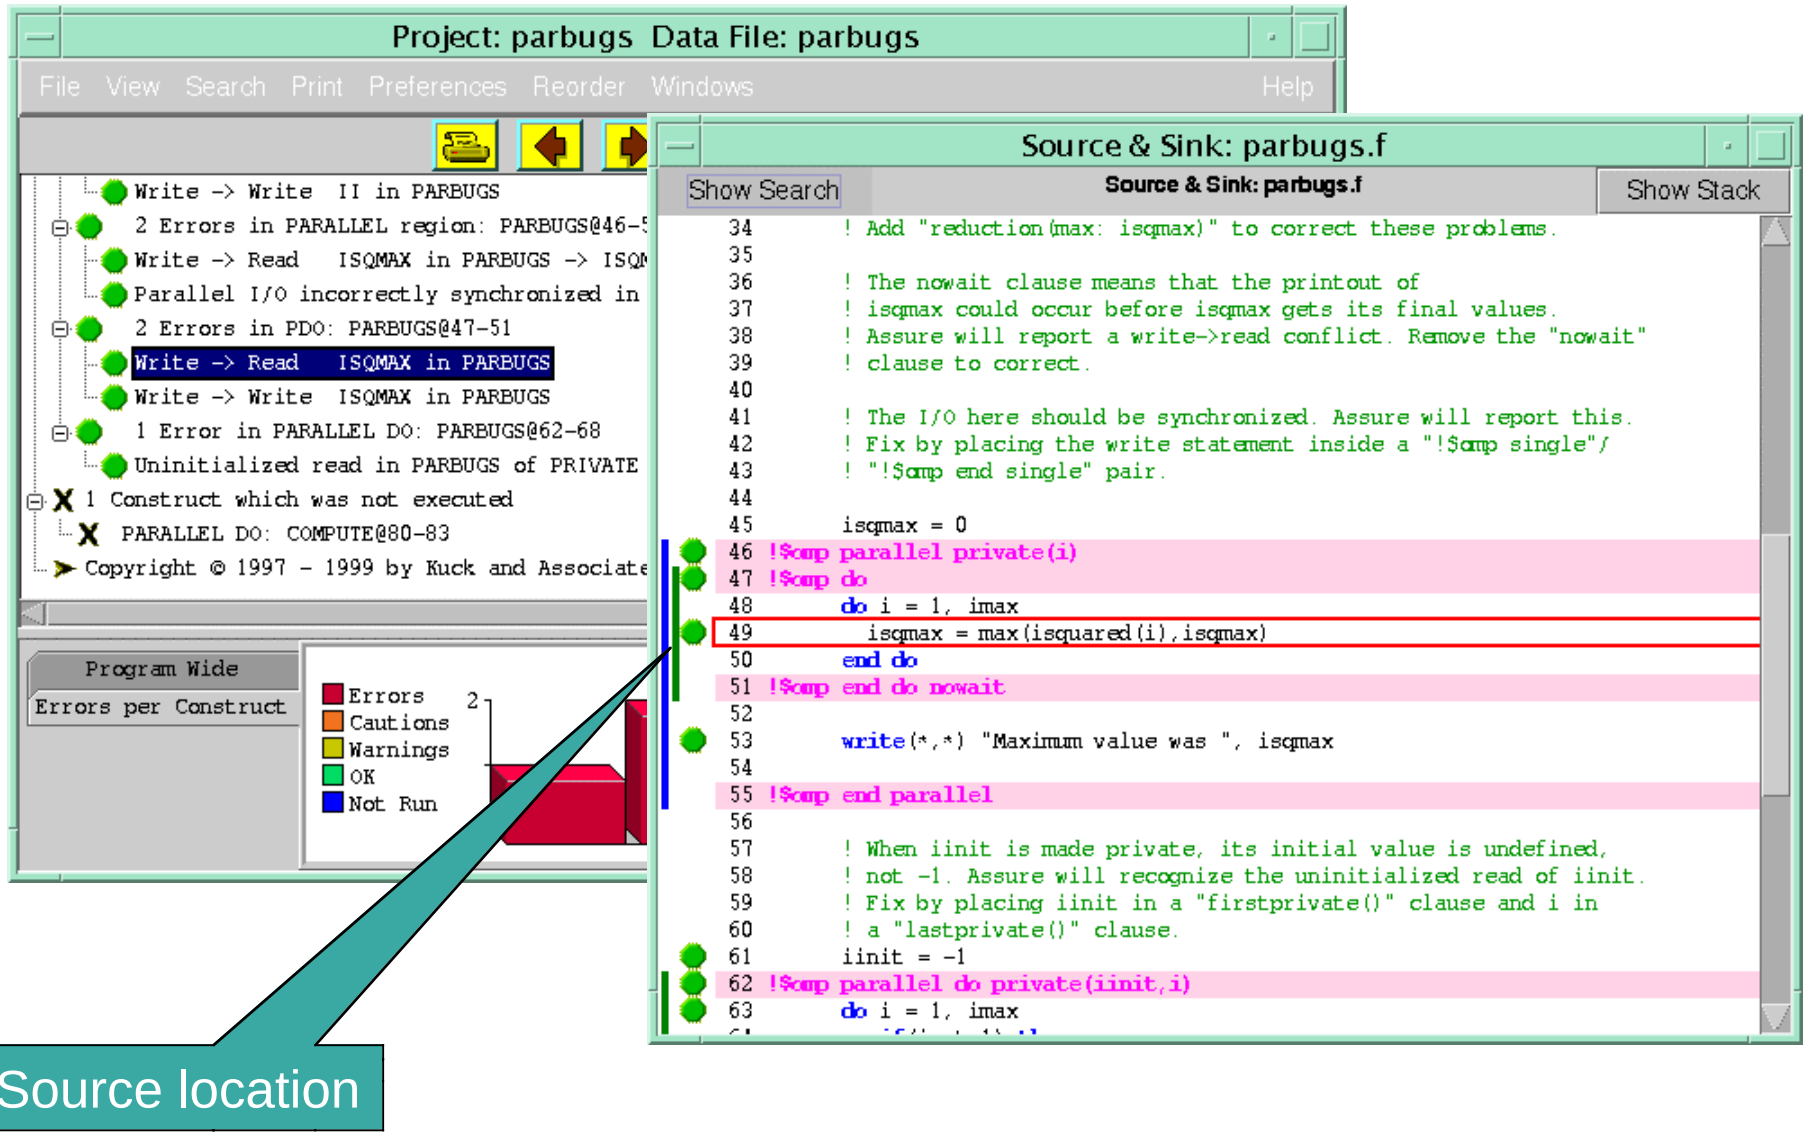

KAP/Pro Toolset - GuideView Example

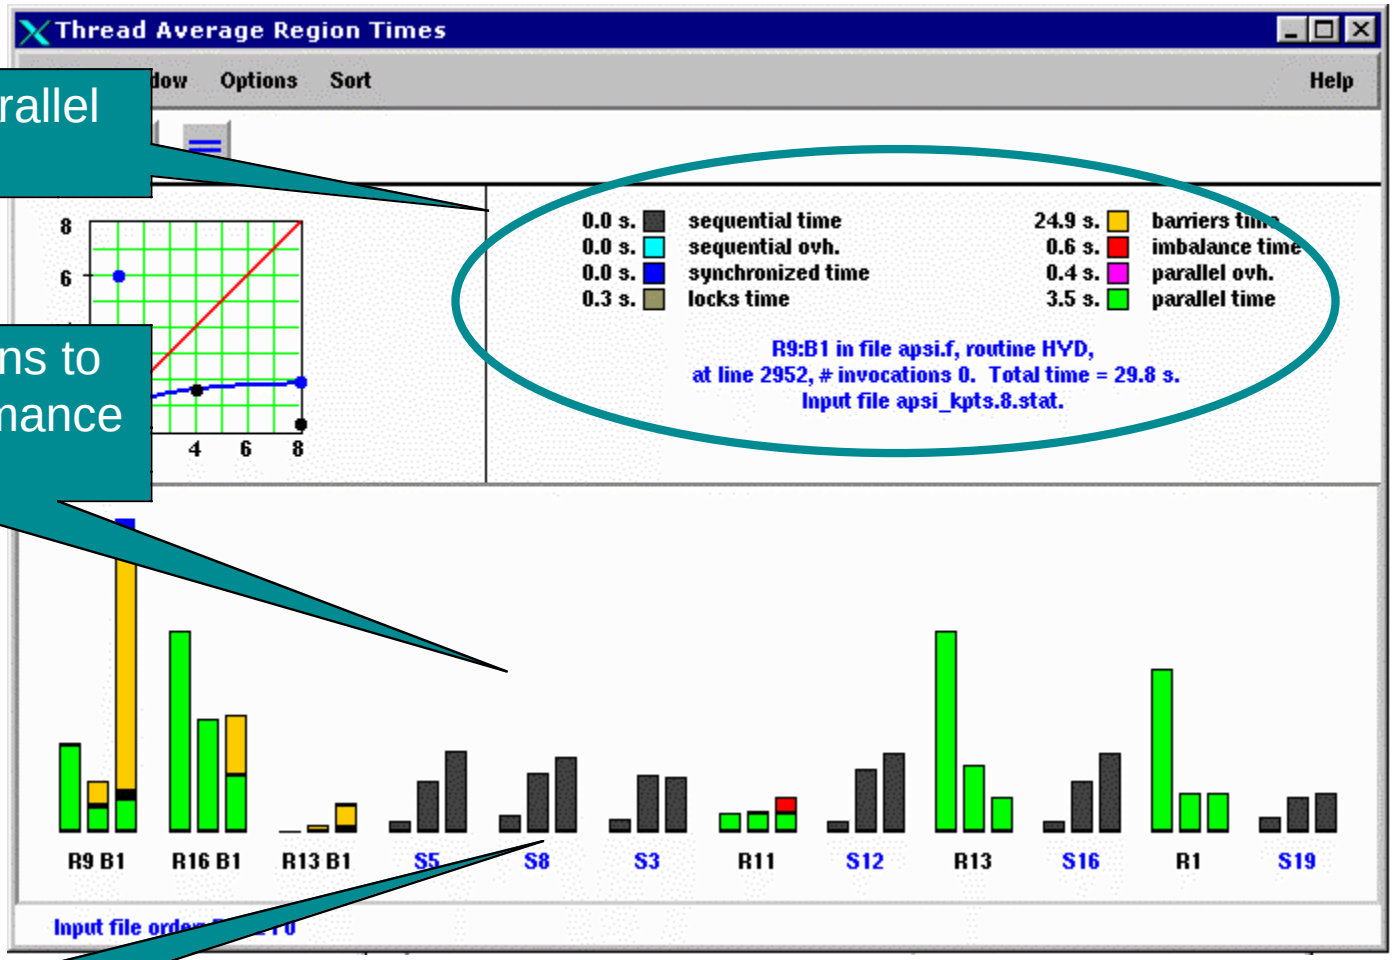

Analyse each parallel region

Sort or filter regions to navigate to performance hotspots

Identify serial regions that hurt scalability

KAI C++

The most modern, best performing, platform independant C++

- ISO C++ standard syntax, including exeptions and member templates
- ISO C++ standard class library
- multi-platform support
- meet C performance requirements
- thread safety (on most platforms)

FORESYS

- Translates FORTRAN code (F77 - F95) into abstract syntax tree (ForLib)
- FORTRAN code consistency checks (definitions of functions, common blocks etc.)
- Interactive visualization & analysis of inconsistencies
- Upgrading from FORTRAN 77 to FORTRAN 90
- Interactive/batch analysis of parallelization possibilities
- Automatic code quality analysis/improvements

pallas

Vampir 2.5

Visualization and
Analysis of
MPI
Programs

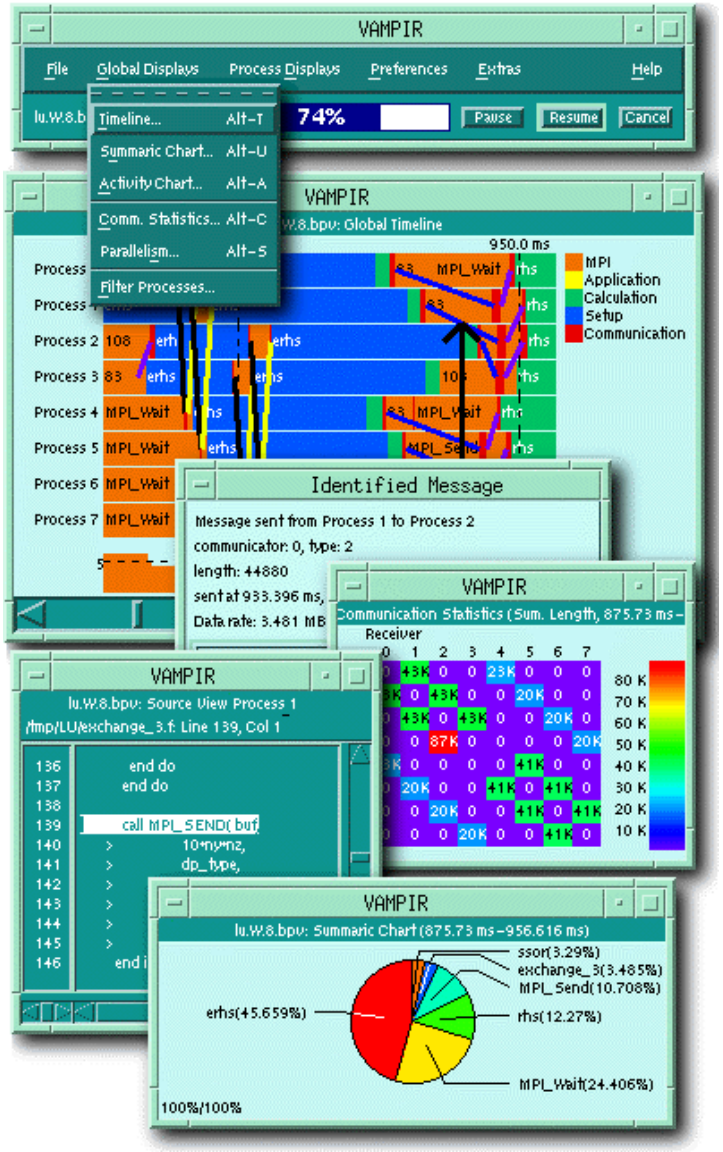

Vampir

- New version: Vampir 2.5

- Significant new features
  - support for collective MPI operations
  - trace comparison
  - tracefile re–write
  - message–length histogram
  - local and global calling trees
  - source–code reference
  - support for MPI–2 I/O operations

Vampir Features

- Offline trace analysis for MPI (and others ...)
- Traces generated by Vampirtrace tool (`ld ... -lVT -lpmpi -lmpi`)
- Convenient user–interface
- Scalability in time and processor–space
- Excellent zooming and filtering
- High–performance graphics
- Display and analysis of MPI and application events:
  - execution of MPI routines
  - point–to–point and collective communication
  - MPI–2 I/O operations
  - execution of application subroutines (optional)

- Easy customization

Vampir Main Window

Vampir 2.5 main window

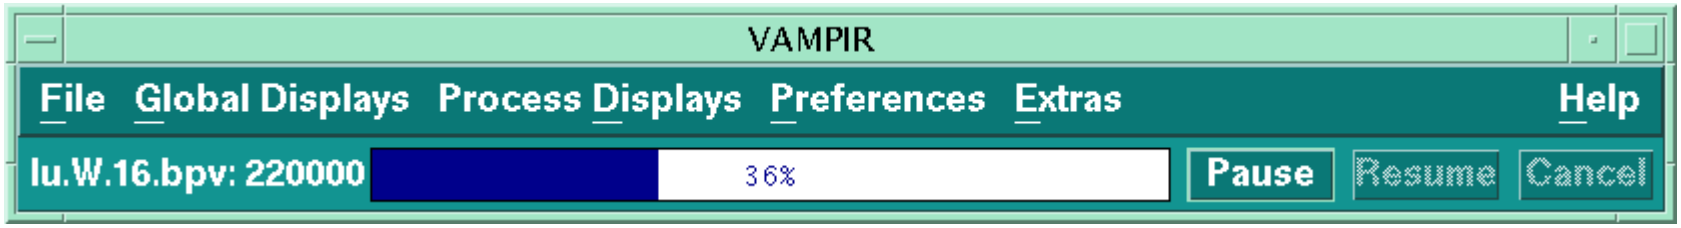

- Tracefile loading can be interrupted at any time
- Tracefile loading can be resumed
- Tracefile can be loaded starting at a specified time offset

- Tracefile can be re–written (re–grouped symbols)

Vampir Displays

- Global displays show all selected processes
  - Summary Chart: aggregated profiling information
  - Activity Chart: presents per–process profiling information
  - Timeline: detailed application execution over time axis
  - Communication statistics: message statistics for each process pair
  - Global Comm. Statistics: collective operations statistics
  - I/O Statistics: MPI I/O operation statistics
  - Calling Tree: draws global or local dynamic calling trees

- Process displays show a single process per window
  - Activity Chart
  - Timeline
  - Calling Tree

Summary Chart

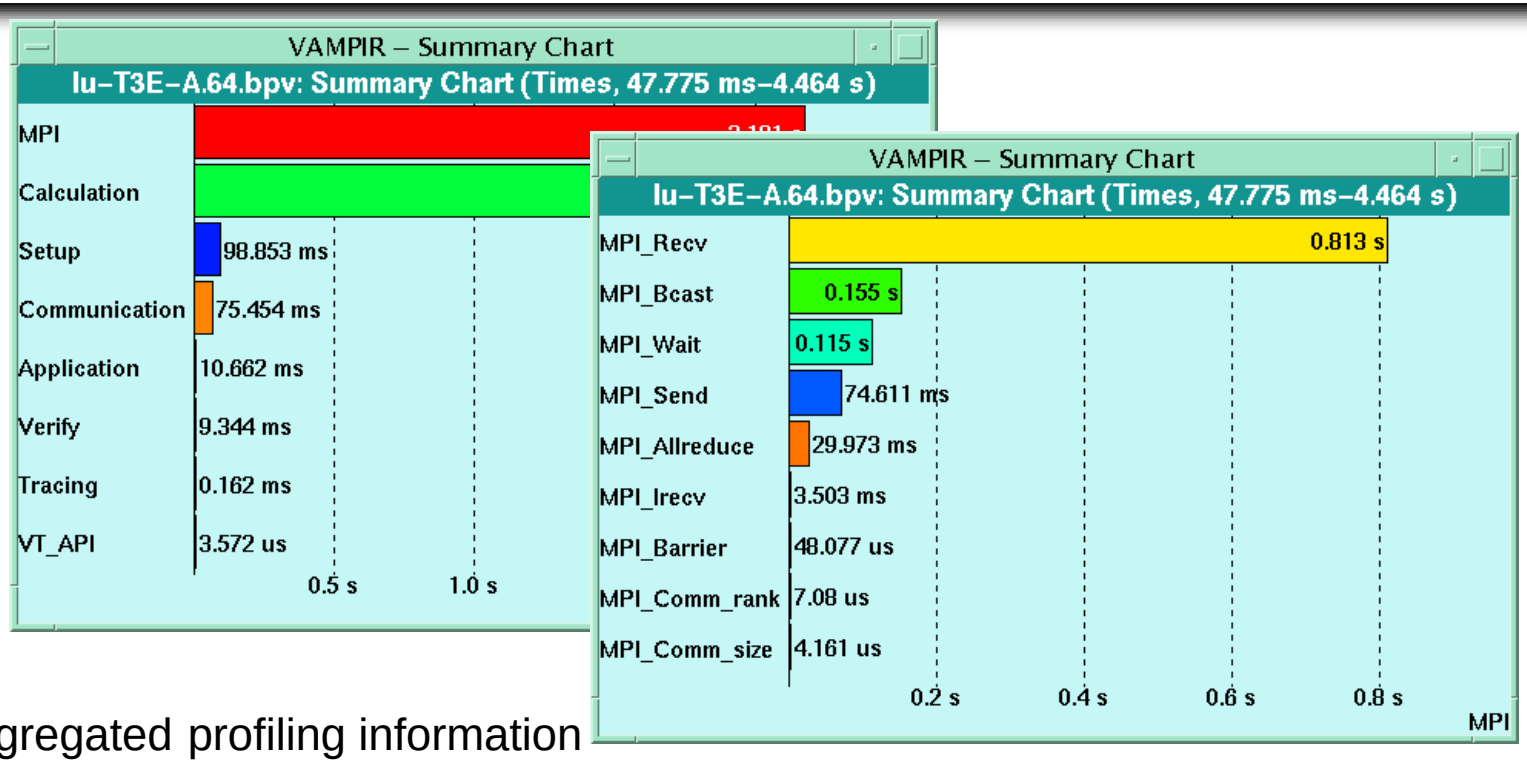

- Aggregated profiling information
  - execution time
  - number of calls
- Inclusive or exclusive of called routines
- Look at all/any category or all states
- Values can be exported/imported
- Tracefiles can be compared1

Timeline Display

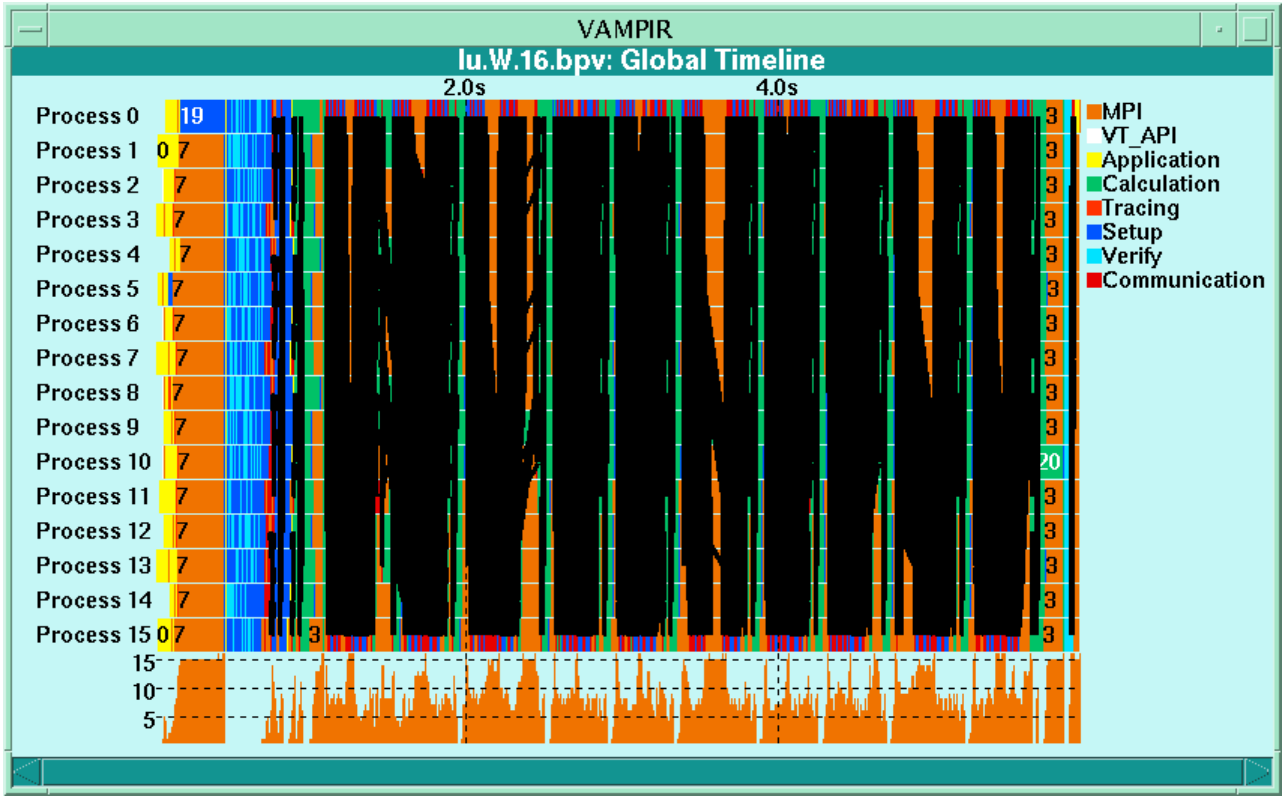

- Now displays MPI collective and I/O operations
- To zoom, draw rectangle with the mouse
- Also used to select sub–intervals for statistics

Timeline Display (Message Info)

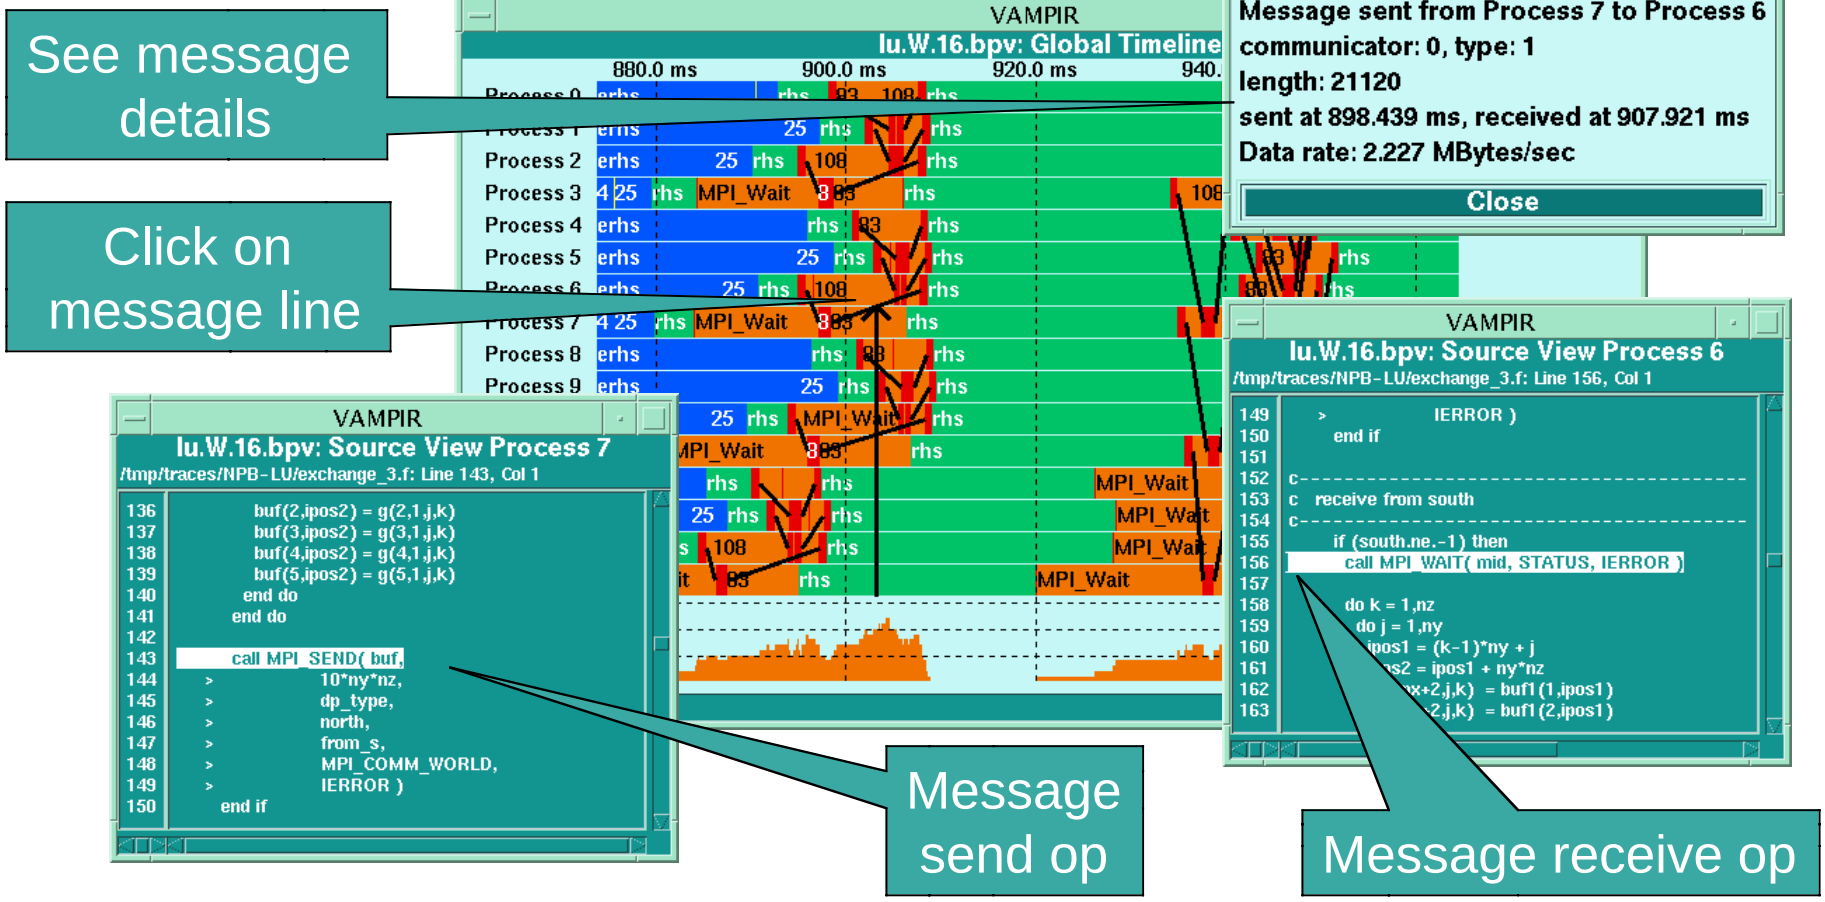

Source–code references are displayed if recorded by

Communication Statistics

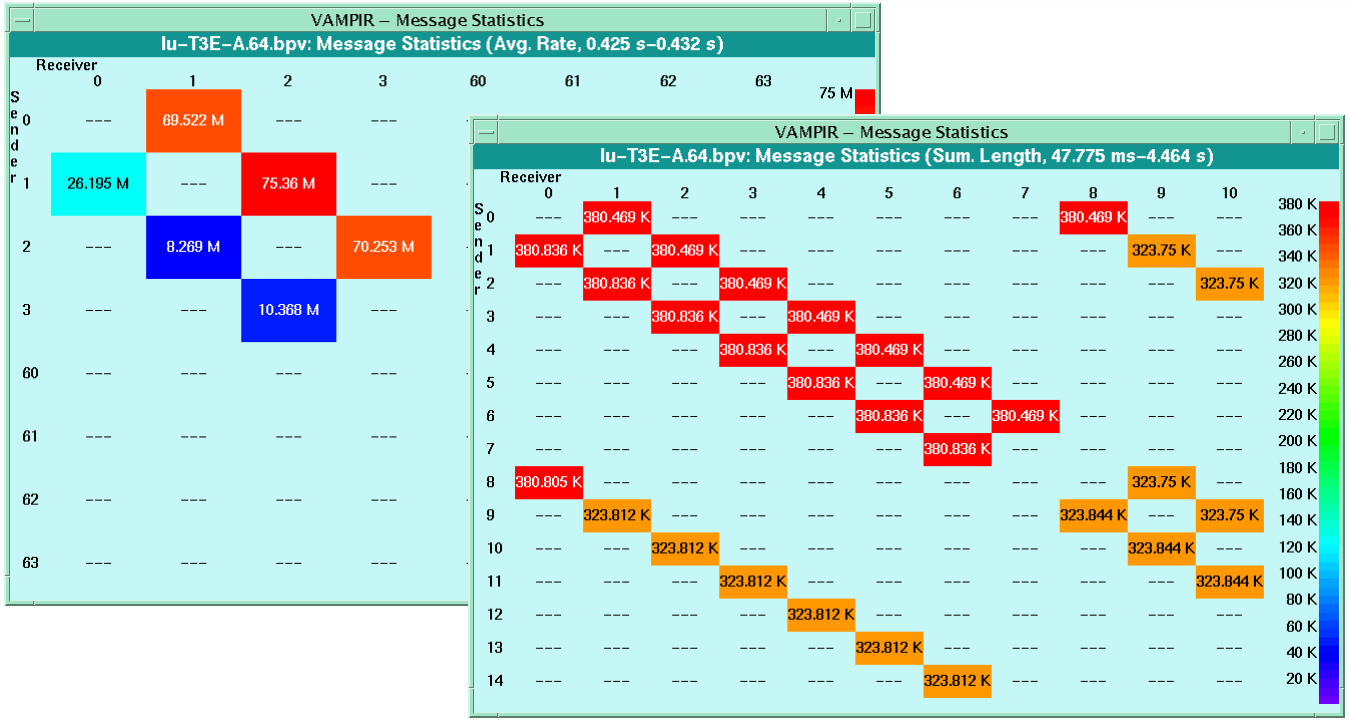

- Message statistics for each process pair:
  - Byte and message count
  - min/max/avg message length
  - min/max/avg bandwidth
- Filter for message tags or communicators

Message Histograms

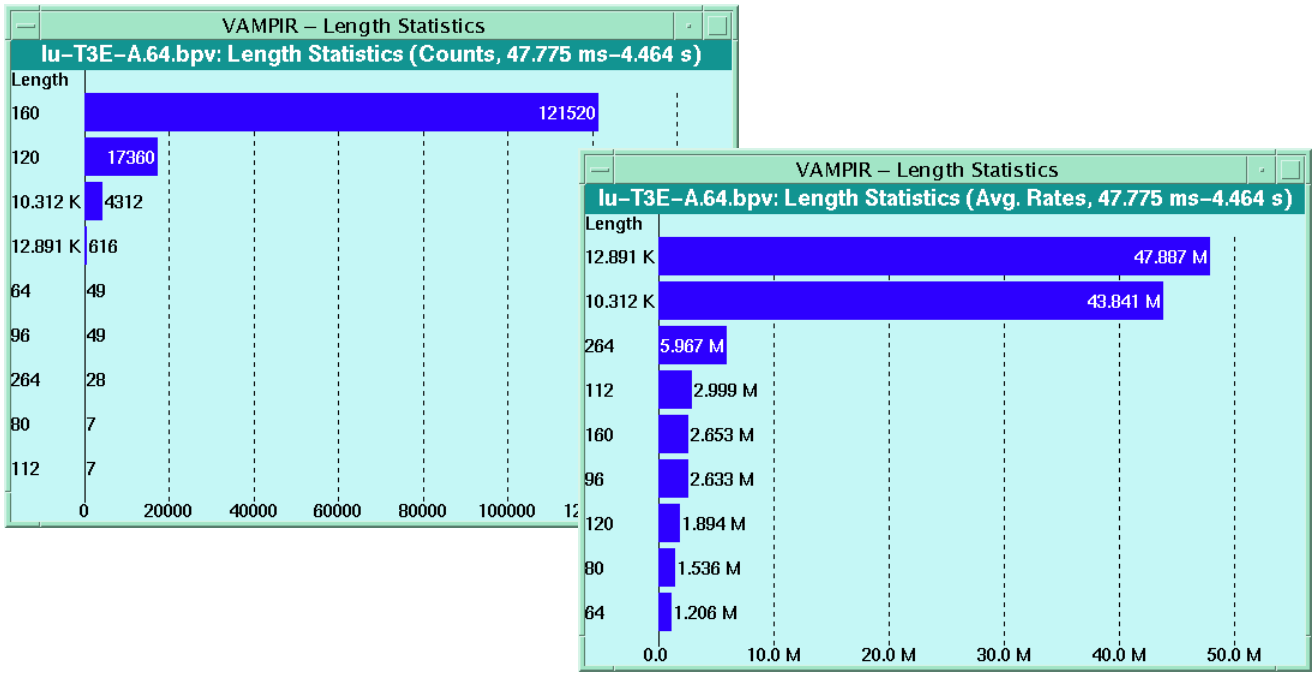

- Message statistics by length, tag or communicator
  - Byte and message count
  - min/max/avg bandwidth
- Filter for message tags or communicators

Collective Operations

- For each process: mark operation locally

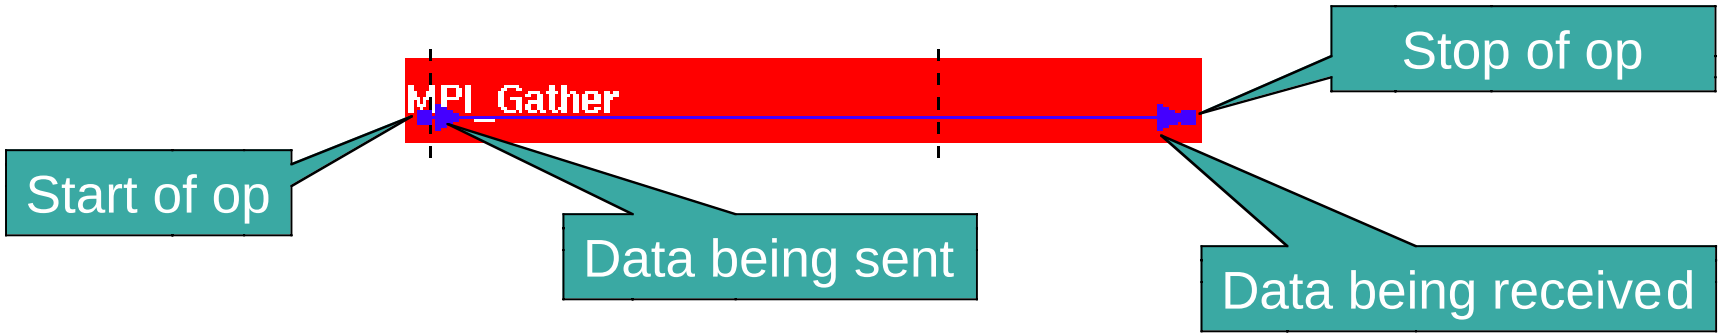

- Connect start/stop points by lines

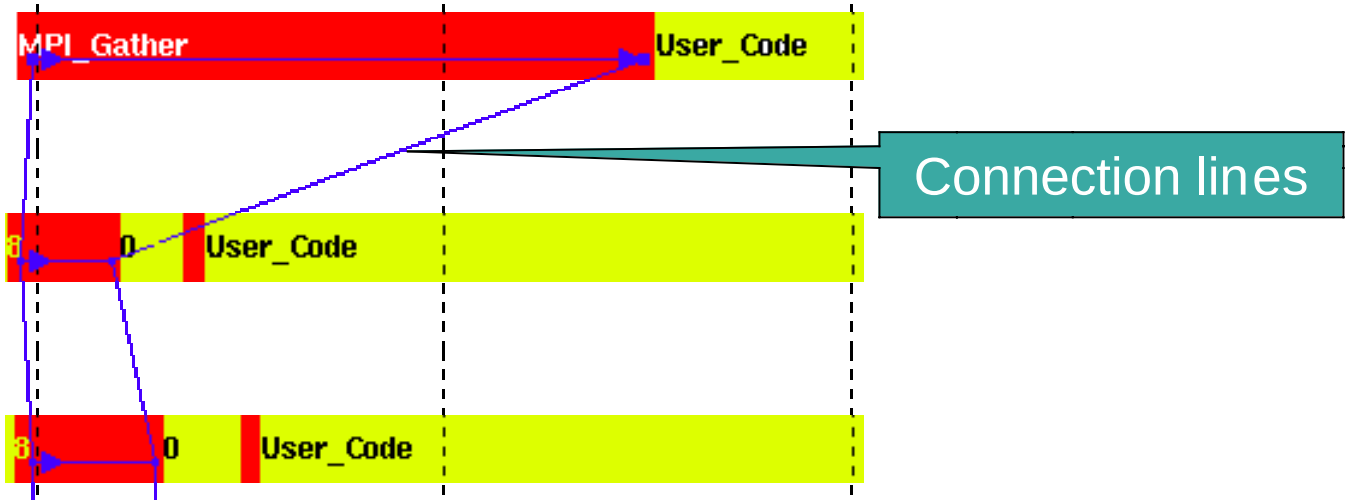

Collective Operations

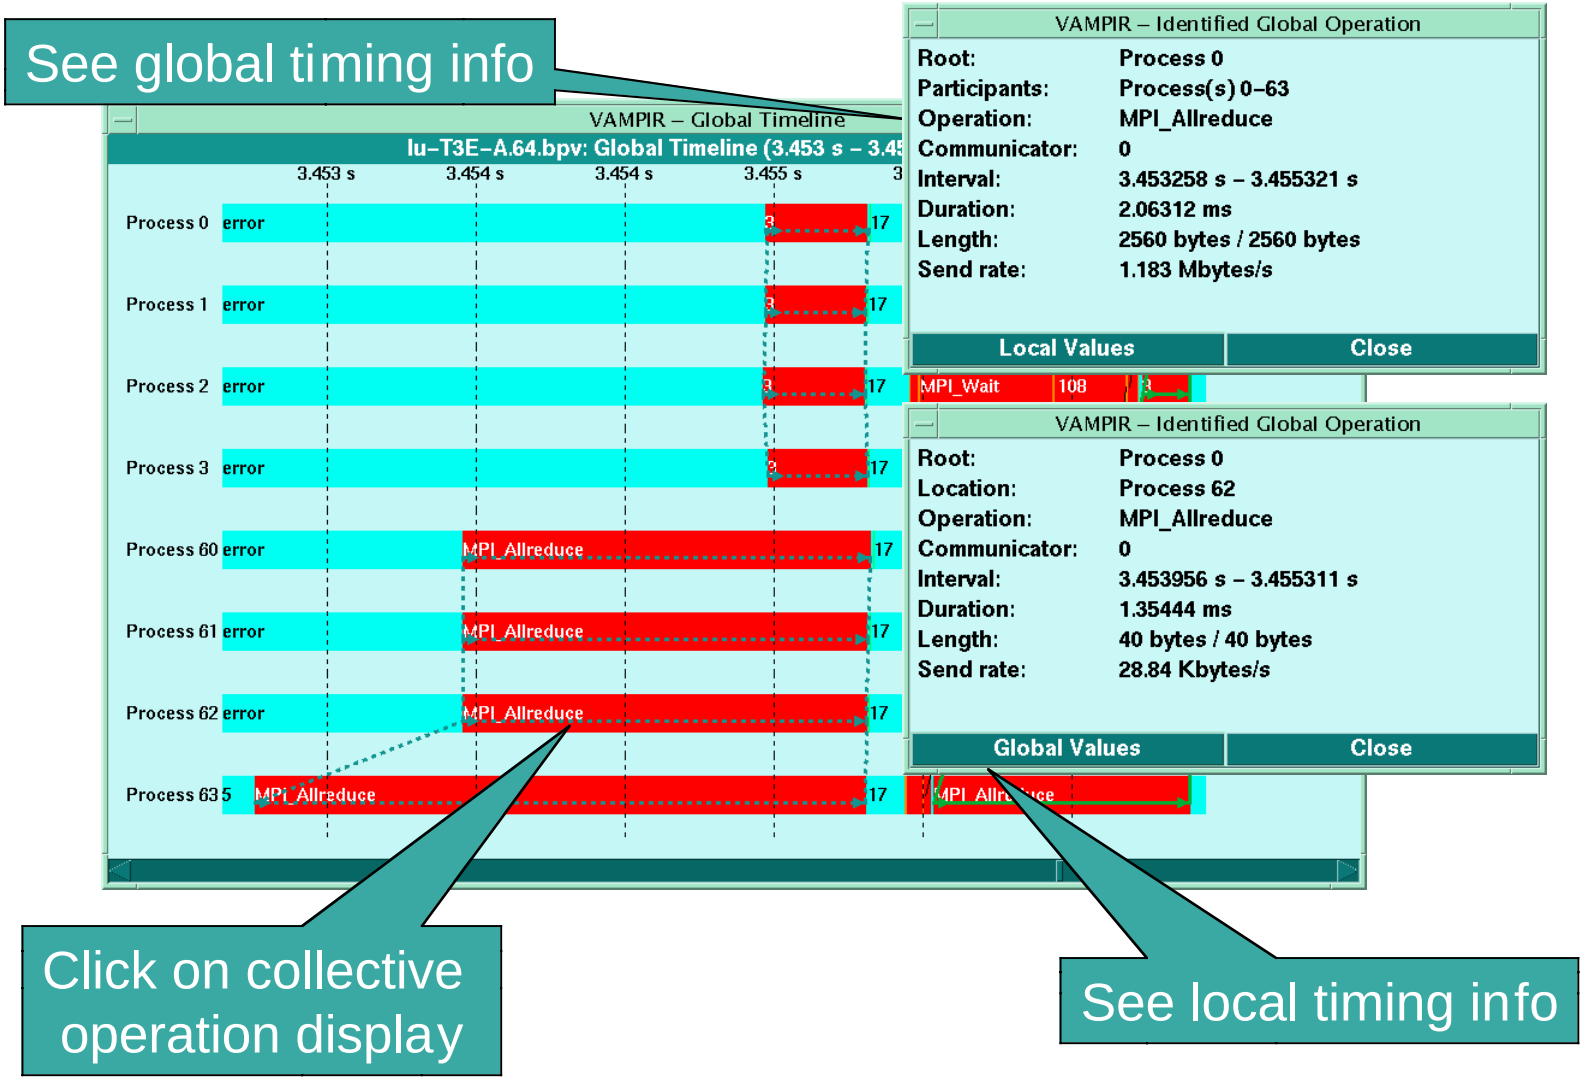

Collective Operations

- Collective operations can be filtered

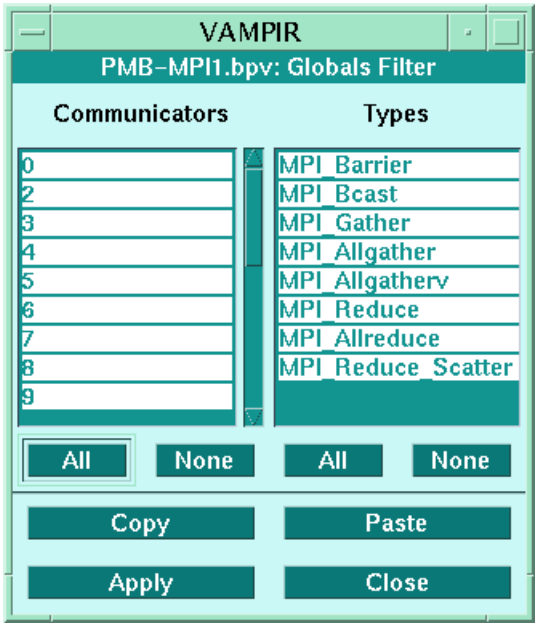

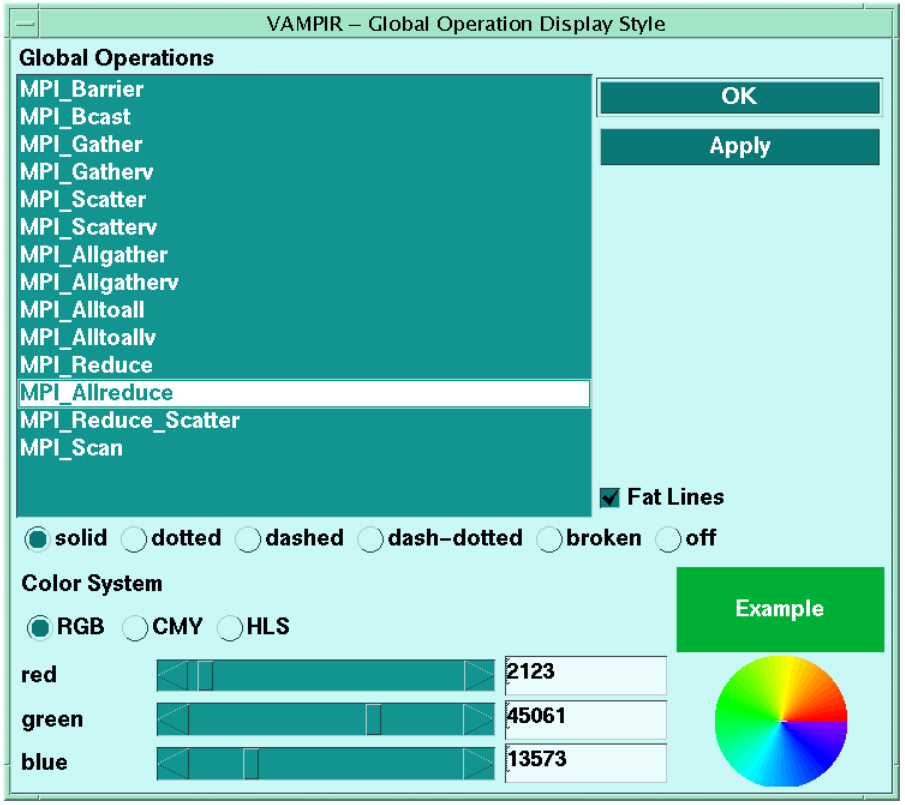

- The display style can be adapted for each collective operation

Global Communication Statistics

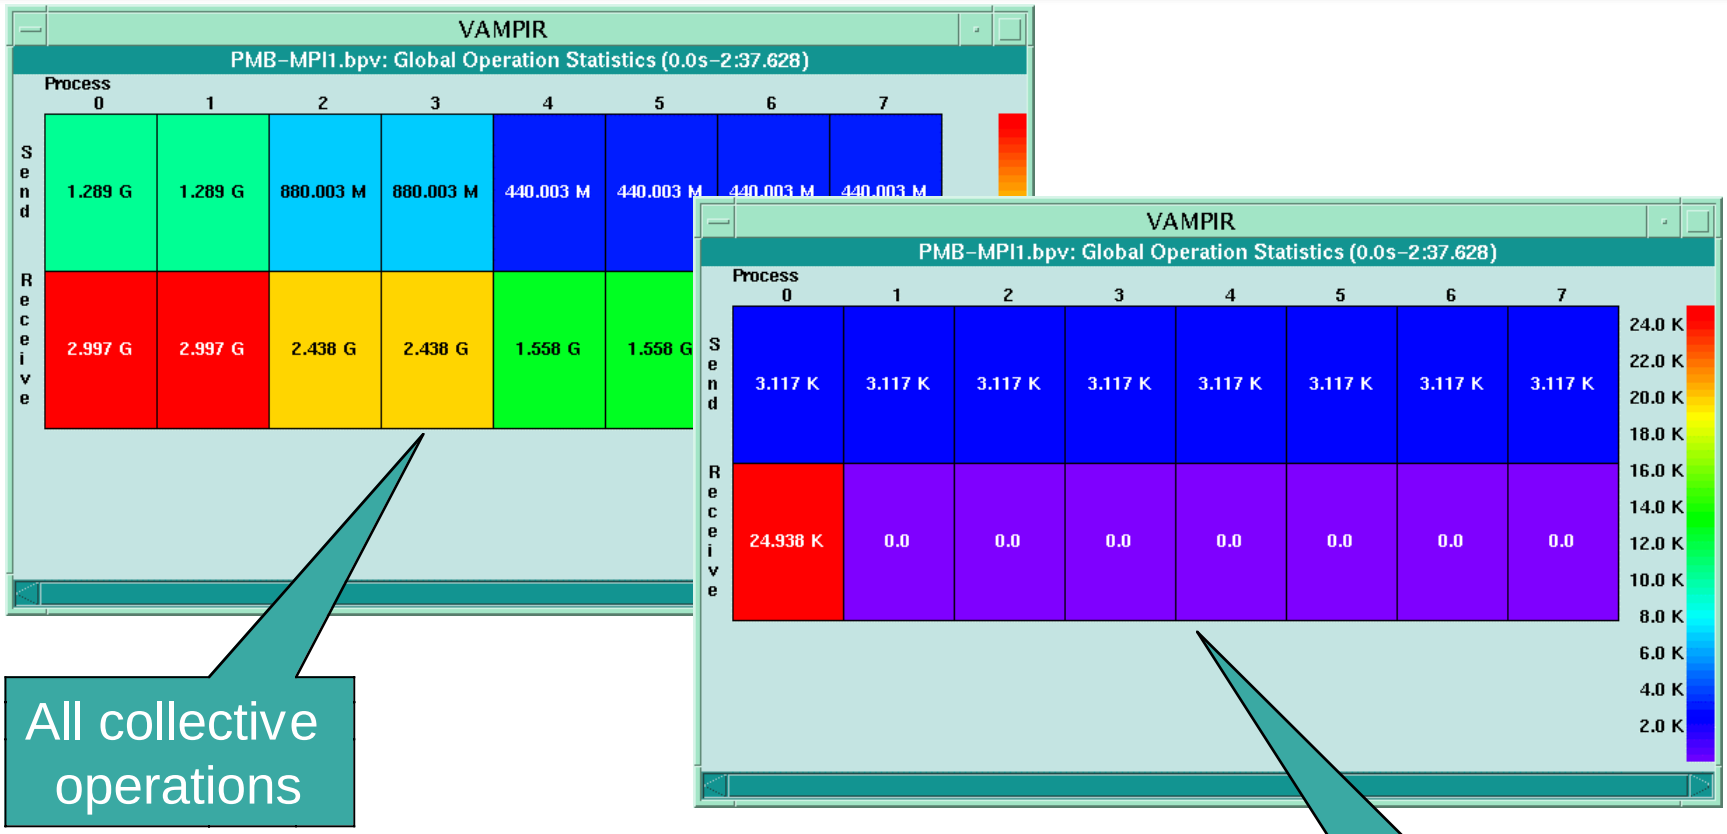

- Statistics for collective operations:
  - operation counts, Bytes sent/received
  - transmission rates
- Filter for collective operation

pallas

Vampirtrace

Tracing of MPI and Application Events

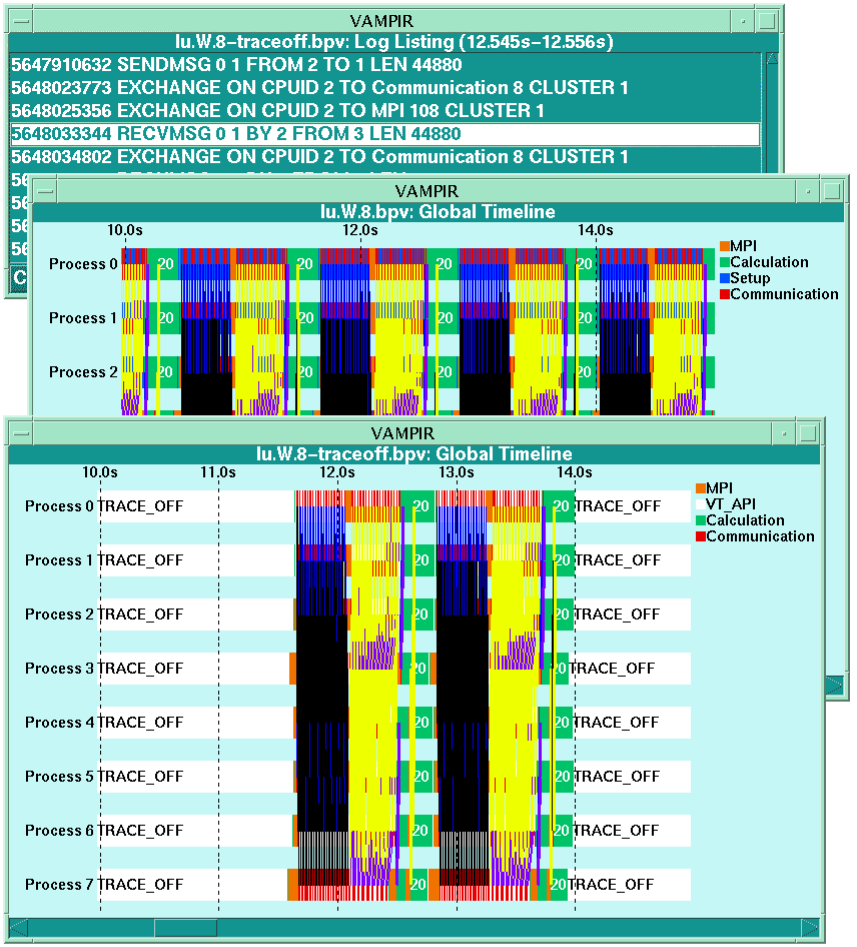

Vampirtrace

- New version: Vampirtrace 2.0
- Significant new features:
  - records collective communication
  - enhanced filter functions
  - extended API
  - records source–code information (selected platforms)
  - support for shmem (Cray T3E)
  - records MPI–2 I/O operations
- Available for all major MPI platforms

MPI–I/O Operations

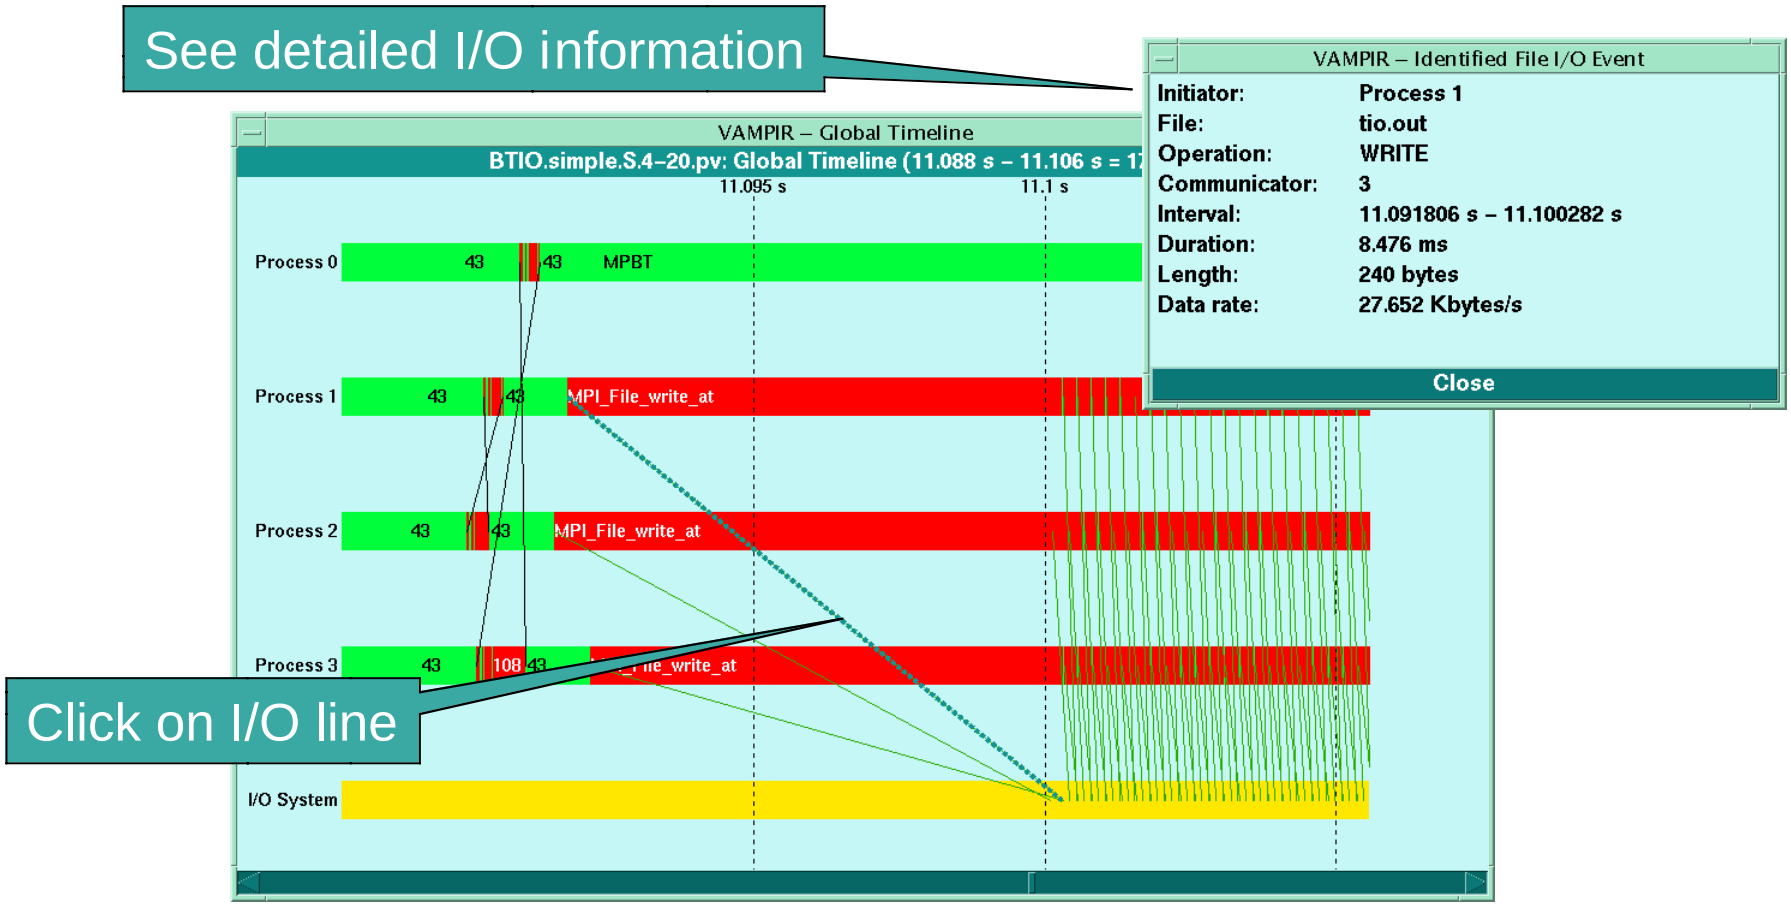

- I/O transfers are shown as lines

shmem Operations

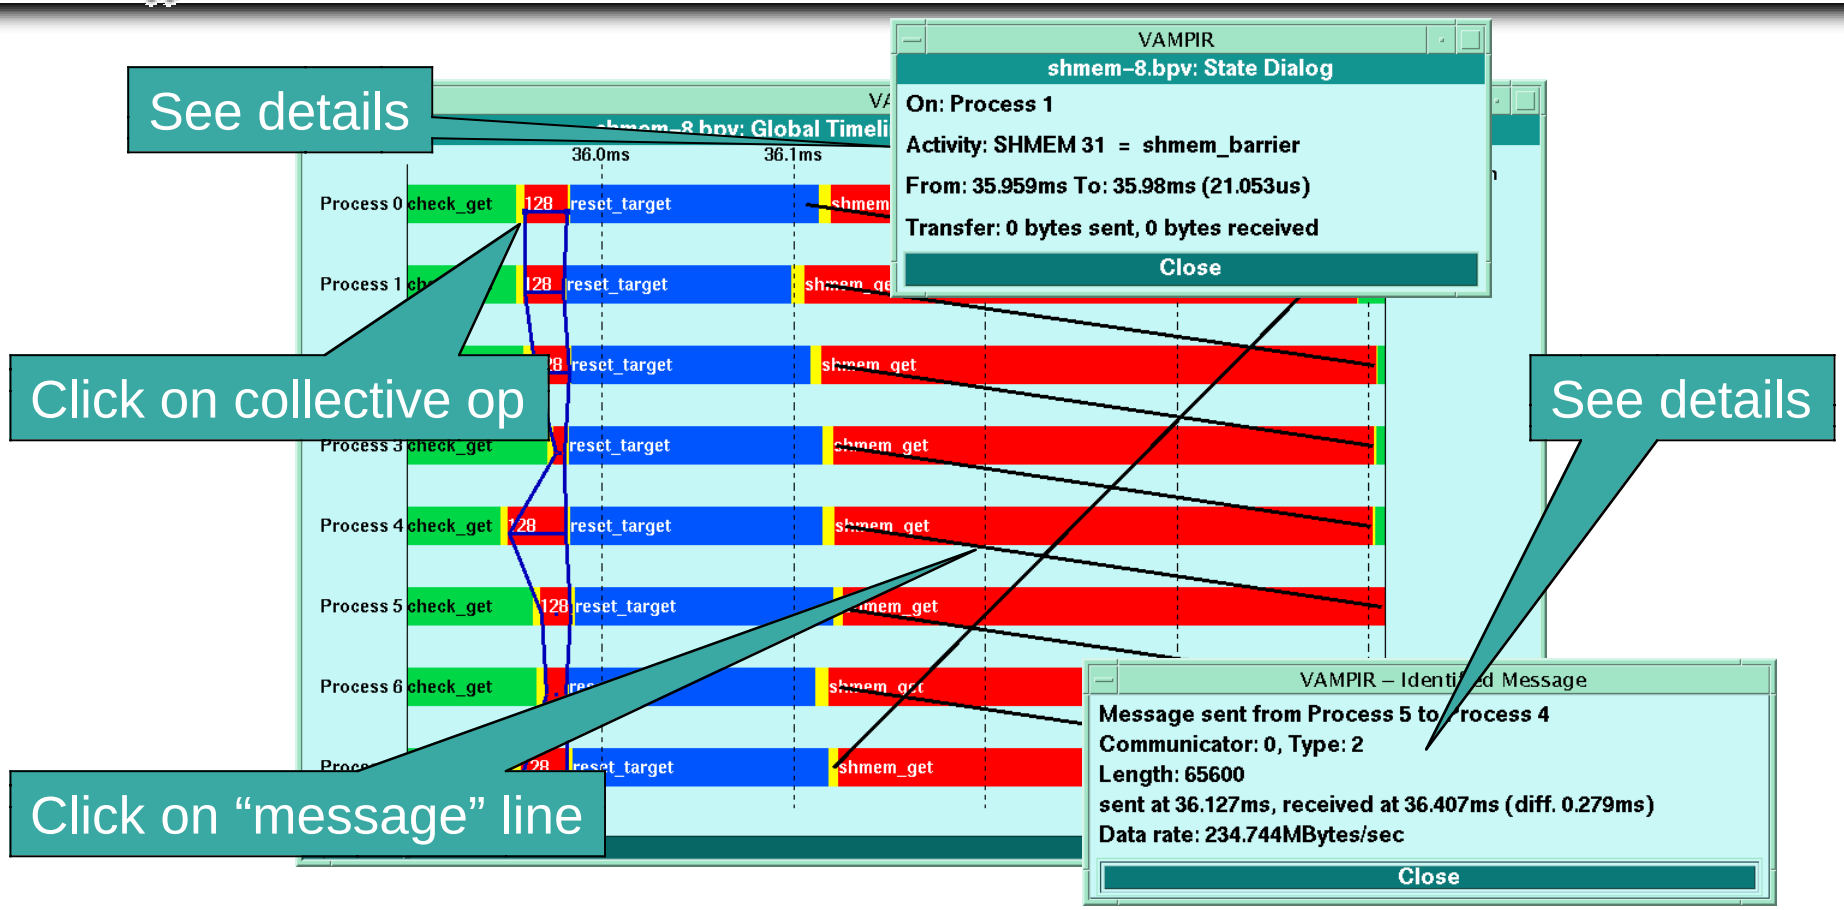

- Display one–sided transfers as messages
- Display shmem global operations

Vampir Track Record

- Reference customers: ARL, ARSC, CEWES, LANL, LLNL, MHPCC, NASA, NERSC, NSA, Cornell TC, Oregon Univ., CEA, DWD, ECMWF, GMD, HLRS, LRZ, PC^2, RUKA, ...

- URLs:
  - www.tc.cornell.edu/Edu/Tutor/Vampir
  - www.llnl.gov/sccd/lc/DEG/vampir/vampir.html
  - www.uni-karlsruhe.de/~Vampir
  - www.lrz-muenchen.de/services/software/parallel/vampir
  - www.hlrs.de/structure/support/parallel_computing/tools/performance/vampir.html

pallas

Dimemas

Performance
Prediction
Made
Easy

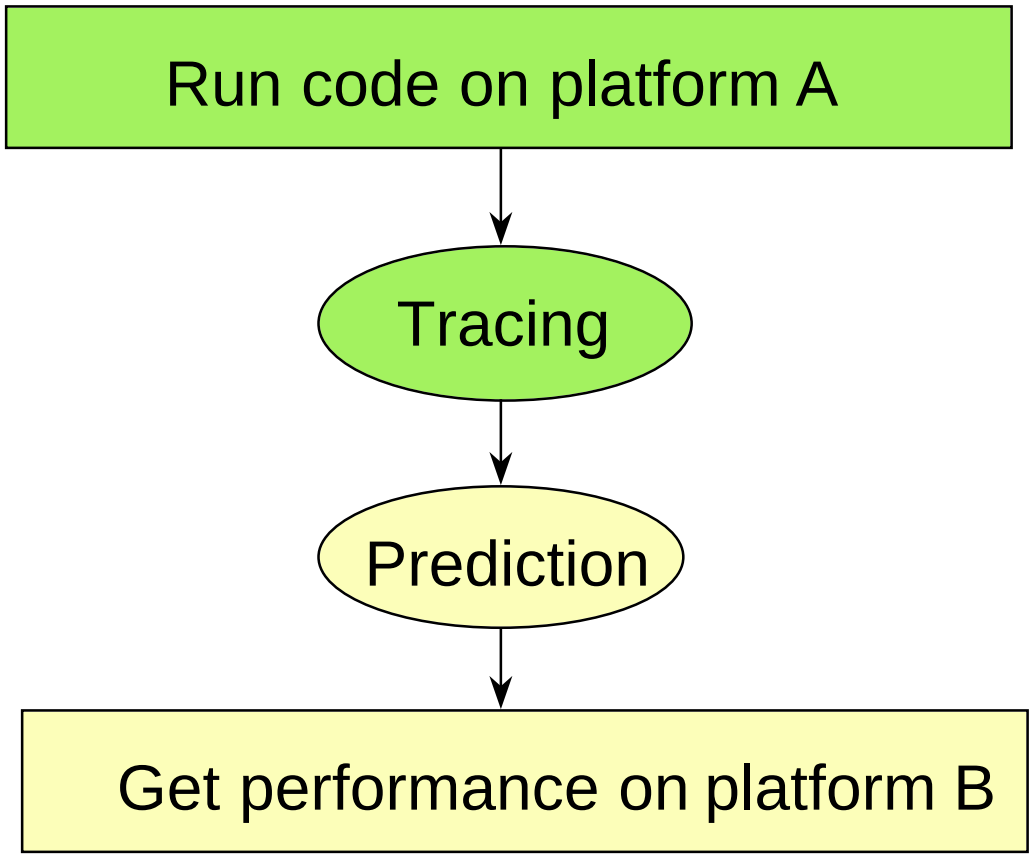

Dimemas Outputs – Vampir Tracefiles

- Actual program run

vs.

- Ideal communication

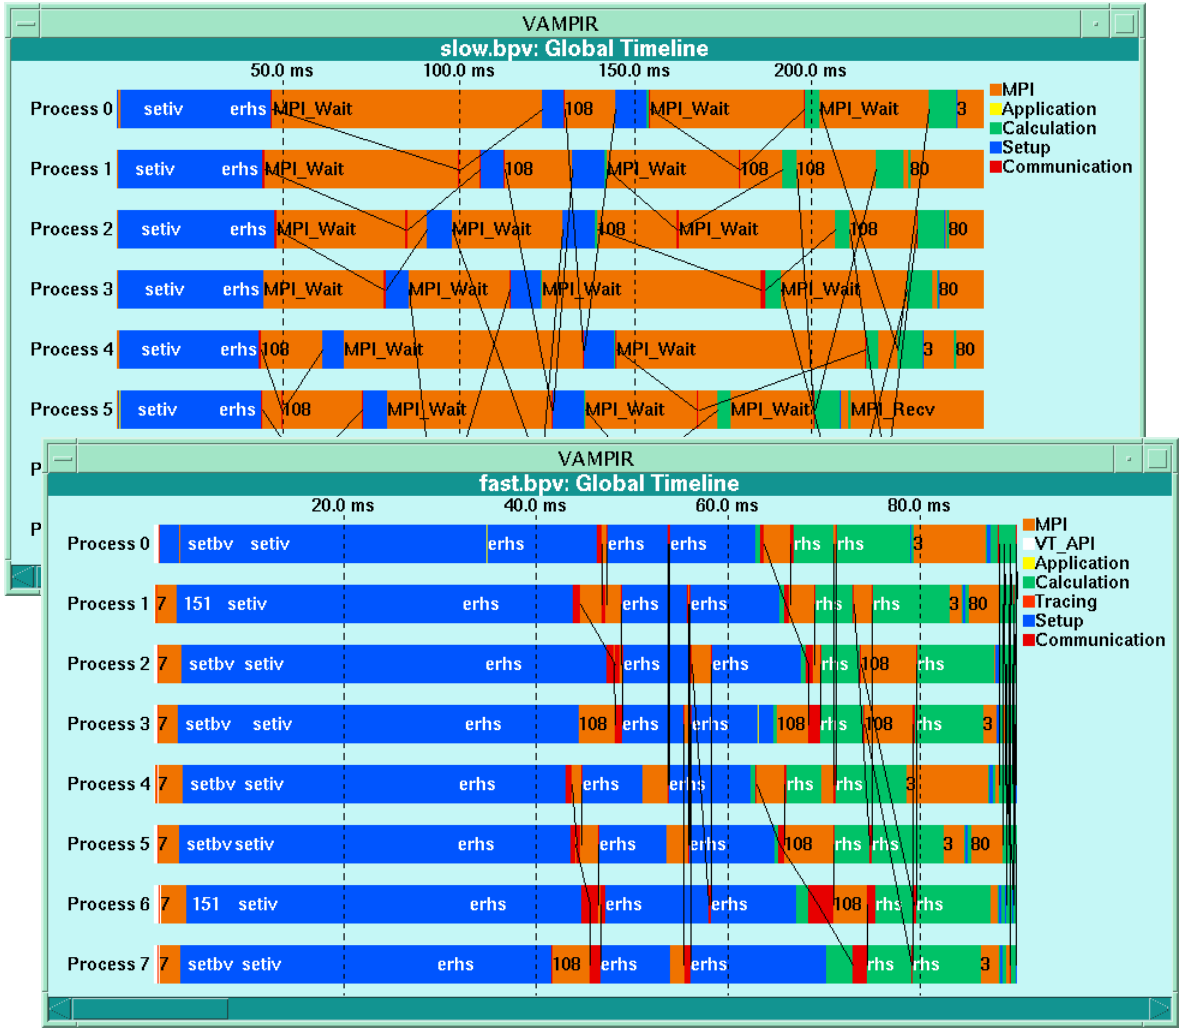

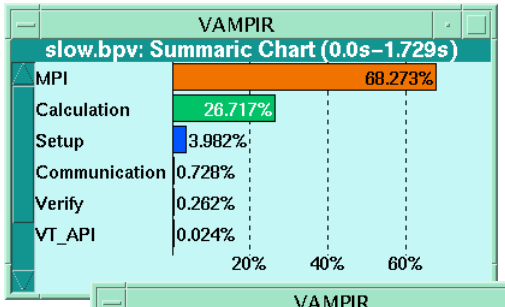

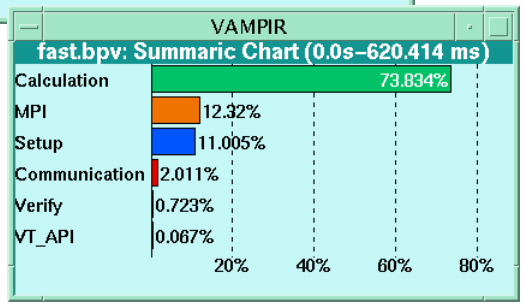

Dimemas - Prediction studies

FFT

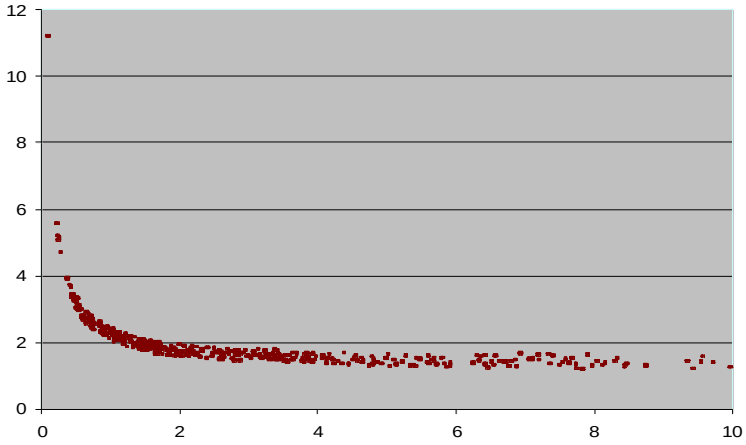

PDE

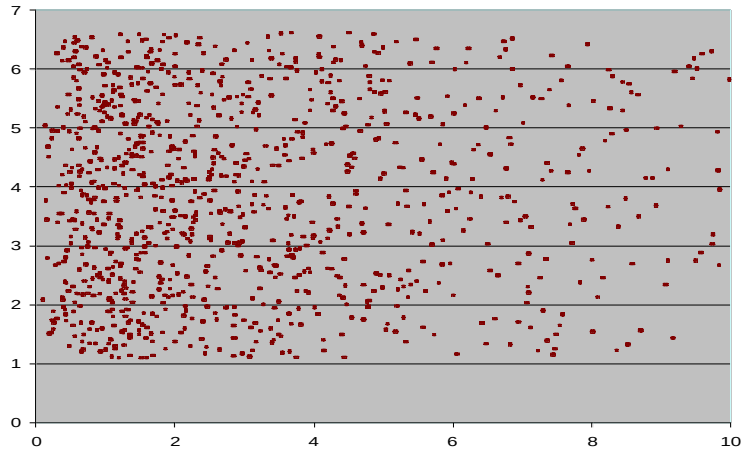

Communication bandwidth (MB/s)

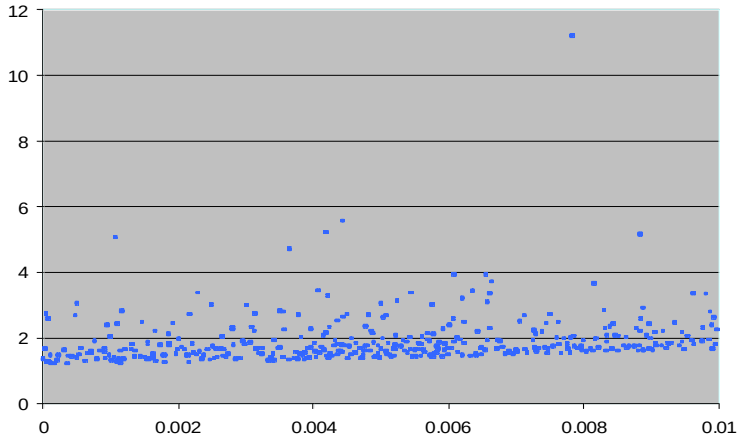

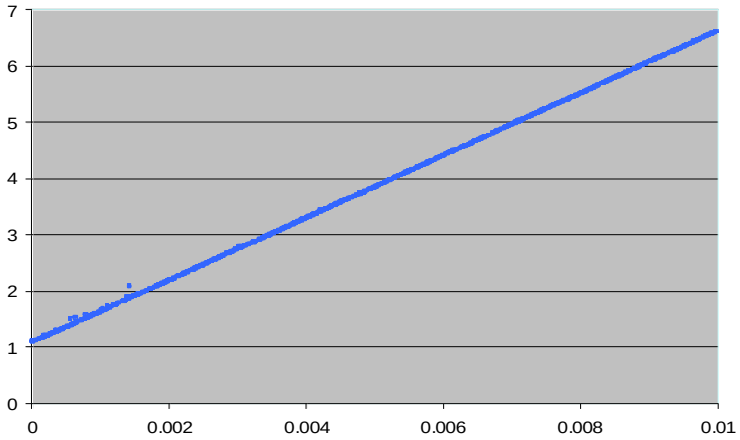

Communication latency (s)

Example Results – Performance Prediction

- NAS LU benchmarks

Benchmark	Class	Time	Simulated	
BT	*A*	316.580000	319.858708	1.04%
BT	*W*	6.410000	5.548973	13.43%
CG	*A*	4.480000	4.210193	6.02%
CG	*B*	342.860000	335.641787	2.11%
CG	*W*	1.130000	1.210509	7.12%
EP	*A*	18.970000	18.843001	0.67%
EP	*B*	75.990000	77.122498	1.49%
EP	*W*	2.380000	2.391074	0.47%
LU	*A*	131.660000	129.139735	1.91%
LU	*B*	793.070000	785.966349	0.90%
LU	*W*	16.550000	17.415359	5.23%
MG	*A*	16.850000	16.290057	3.32%
MG	*B*	59.480000	64.974789	9.24%
MG	*W*	1.052000	1.112022	5.71%
SP	*A*	152.020000	151.821937	0.13%
SP	*W*	16.950000	15.237959	10.10%

pallas

Future plans - Vampir, Dimemas

- Towards automatic performance analysis
  - improve user guidance in Vampir and Dimemas
  - add “assistant” module for inexperienced users
- Support for clustered shared–memory systems
  - support shared–memory programming models (threads, OpenMP)
  - expose cluster structure
  - aggregate information on SMP nodes
- Support for (very) large systems
  - new structured tracefile format
  - fine–grain interactive control over tracing
  - scalable displays
  - new Vampir structure (can exploit parallelism)

Access to Pallas Tools

Download free evaluation copies

http://www.pallas.com

Thanks for your attention!

Pallas GmbH
Hermülheimer Straße 10
D–50321 Brühl,
Germany

info@pallas.com
www.pallas.com