BurstClustering
Cluster-based analysis of Paraver traces

Lau Mercadal, Sandra Méndez, Germán Llort
(tools@bsc.es)
Barcelona Supercomputing Center

Running BurstClustering

```
alogin1> module load clustering_suite
alogin1> cd $HOME/bsctools/clustering
alogin1> BurstClustering -d cluster.xml \
        -i lulesh2.0_64mpi_cpu-bind16numa.chop.prv \
        -o lulesh2.0_64mpi_cpu-bind16numa.chop.clustered.prv
...
alogin1> ls *clustered*
… *.clustered.clusters_info.csv
… *.clustered.prv
… *.clustered.IPC.PAPI_TOT_INS.gnuplot
```

Statistics per cluster in a plain CSV-text file

Paraver trace enriched with cluster information

Scatter plot to visualize cluster structure

Cluster-based analysis

- Check the resulting scatter plot

```
alogin1> gnuplot lulesh2.0_64mpi_cpu-bind16numa.chop.clustered.IPC.PAPI_TOT_INS.gnuplot
```

- Identify main computing trends
  - More work: Clusters 1 and 5
  - Higher IPC: Clusters 3 and 4
- Mostly constant work (fixed Y) & variable speed (X)
  - Except Cluster 2, which also shows variable work (Y)
- Look at the clusters shape
  - Compact vs. elongated and widespread
  - Variability in any axis highlights **potential imbalances**

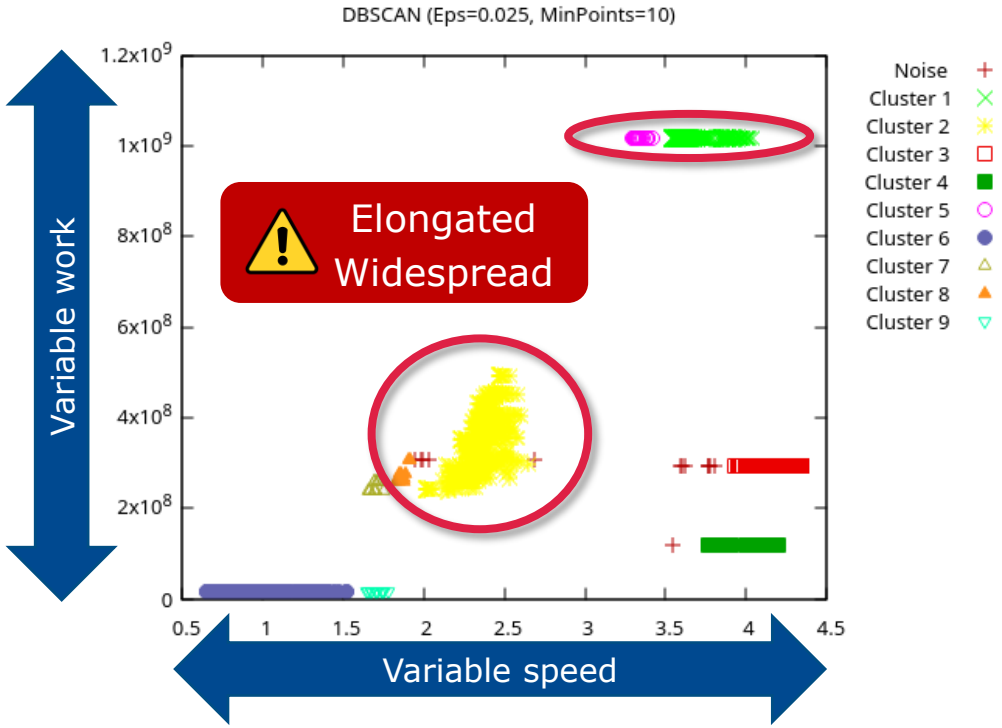

Correlating scatter plot and time distribution

- Open the clustered trace with Paraver and look at it

```
laptop> $HOME/paraver/bin/wxparaver <path-to>/lulesh2.0_64mpi_cpu-bind16numa.chop.clustered.prv
```

- Display the distribution of clusters over time
  - File → Load configuration → `$HOME/paraver/cfgs/clustering/clusterID_window.cfg`

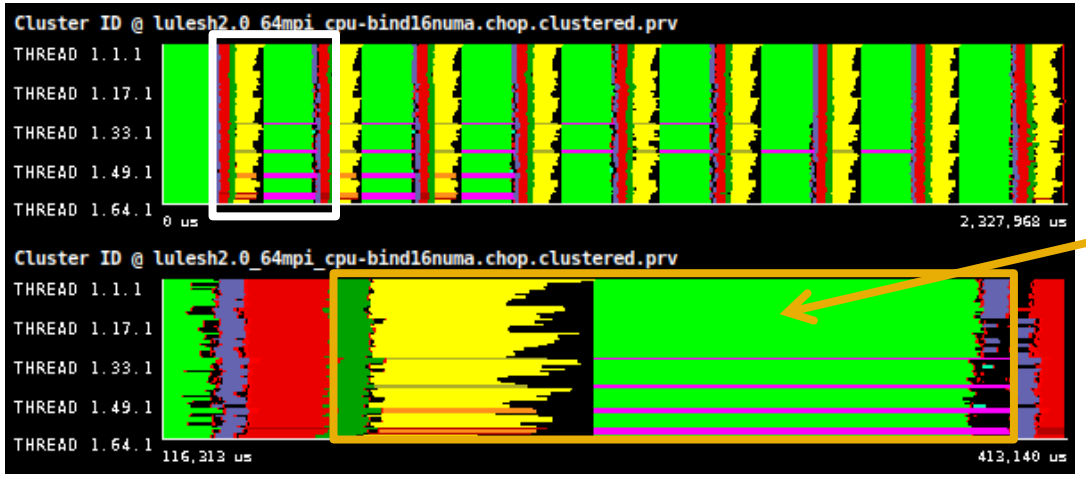

Variable work / speed + Simultaneously @ different processes = Imbalances

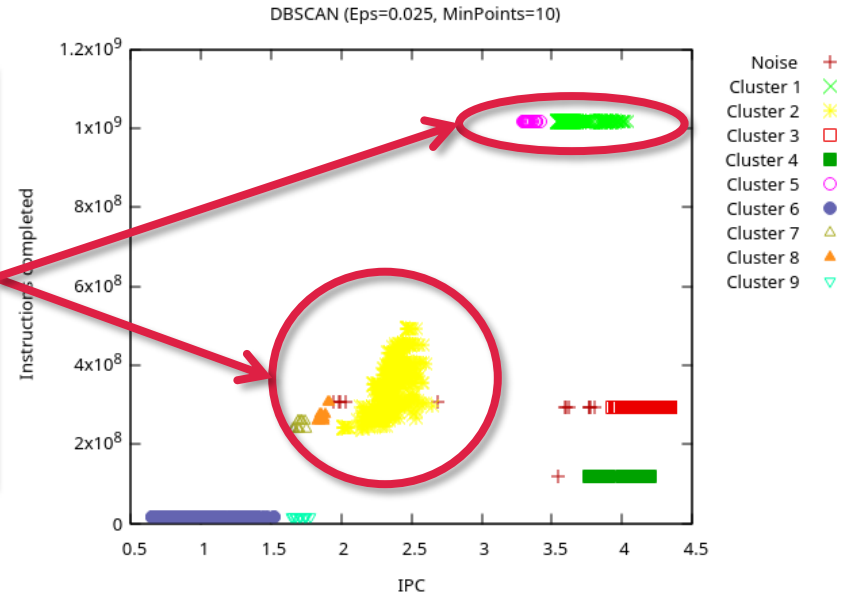

BurstClustering configuration (cluster.xml)

```
<clustering_definition use_duration="no" apply_log="yes"
 normalize_data="yes" duration_filter="0" threshold_filter="0">

  <clustering_algorithm name="DBSCAN">
    <epsilon>0.025</epsilon>
    <min_points>10</min_points>
  </clustering_algorithm>

  <clustering_parameters>
    <single_event apply_log="yes" name="PAPI_TOT_INS">
          ...
    </single_event>
    <mixed_events apply_log="no" name="IPC" operation="/">
          ...
    </mixed_events>
  </clustering_parameters>

  <output_plots all_plots="no">
          ...
  </output_plots>
</clustering_definition>
```

DBSCAN parameters

Trace metrics used as clustering dimensions

Plots with configurable axes

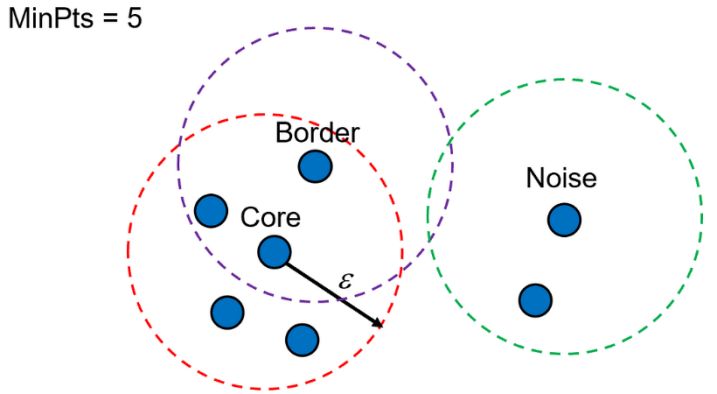

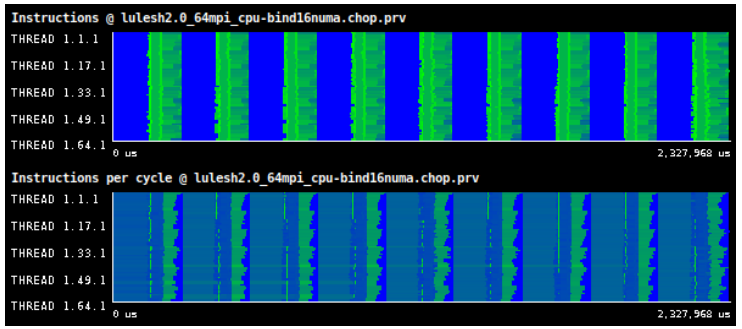

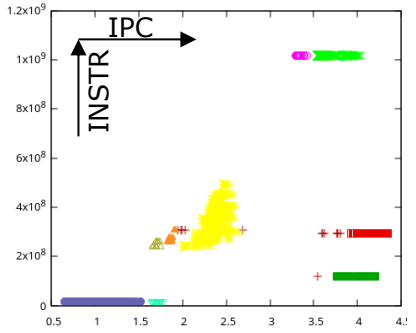

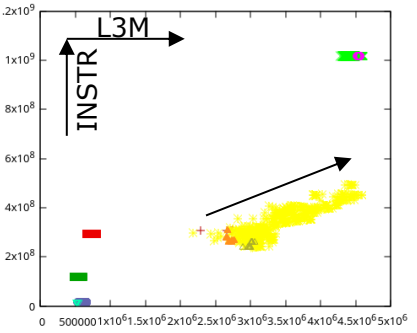

Epsilon

- Which one is better?

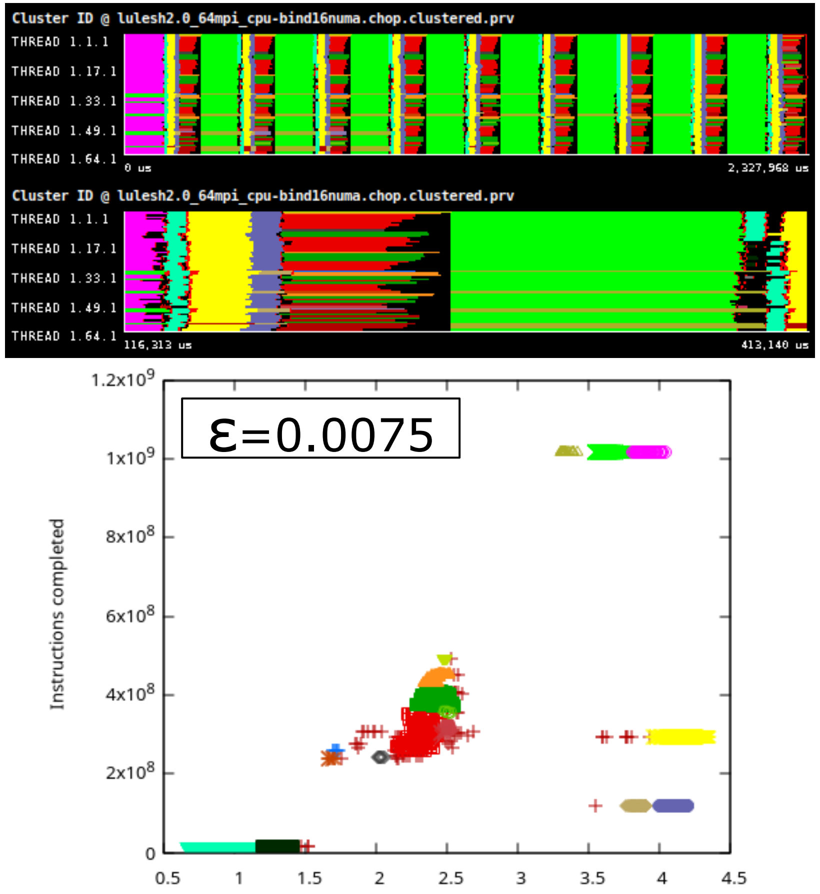

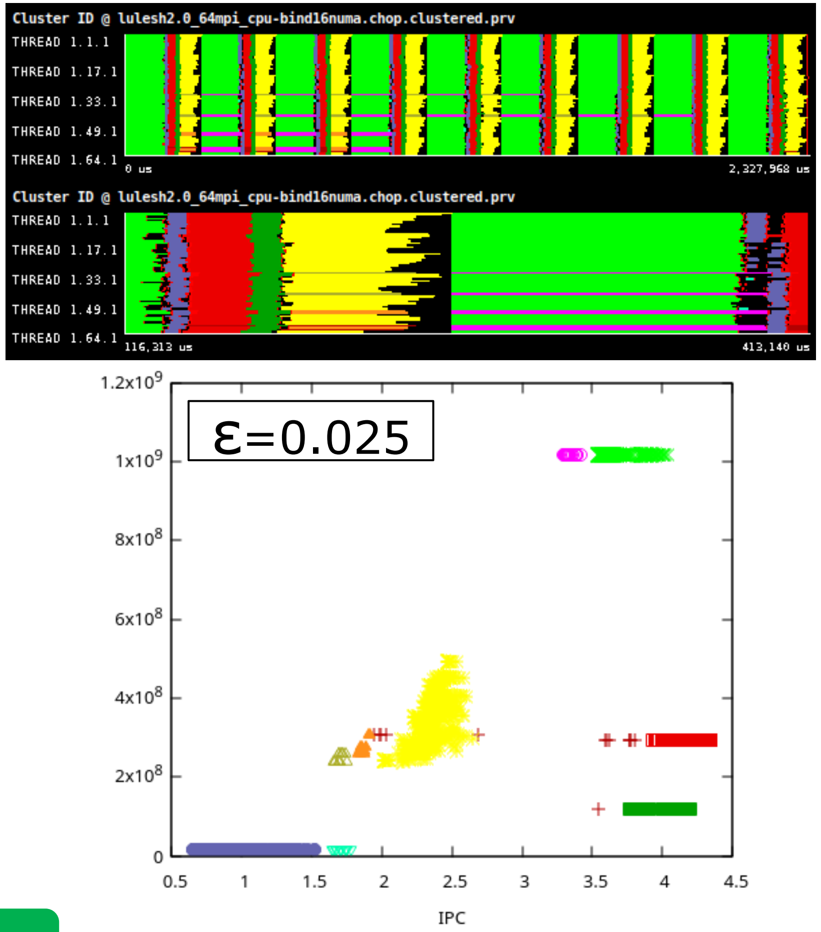

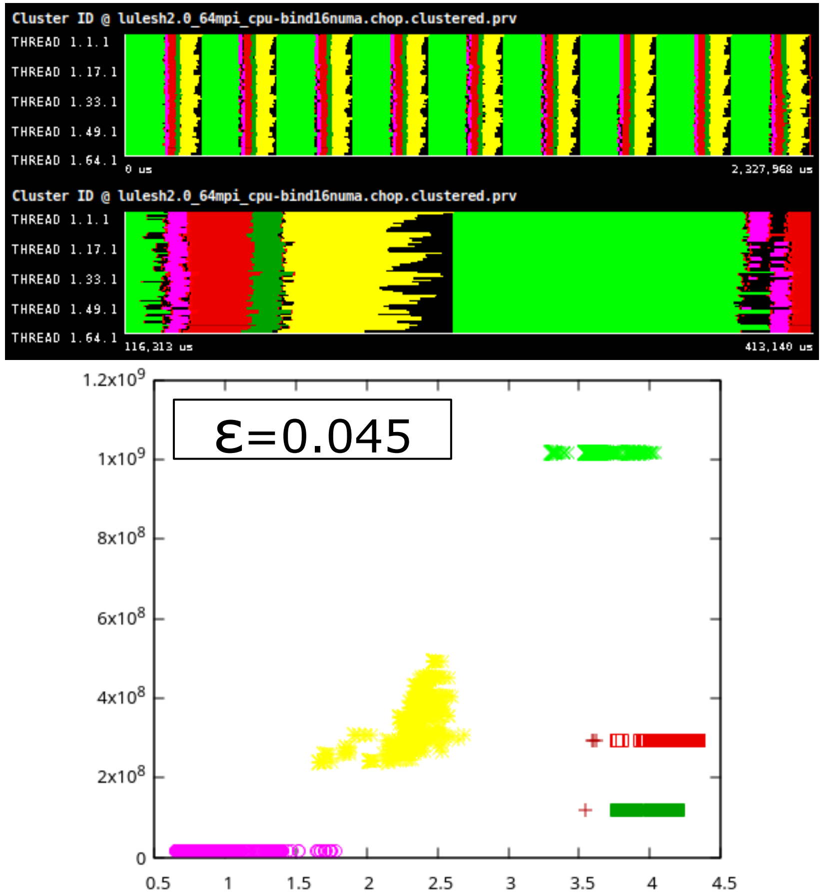

Highlights sub-structure & variability

Highlights SPMD behavior

Comparing scales

- Quick overview of the evolution of the different code regions
  - Clusters keep elongating in the X-axis ➔ IPC variability increases
  - Main computing regions 1 and 2 move to the left ➔ IPC degrades
  - Resource sharing affects Cluster 2 more, yet Cluster 1 remains more relevant