www.bsc.es

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC Demos

Interested in any of these topics?
Ask for a demo

SC14 New Orleans, LA hpc matters.

OmpSs Development Environment and Training Material

- OmpSs Examples:
- Debugging
- What to taskify?
- Paraver

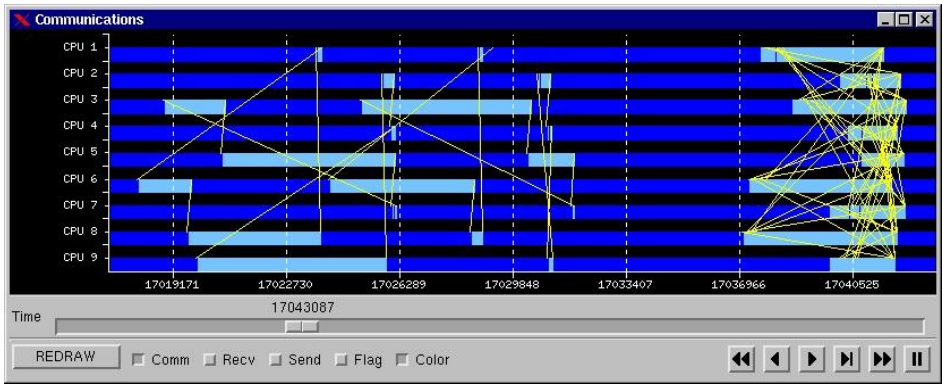

Parallel Python at Heterogenous Architectures

- PyCOMPSs Automatic parallelization
- Xeon and Xeon Phi
- Neuroscience application

Tasks definition

```
@task(originals_file = FILE_INOUT, surrs_file = FILE_INOUT)
def gather(result, originals_file, surrs_file):
#   Accumulate results

@task(returns = list)
def cc_surrogate_range(start_idx, end_idx, seed, num_neurons, num_surrs, num_bins,
maxlag):
#   Compute neuron correlations
```

Main program

```
originals_file = 'result_cc_originals.dat'
surrs_file = 'result_cc_surrogates_conf.dat'
for frag in range(num_frags):
    start_idx = end_idx
    end_idx = int(min((frag+1)*step,num_ccs))
    result = cc_surrogate_range(start_idx, end_idx, seed, num_neurons,
num_surrs, num_bins, maxlag)
    gather(result, originals_file, surrs_file)
    seed = seed + delta
```

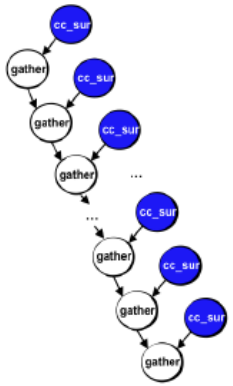

Interested in any of these topics? Ask for a demo

- Transparent resilience at OpenMP
  - NanoFT: Task based fault tolerance

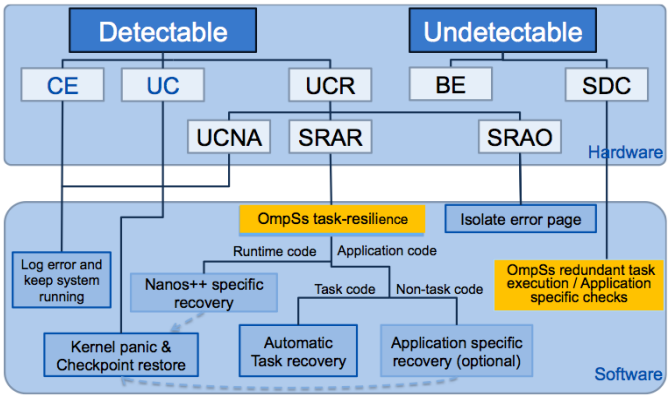

- Workflow development environment
  - COMPSs workflow in Java
  - IDE for development and deployment

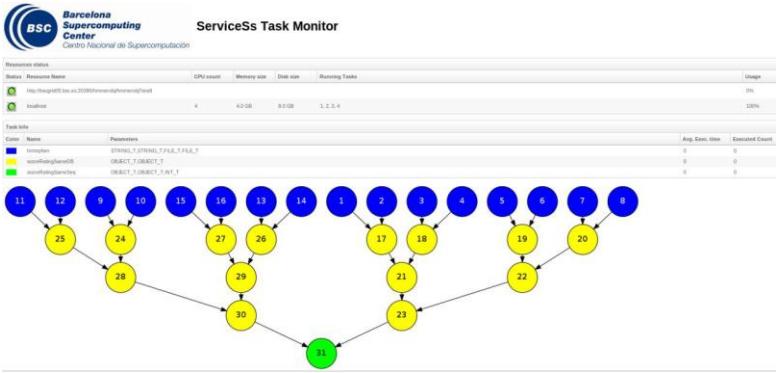

Interested in any of these topics? Ask for a demo

- Integration of workflow and Data Infraestructures
  - Persistent objects in Parallel Python
    - Cassandra and data clay

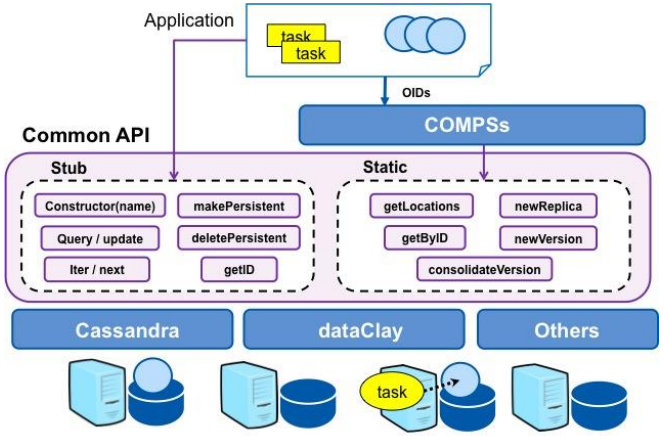

- Instantaneous performance metrics
  - With no noticeable overhead
  - Example uses
    - Improving production codes
    - Analysing impact of multicore sharing

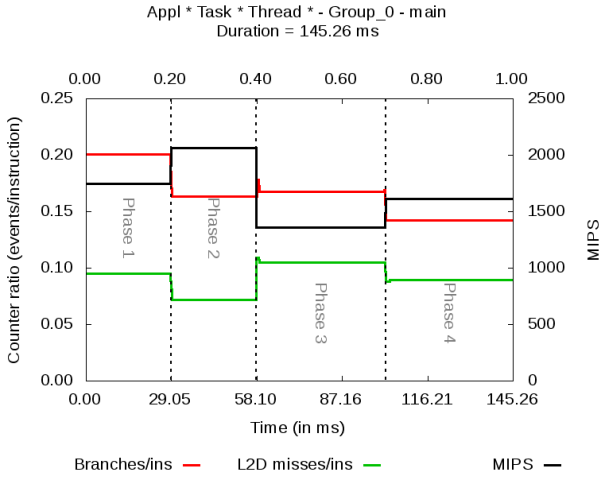

Interested in any of these topics? Ask for a demo

- Multilevel architectural simulation
  - Integration Tasksim and Dimemas
  - Trace-driven simulation of communications between nodes
  - Execution driven simulation of node level runtime
  - Trace driven simulation of core

- Homogenizing heterogeneity
  - Intra-node heterogenity
    - OpenCL, CUDA, OpenMP 4.0, FPGAs, etc
  - Cluster heterogeneity
    - OmpSs @ Cluster
    - OmpSs MPI offload

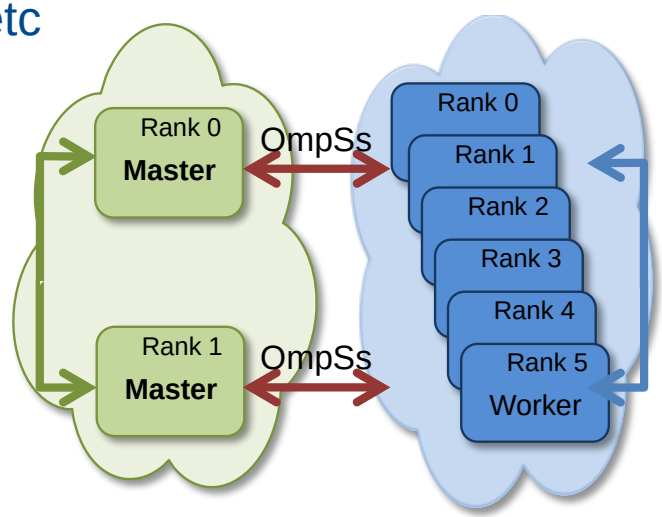

- Generic DSL infrastructure
  - Common framework to speedup development of HPC DSLs
    - OmpSs + Lightweight Modular Staging (LMS) + OpenCL
  - Uses case: Saiph: A DSL for solving Partial Differential Equations

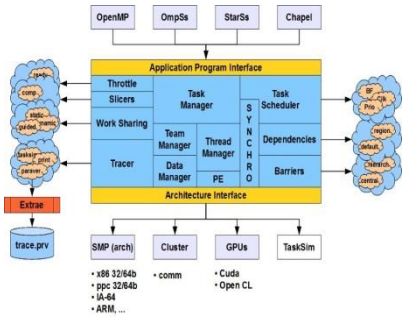

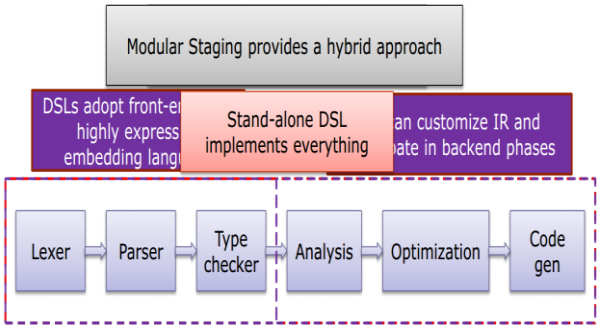

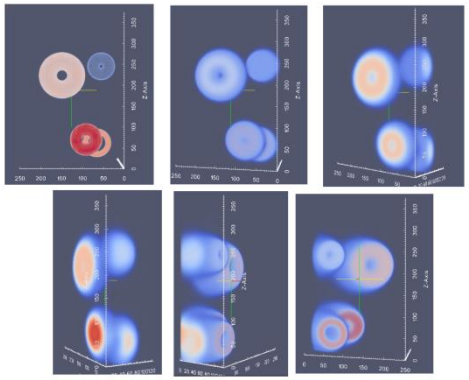

Interested in any of these topics? Ask for a demo

- Tracking the performance of computation regions
  - Changing core count …
  - … problem size …
  - … processes per node

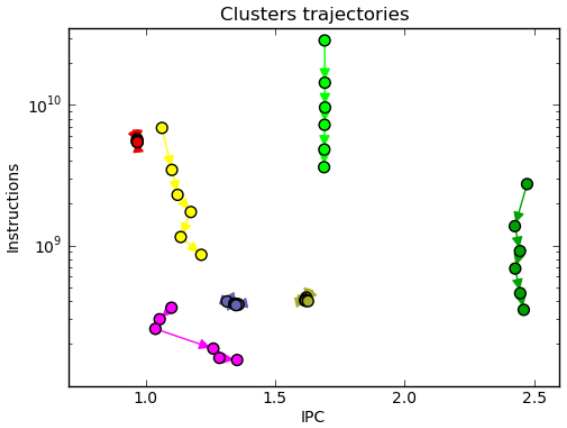

- Performance models
  - Characterizing fundamental efficiency factors
  - Scalability prediction for production codes