

Visual Analytics for Studying Dragonfly Network Performance

Kelvin Li

University of California, Davis

July 12, 2016

Motivation

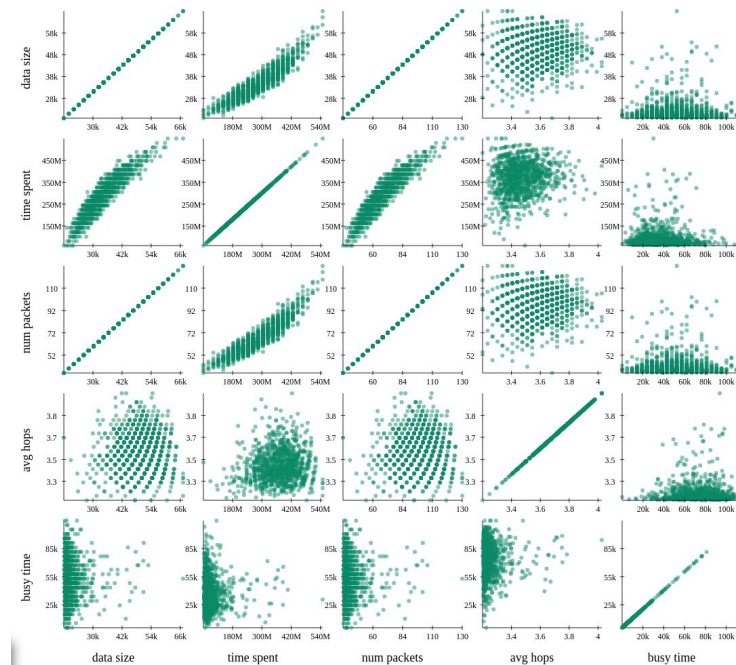
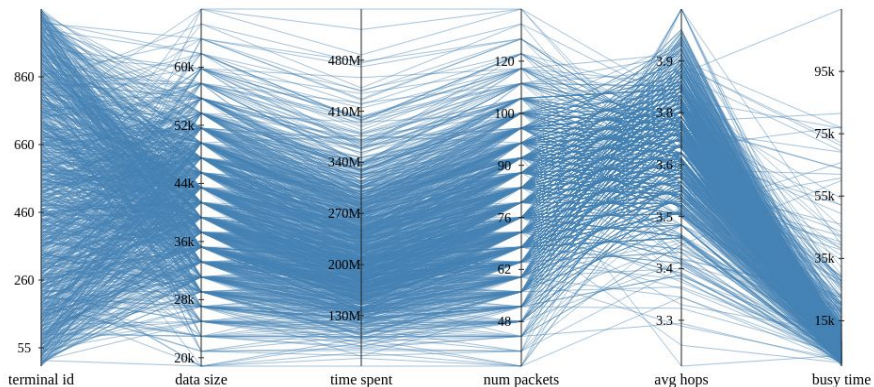
- Our ability to perform simulations has outpaced the techniques for analyzing the simulation results
- Develop a visual analytic tool for scalable data analytics and visualization of large-scale network simulations

Visual Analytics for Exploring Network Simulations

- Understand, analyze, and explore simulation results
- Visualization directed, interactive analysis for studying large-scale network
- Explore the design space of HPC networks
 - Evaluate network performance with different workloads
 - Compare design parameters

CODES Simulation Data

- Multidimensional
- Large
- Complex
- Time-varying data are also available



1056 Nodes Dragonfly Simulation

Dragonfly Network Topology

- g groups (fully-connect), a routers per group, p terminals per router
- Metrics collected from CODES:


Entity	Metric (Unit)
global link	busy time (ns) traffic (byte)
local link	busy time (ns) traffic (byte)
terminal	packets finished data size (byte) busy time (ns) avg. packet latency (ns) avg. hops
router	total local busy time (ns) total global busy time (ns) total local traffic (byte) total global traffic (byte)

Busy time = Saturation time


Circular Layouts for Analyzing Networks

- Effective for showing communication patterns
- High degree of symmetry
 - Check load balancing
 - Detect congestions and identify bottlenecks

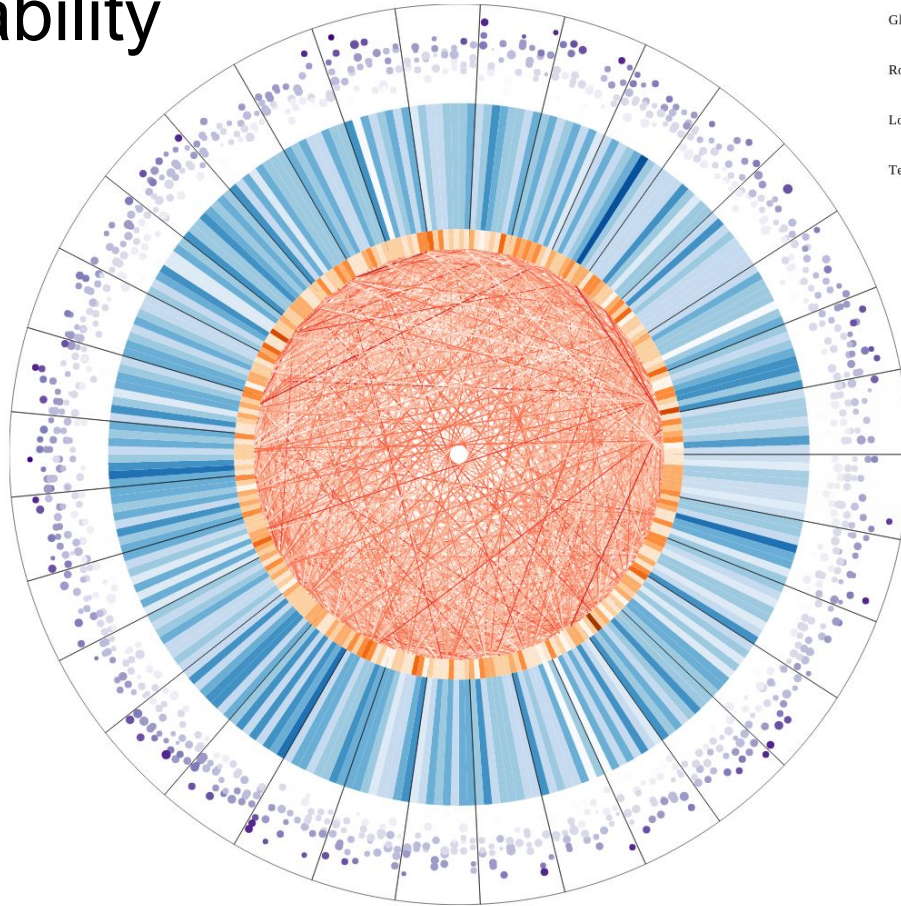
Sigovan et al. 2013



Bhatele et al. 2015



Scalability



- 33 Groups
- 264 Routers (8 per group)
- 1056 Terminals (4 per router)

Visual Analysis of Dragonfly Networks

- Visual aggregate for perceptual scalability
- Interactive visualizations for directing the process of analysis and exploration
- Explore hierarchical structures in Dragonfly Networks
 - Reserve structural properties in data aggregation (Topology-aware)

Aggregation by Groups

- Aggregate multiple Dragonfly groups into one
- Routers and terminals with the same rank in the groups are aggregated
- Performance metrics can be summed up or averaged

Useful for checking global link traffic and load balancing between and within groups.


1056 Nodes, Non-Minimal Routing Uniform Random Traffic

global link traffic




(width of the ribbons = number of routers)

global link busy time




local link busy time




(height of the bars = local link traffic)

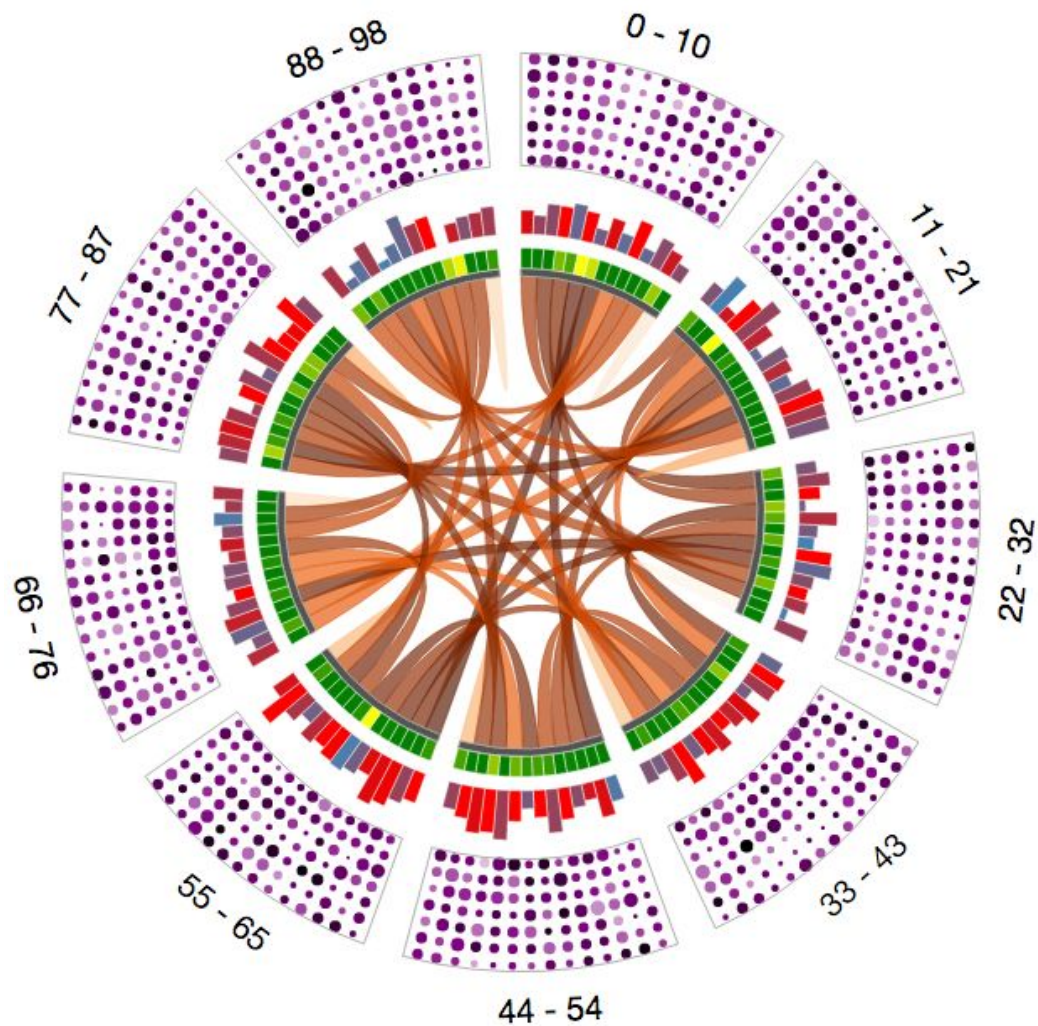
avg. packet latency



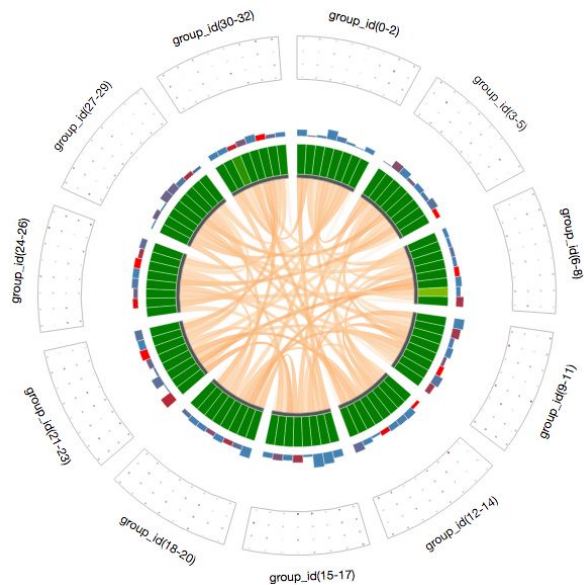
(size of the dots = avg. hops)




9K Nodes
Minimal Routing
Uniform Random Traffic




Comparing Routing Methods with Uniform Random Traffic



Minimal Routing





Non-Minimal Routing




Adaptive Routing

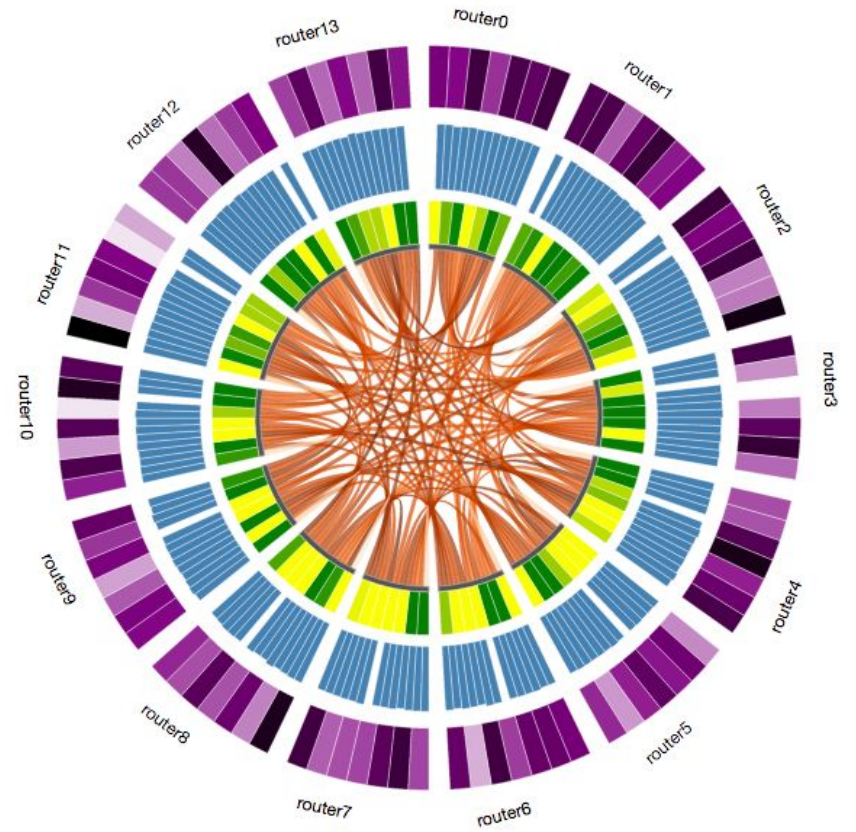
Aggregation by Router Rank



1056 Nodes, Non-Minimal Routing
Uniform Random Traffic



5K nodes



9K nodes


User Defined Visual Mapping

Data Sets:

Group By:

Number of Groups: (8)

Entity	Visual Encoding	Color Mapping
<input type="text" value="router"/> <input type="checkbox"/> aggregate	Color: <input type="text" value="global_traffic"/> Size: <input type="text" value="-----"/> Angular (x): <input type="text" value="-----"/> Radial (y): <input type="text" value="-----"/>	 0 700k 1.4M 2.1M 2.8M
<input type="text" value="router"/> <input type="checkbox"/> aggregate	Color: <input type="text" value="global_traffic"/> Size: <input type="text" value="-----"/> Angular (x): <input type="text" value="-----"/> Radial (y): <input type="text" value="-----"/>	 280k 300k 330k 350k 380k
<input type="text" value="router"/> <input type="checkbox"/> aggregate	Color: <input type="text" value="local_busy_time"/> Size: <input type="text" value="local_traffic"/> Angular (x): <input type="text" value="-----"/> Radial (y): <input type="text" value="-----"/>	 0 20 40 60
<input type="text" value="terminal"/> <input type="checkbox"/> aggregate	Color: <input type="text" value="data_size"/> Size: <input type="text" value="busy_time"/> Angular (x): <input type="text" value="-----"/> Radial (y): <input type="text" value="avg_packet_latency"/> <input type="text" value="avg_hops"/>	 0 20k 40k 60k 80k



Visualizing Network Simulation


- Focus on the aspects based on the analysis tasks in hand
- Customize visual components by choosing entities and metrics
- Better use of visualization to direct data exploration and analyses

Visual Analysis of Time-varying Data


- Analyze temporal patterns of the workload traffic and its effect on network performance
- Exploring both structural and temporal properties
 - Interactive visual analysis for highlighting connections

Time-varying Data


- Stacked time-series plots for comparing different performance metrics
- Aggregate data to provide different levels of granularity




System



Dragonfly Group

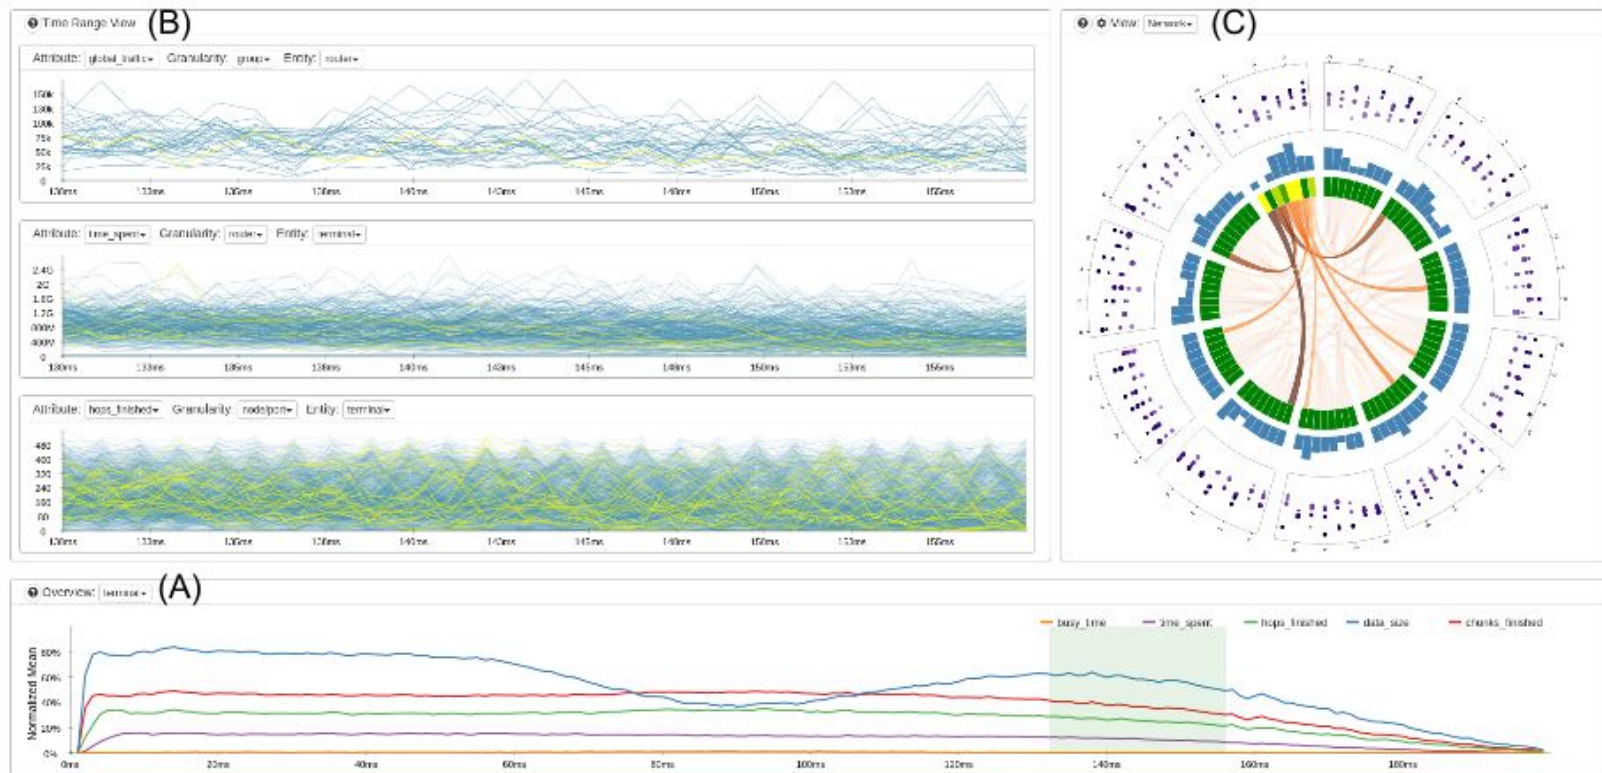


Router Group

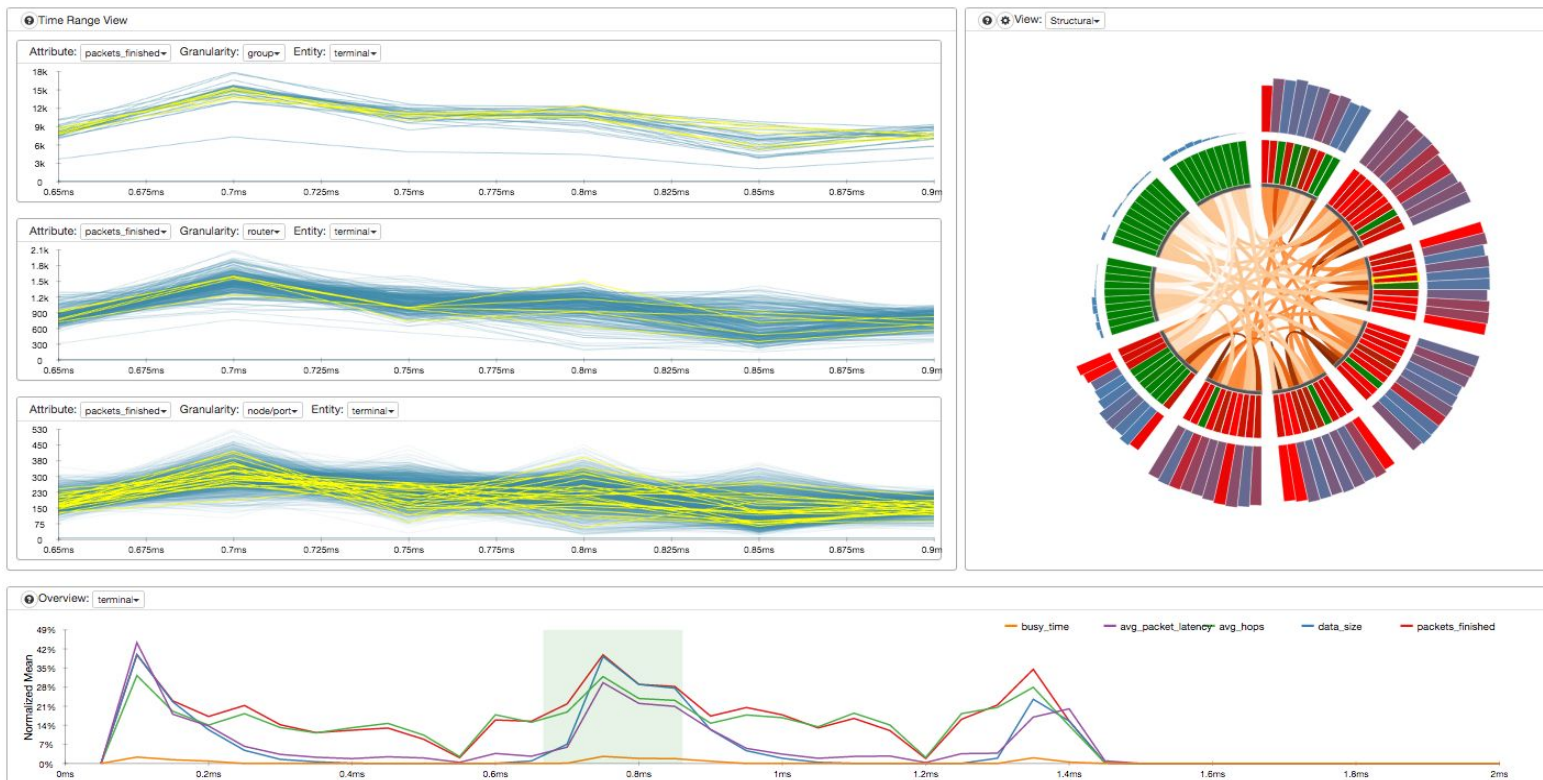


Compute Node

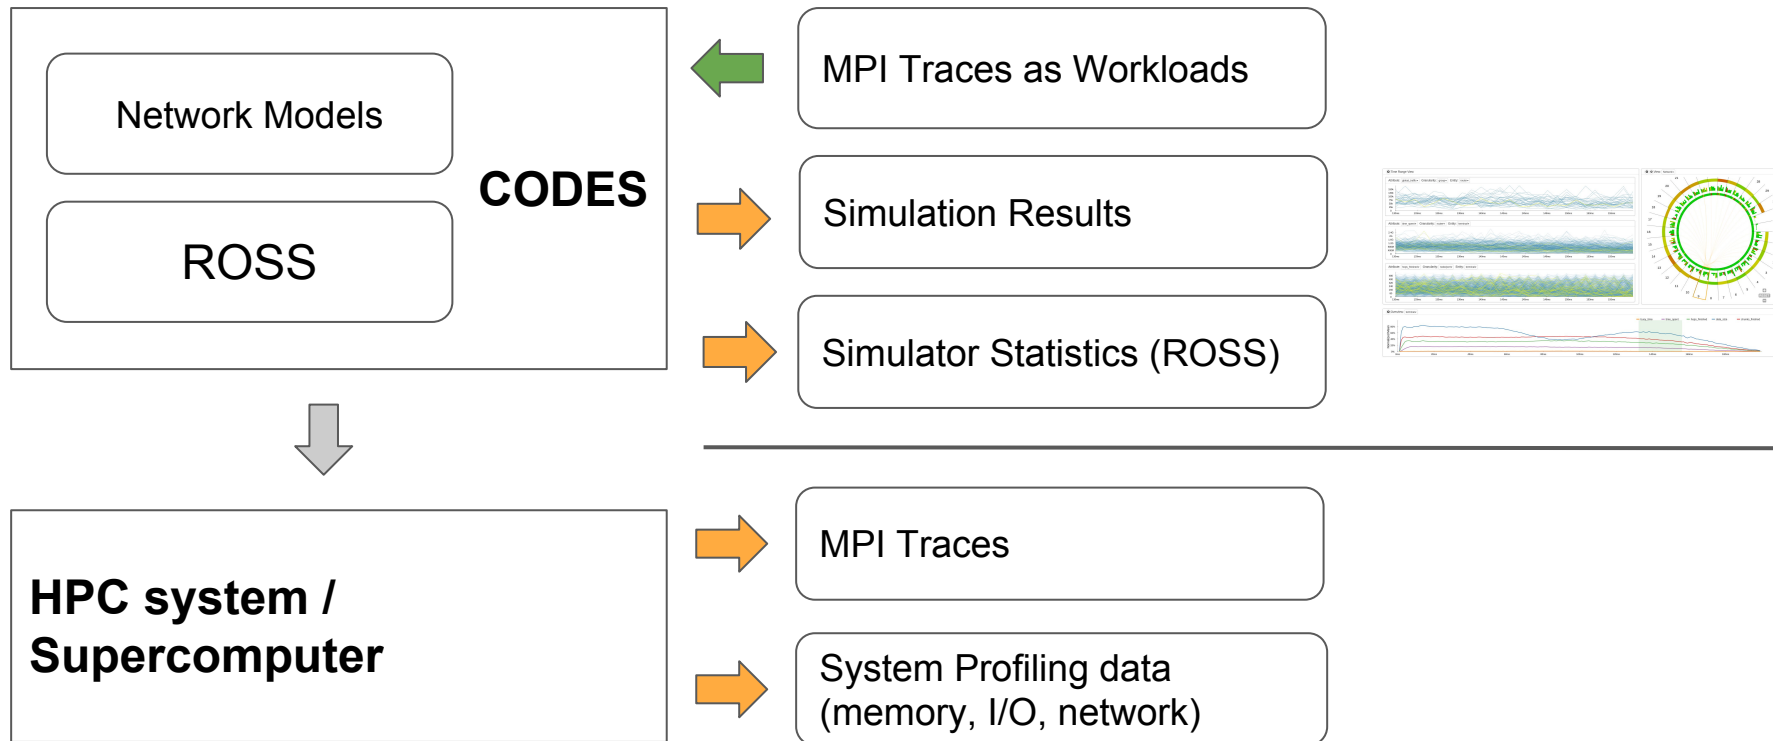
Visual Interface for Exploring Time-varying Data



AMG 1728 MPI rank on 2.5K Dragonfly network



Network and Performance Data in CODES



Future Development

- Data analytics and visualization:
 - Alternative layouts for more network models
 - Improve the user interface
 - More data mining techniques
- Network and Performance Analysis:
 - Analyze simulator performance data in ROSS
 - Study other network models
 - Support analyses and explorations for different levels of data in CODES
 - Explore the connections between model-level and simulator-level properties

Thank You!