IEEE PACIFIC VIS

A Visual Analytics Framework for Reviewing Streaming Performance Data

Suraj P. Kesavan[1] , Takanori Fujiwara[1], Jianping Kelvin Li[1], Caitlin Ross[2], Misbah Mubarak[3], Christopher D. Carothers[2], Robert R. Ross[3], Kwan-Liu Ma[1]

University of California, Davis, USA[1]
Rensselaer Polytechnic Institute, Troy, USA[2]
Argonne National Laboratory, Chicago, USA[3]

Rensselaer

Argonne NATIONAL LABORATORY

VISUALIZATION & INTERFACE DESIGN INNOVATION

Performance Analysis of Large-Scale Scientific Simulations

HPC compute nodes

Ensure Performance Efficiency

Maximize Utilization of Resources

Post-Hoc data analysis workflow

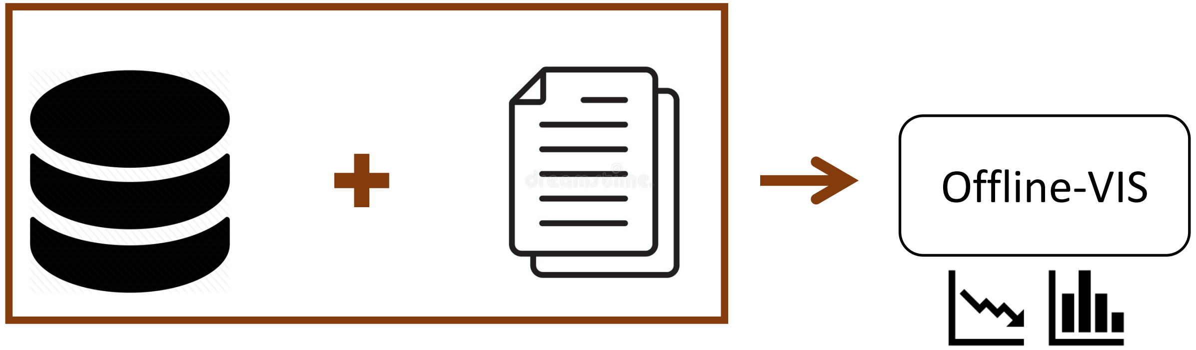

Unviable for high **volume** and **variety**

Performance Data

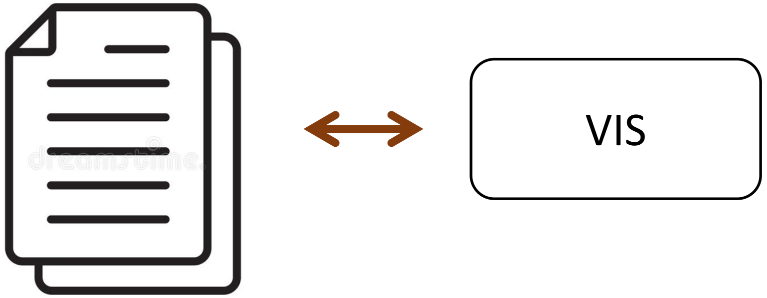

In-Situ data analysis workflow

In-situ Performance Analysis for HPC is challenging.

Data Streams Complexity

Multiple performance metrics (heterogeneous) are collected from **multiple compute nodes**.

Active Monitoring

Key changes and patterns can occur anytime.

Situational Awareness

Heterogeneous data streams make it difficult to understand **causal relationships**.

Contributions

Build a **progressive visual analytic framework** that helps perform real-time analysis of streaming performance metrics and communication data

Data Management Module

- Handle **high volume** and **velocity** HPC streaming data.

Analysis Module

- Identify similarities and (dis)similarities from **temporal behavior patterns**.
- Detect **key changes** that deviate from a baseline behavior.
- Derive **causal relationships** among performance metrics.

Visualization Module

- Enable users to analyze the performance metrics **interactively**.
- Provide visualizations to analyze **communication data** along with **performance behaviors**.

Related Work

Performance of large-scale networks in Supercomputing systems [Aaditya et. al, 2012; Fujiwara et. al, 2018]

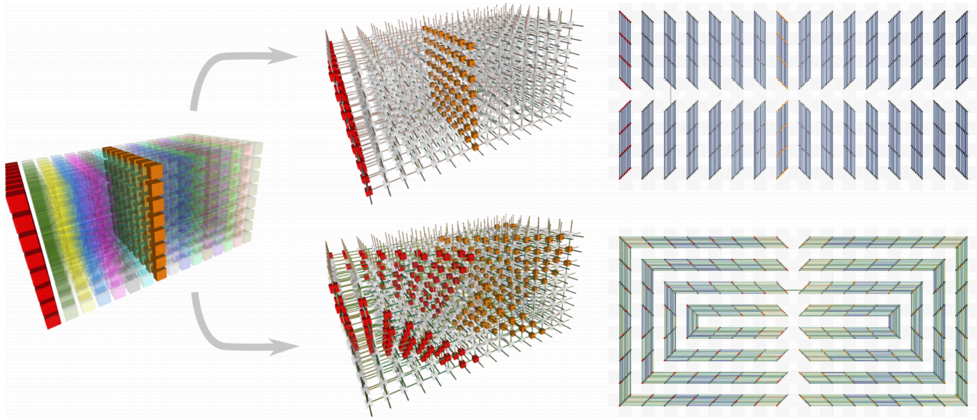

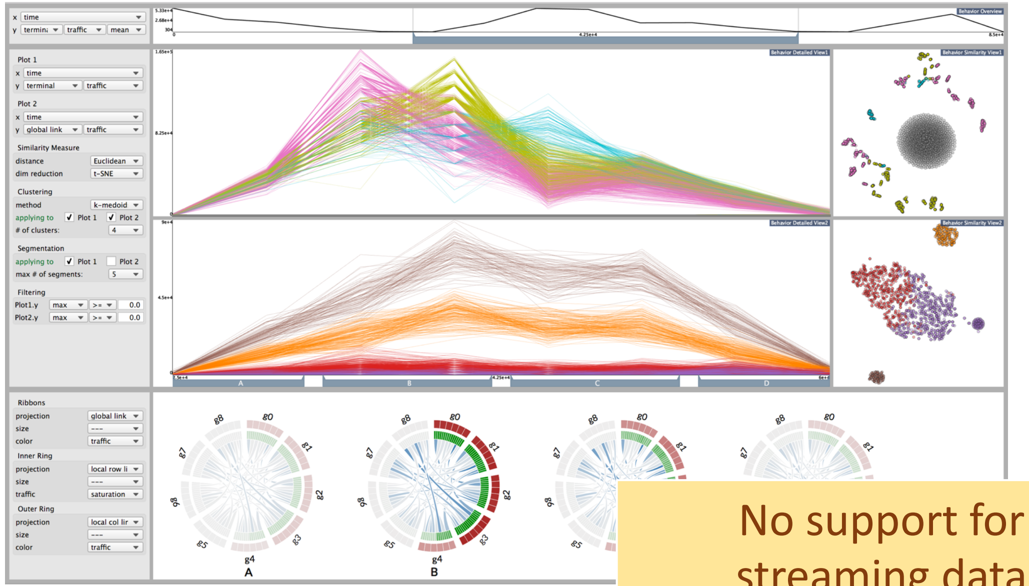

No support for streaming data

Progressive Visual Analytics Approximate-tSNE [Pezzotti et. al, 2017]

Characteristics of HPC Performance Data

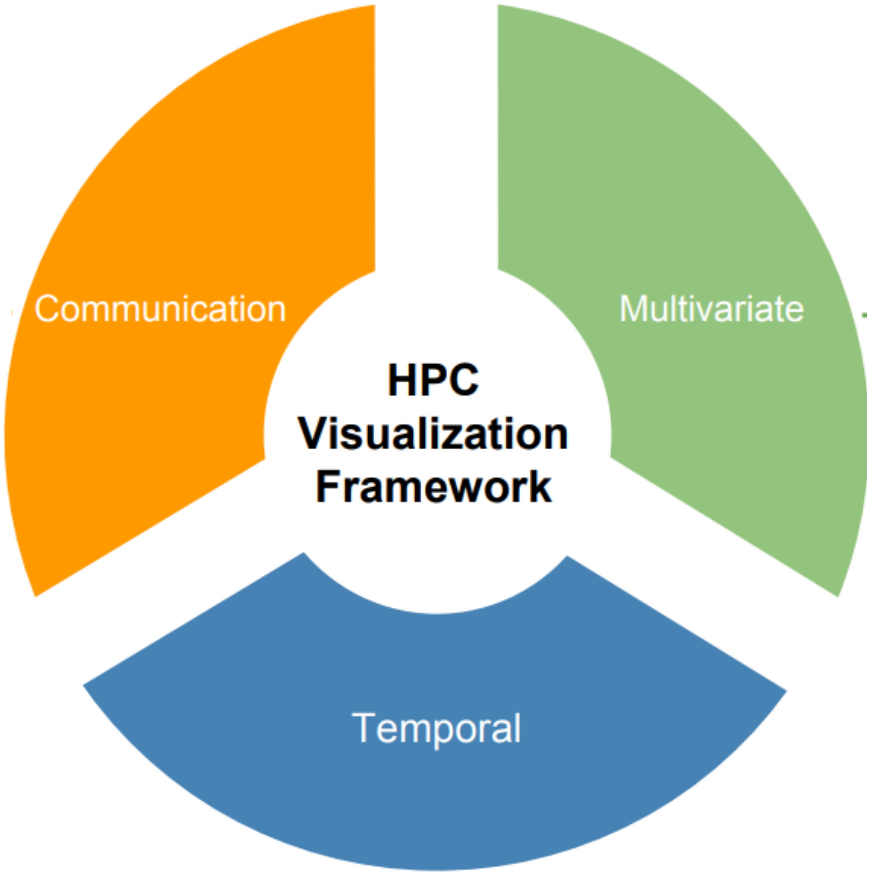

Data Management module

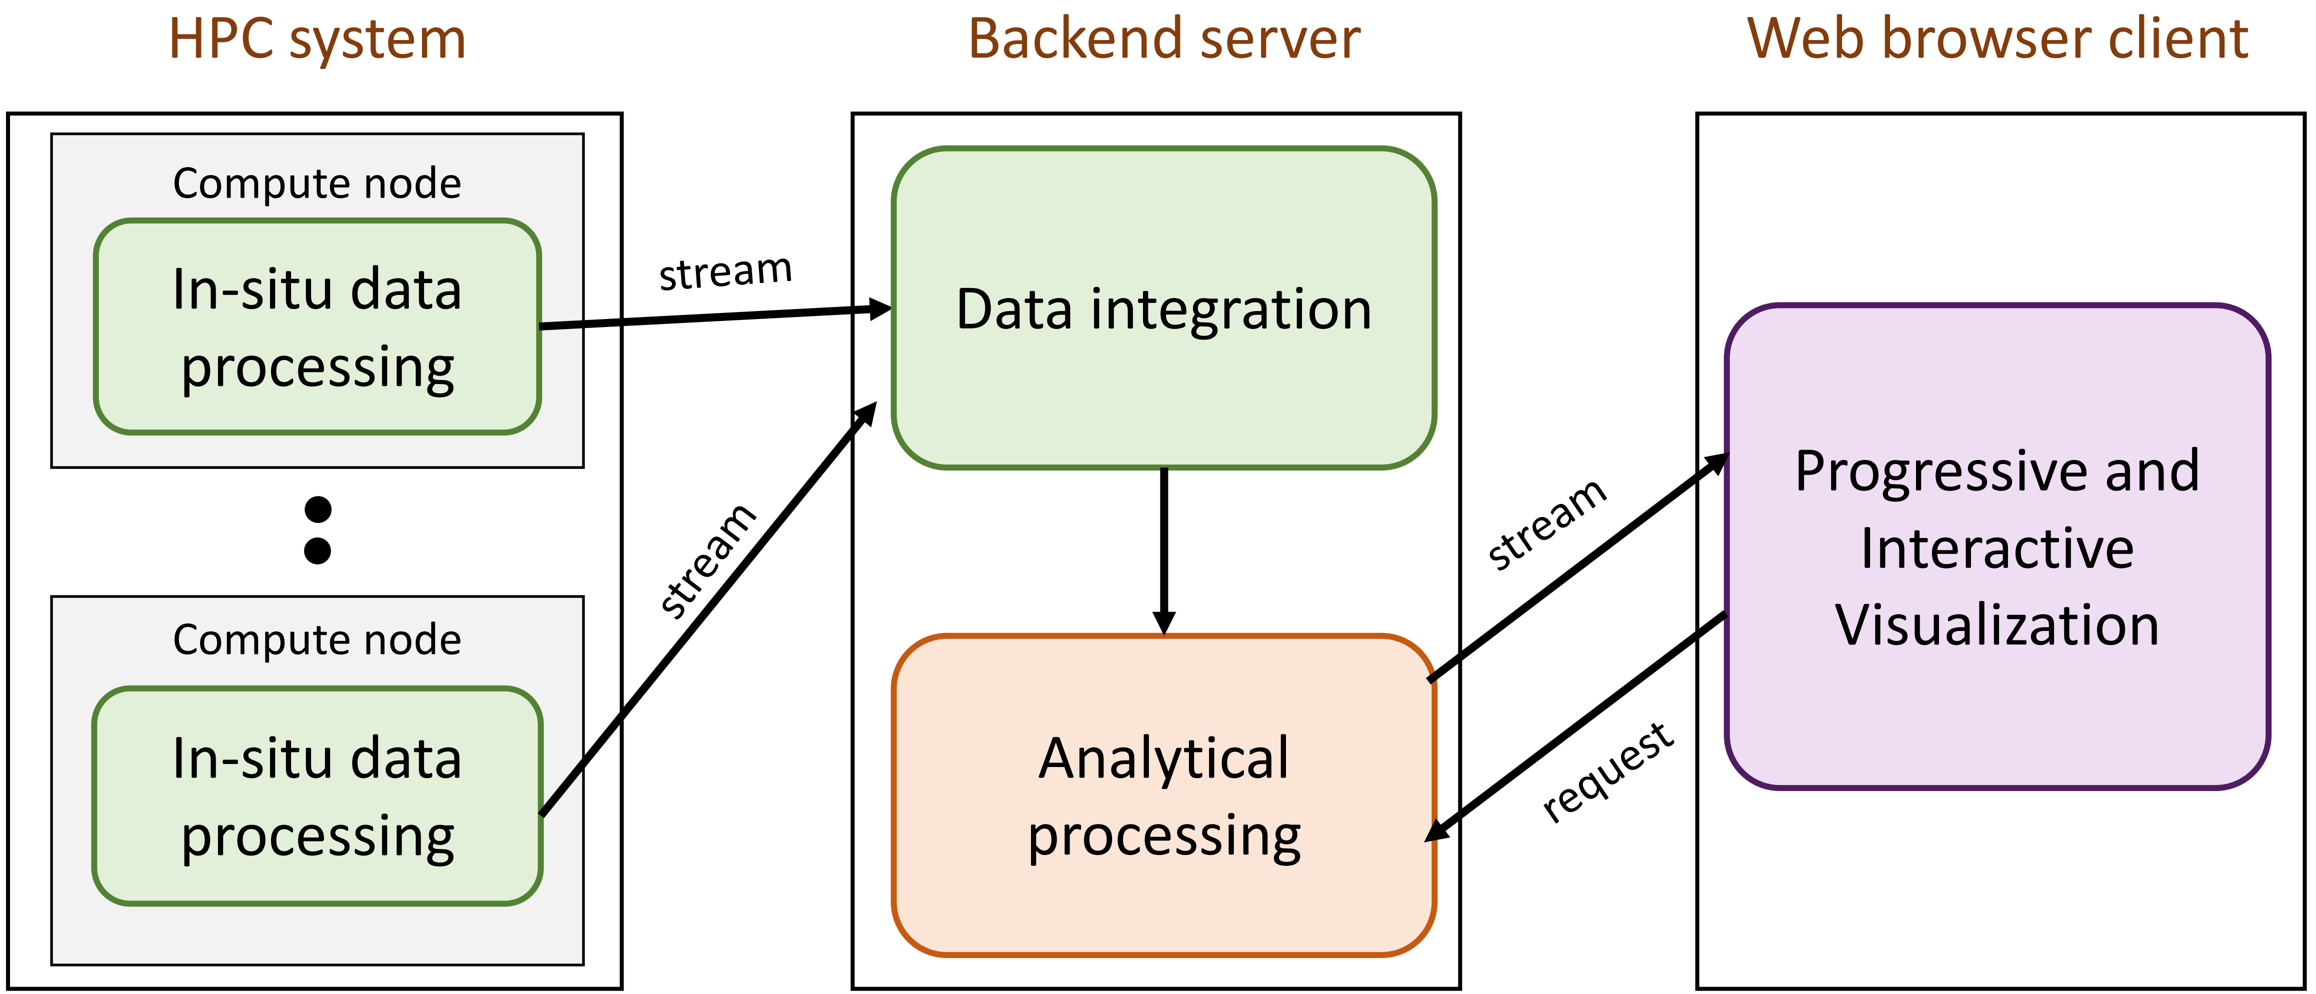

Analysis Module: Identify **similarities** among temporal behaviors.

Sustain active monitoring

Enable situational awareness

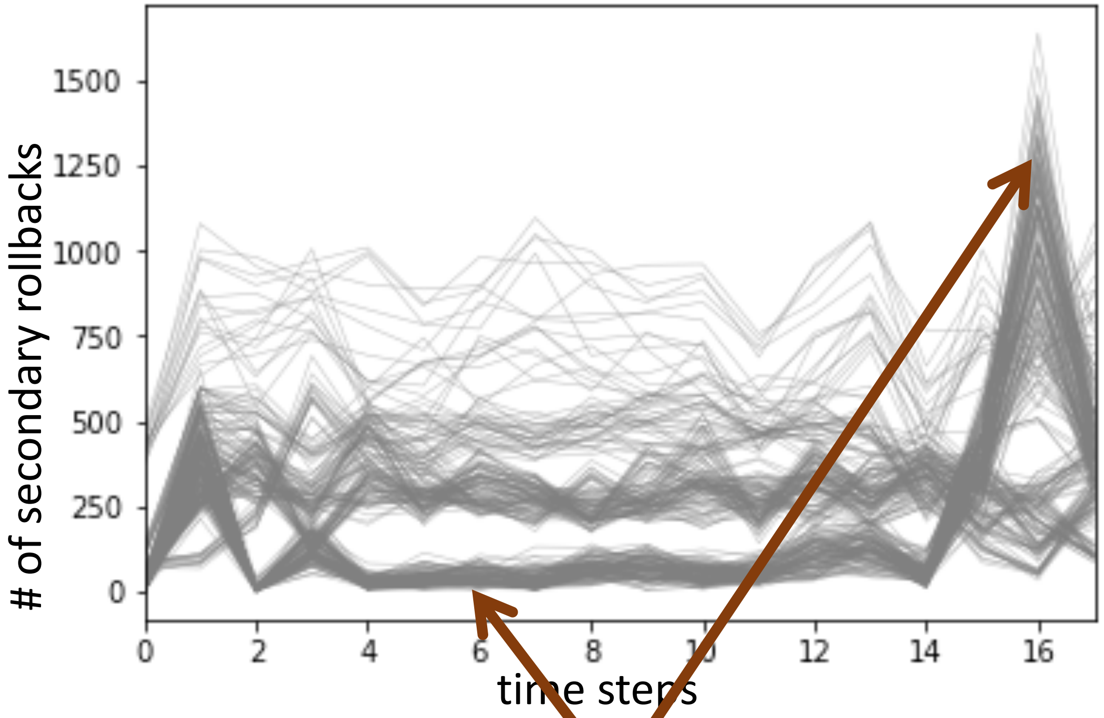

Mini-batch K-means

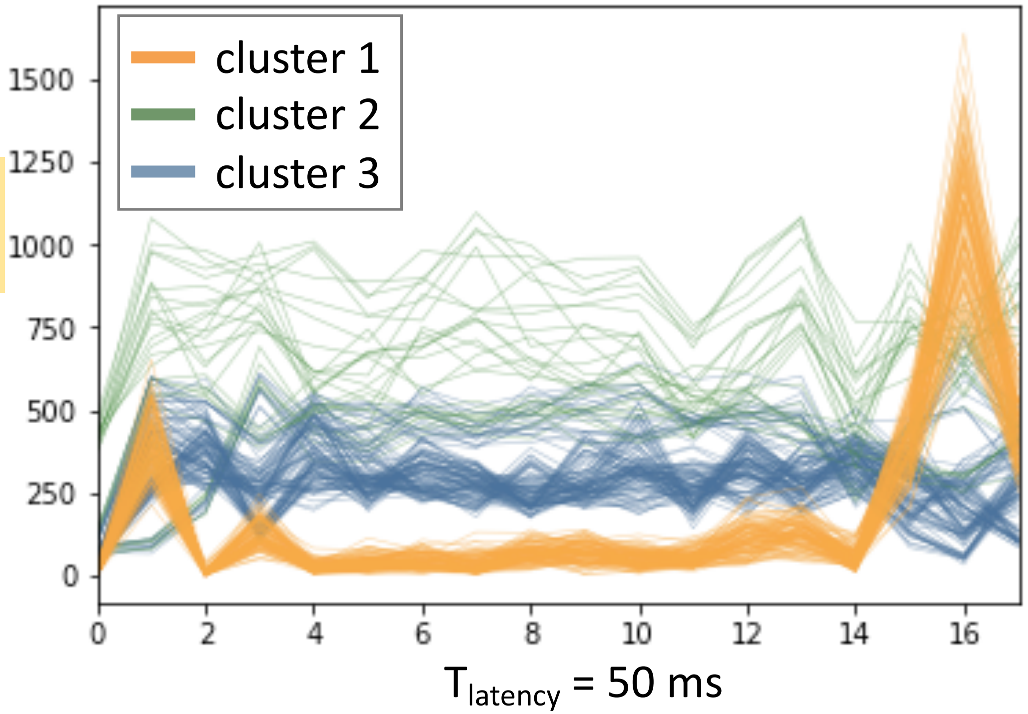

Performance behavior change

This handles latency issue.
“m” entities (m << n) are randomly selected to create ”k” clusters.

Analysis Module: Consistent Cluster Assignment.

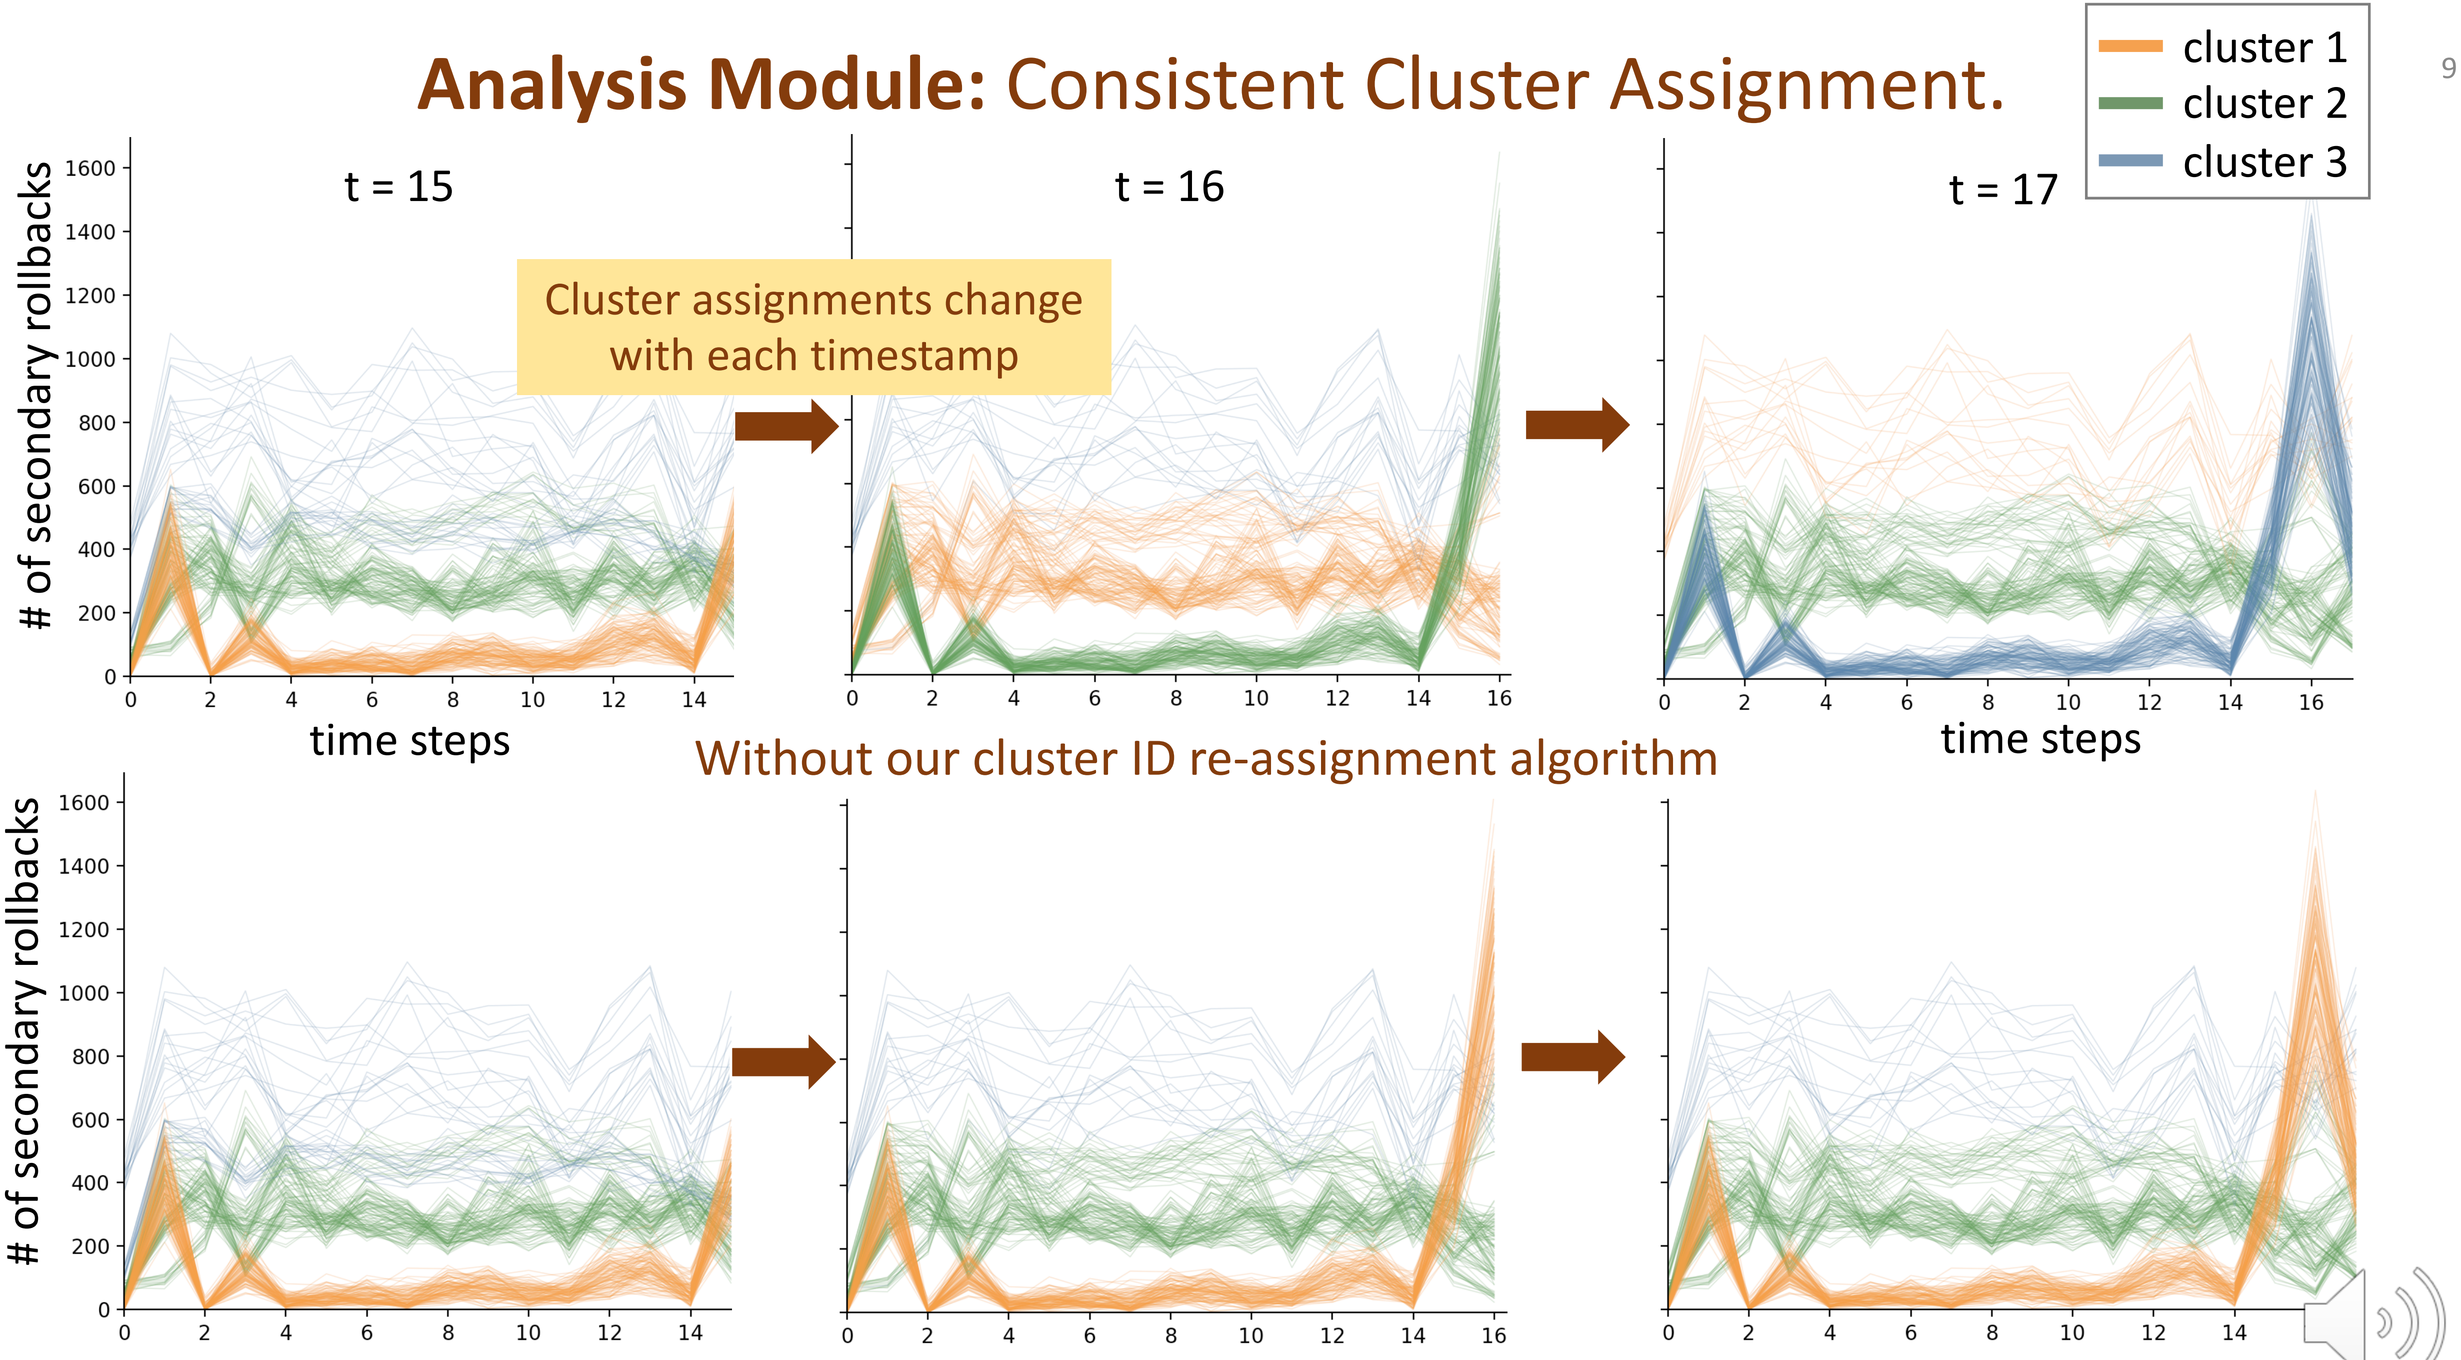

With our cluster ID re-assignment algorithm

Analysis Module: Identify **(dis)similarities** among temporal behaviors.

Progressive Dimensionality Reduction (IPCA by [Ross et al., 2008])

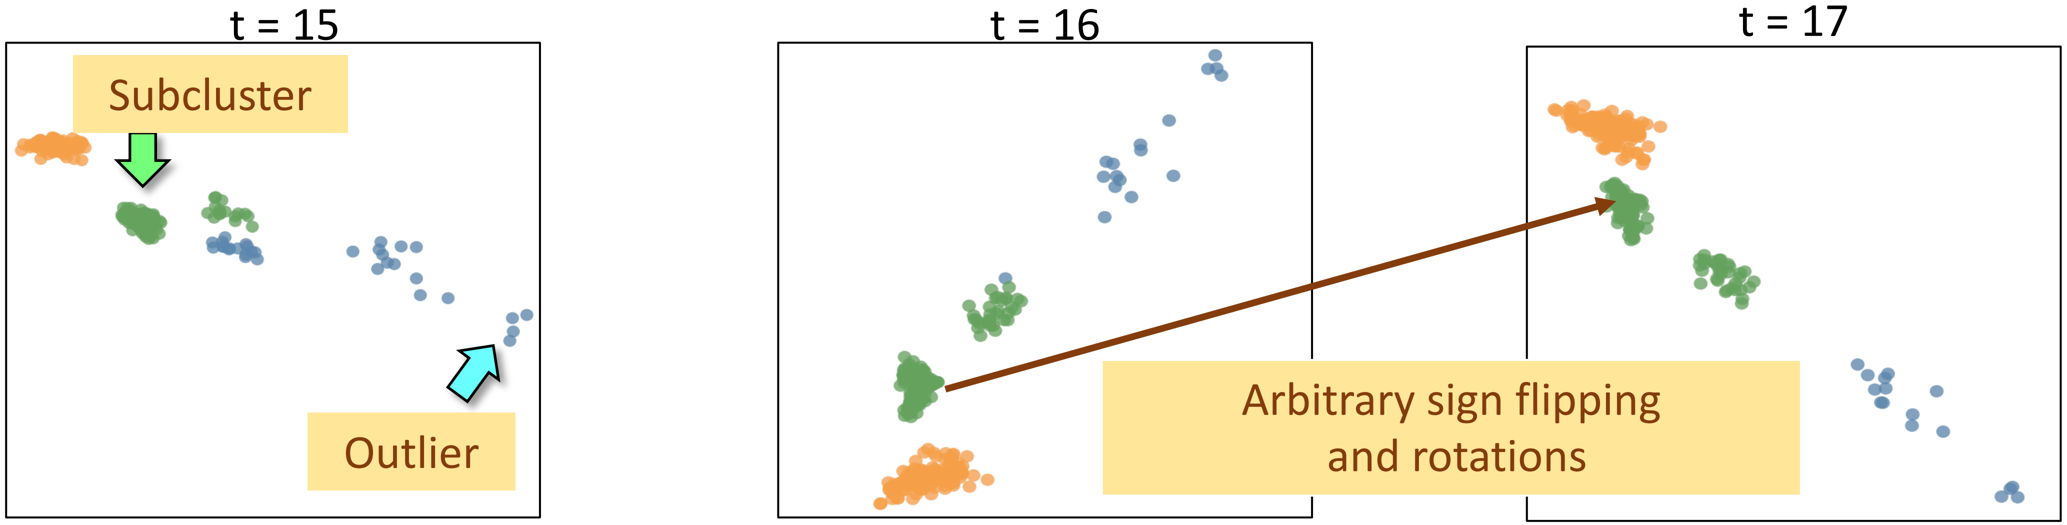

Without the Procrustes Transformation

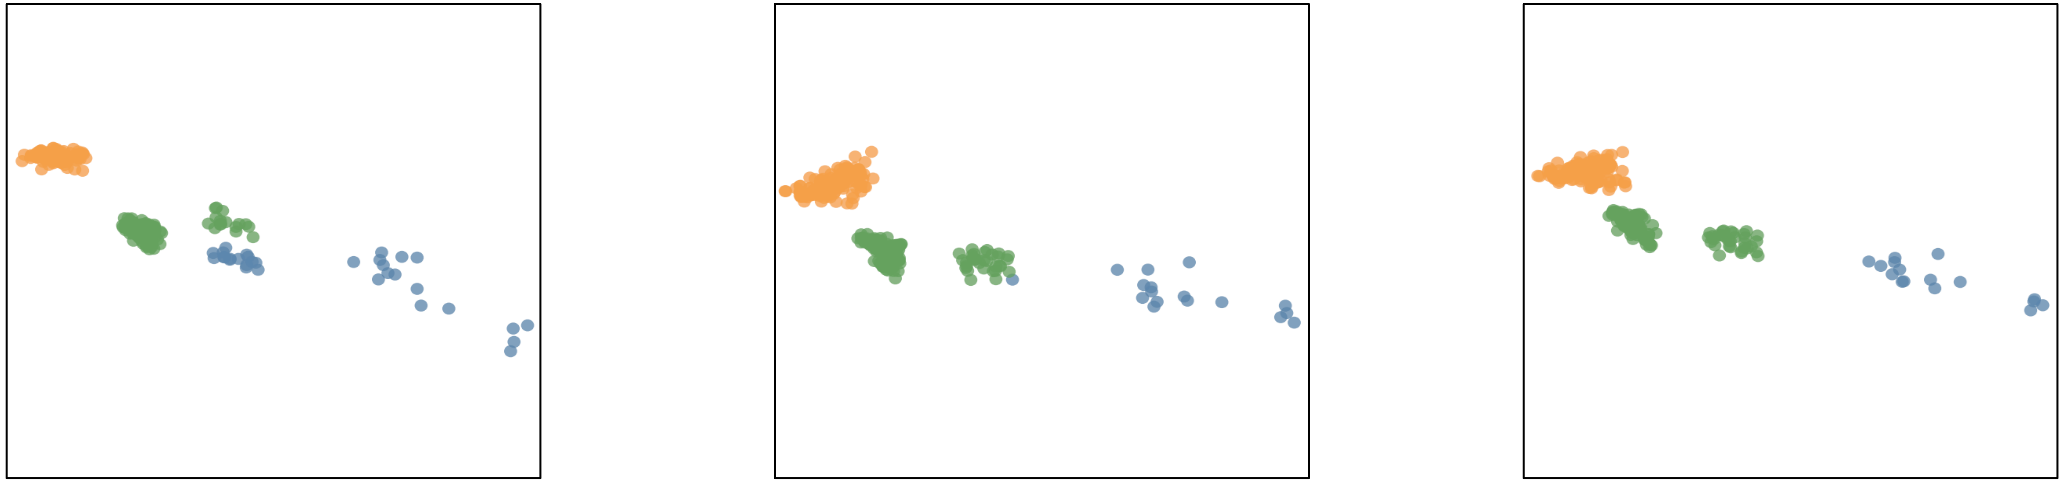

With the Procrustes Transformation

Analysis Module: Detect **key changes** that deviate from a baseline behavior.

Change Point Detection using Adaptive Estimation with Forgetting Factor (AFF)

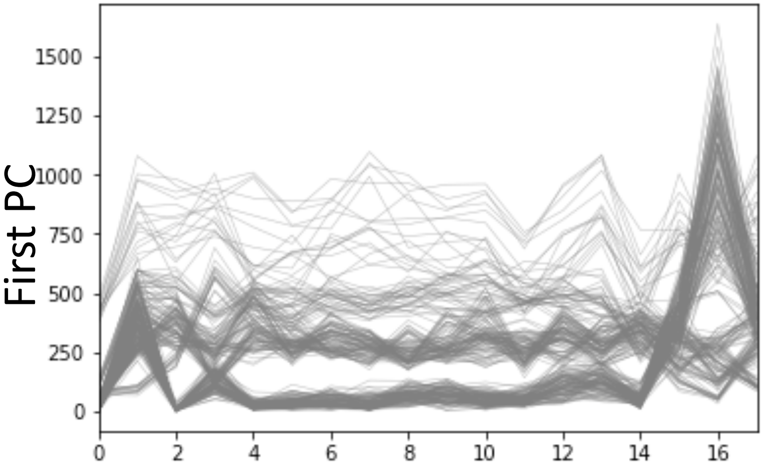

~~PCA~~

IPCA

Representative time series D_t
[Qahtan et. al., 2015]

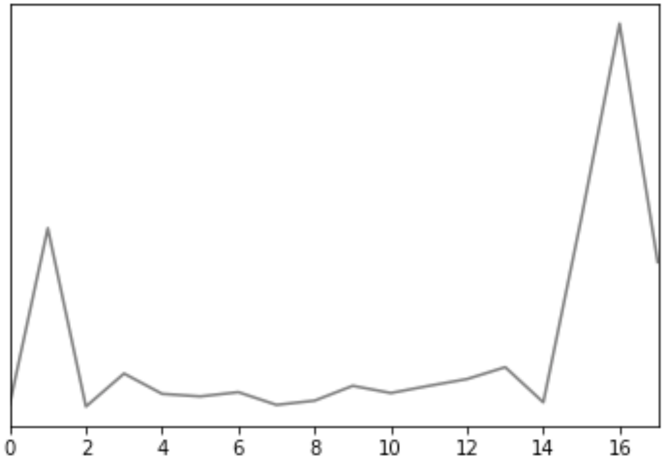

AFF

Significance level α = 0.01
[Bodenham and Adams, 2017]

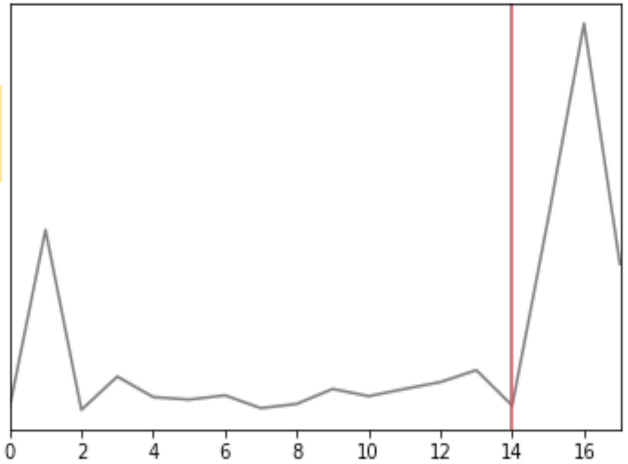

PCA cannot update the PC in real-time

Minimal number of parameters to be specified

Analysis Module: Detect key changes that deviate from a baseline behavior.

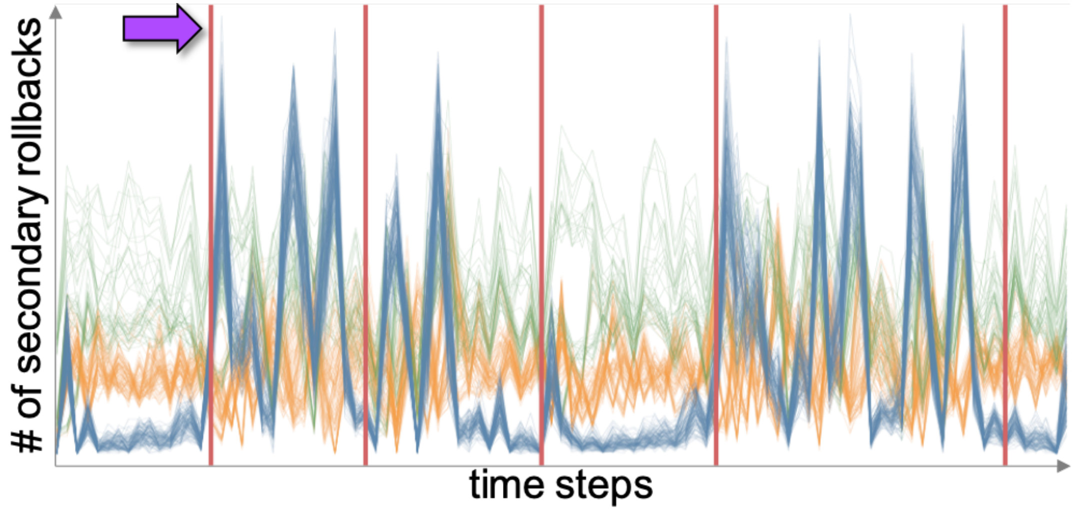

Original Multidimensional Data Streams

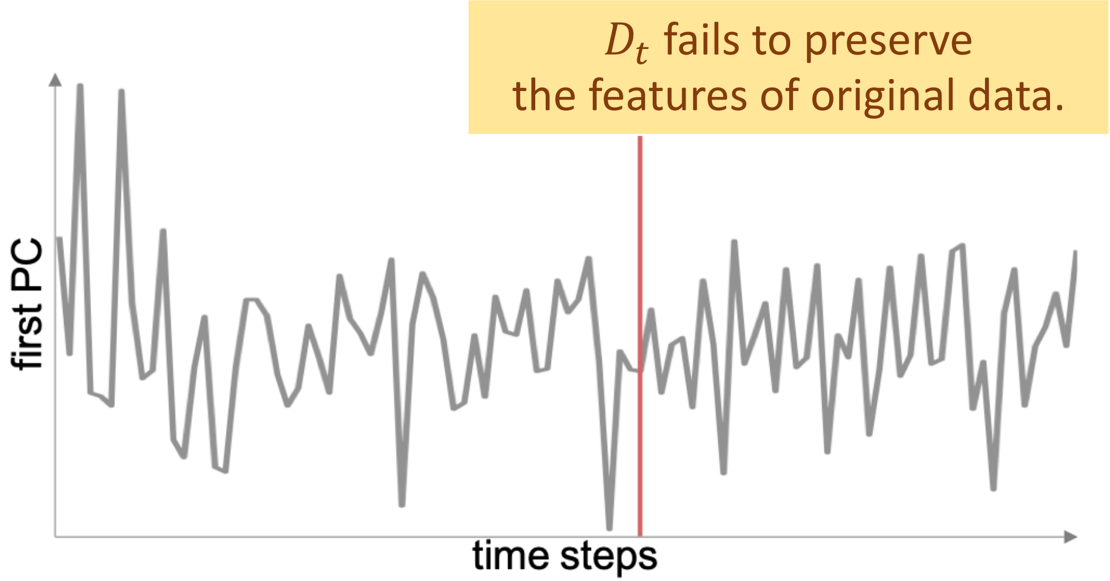

With ordinary PCA

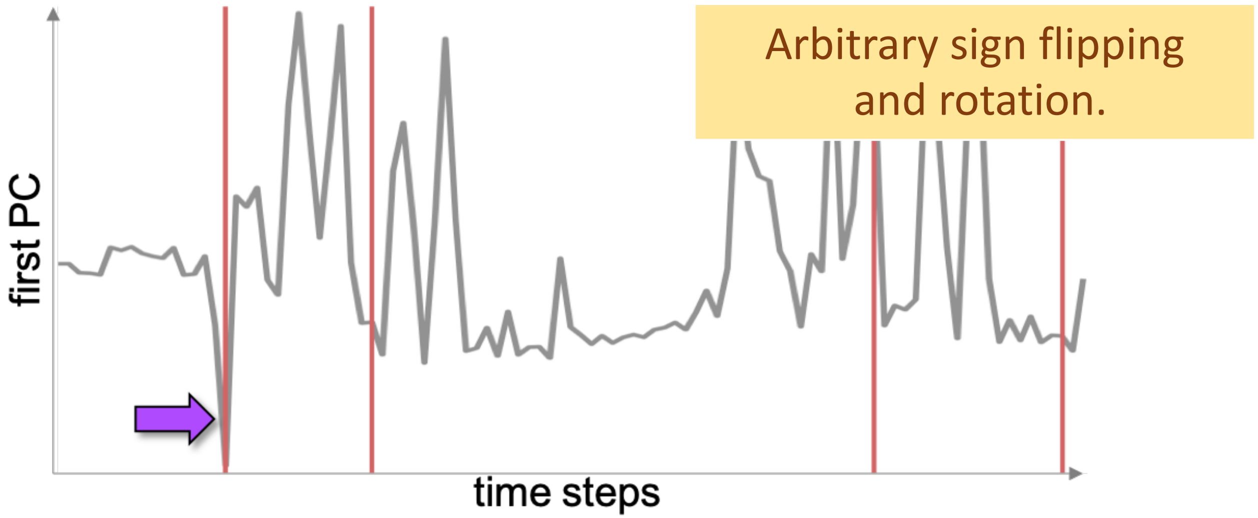

With incremental PCA (IPCA)

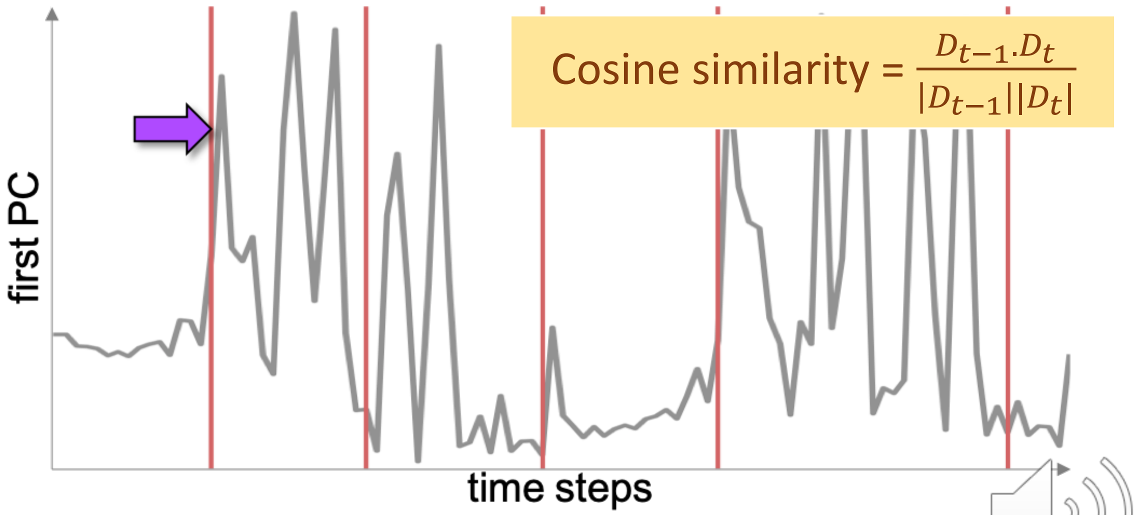

With IPCA and sign adjustment

Analysis Module: Derive **causal relations** among performance metrics.

Granger causality test provides whether a Granger causality between two time series with **p-value**.

"A time series X_t **causes** another time series Y_t, if **present Y can be predicted** better by **using past values of X** than by not doing so"

Granger's causality cannot measure how much one time series affects another time series

Impulse Response function (IR) determines how much **shock to a variable of interest** can affect other variables.

Variance Decomposition (VD) determines the **contribution of a shock** to the **variance of the forecast error** of other variables.

Vector Autoregression (VAR) model [Hamilton, 1994]

VAR fitting is computationally expensive ($O(d^2)$)

↓

$$s = d\sqrt{{(t_{latency} - t_{completion})}/{t_{completion}}}$$

$s \ll d$ (s = 10)

Visualization Module: System overview

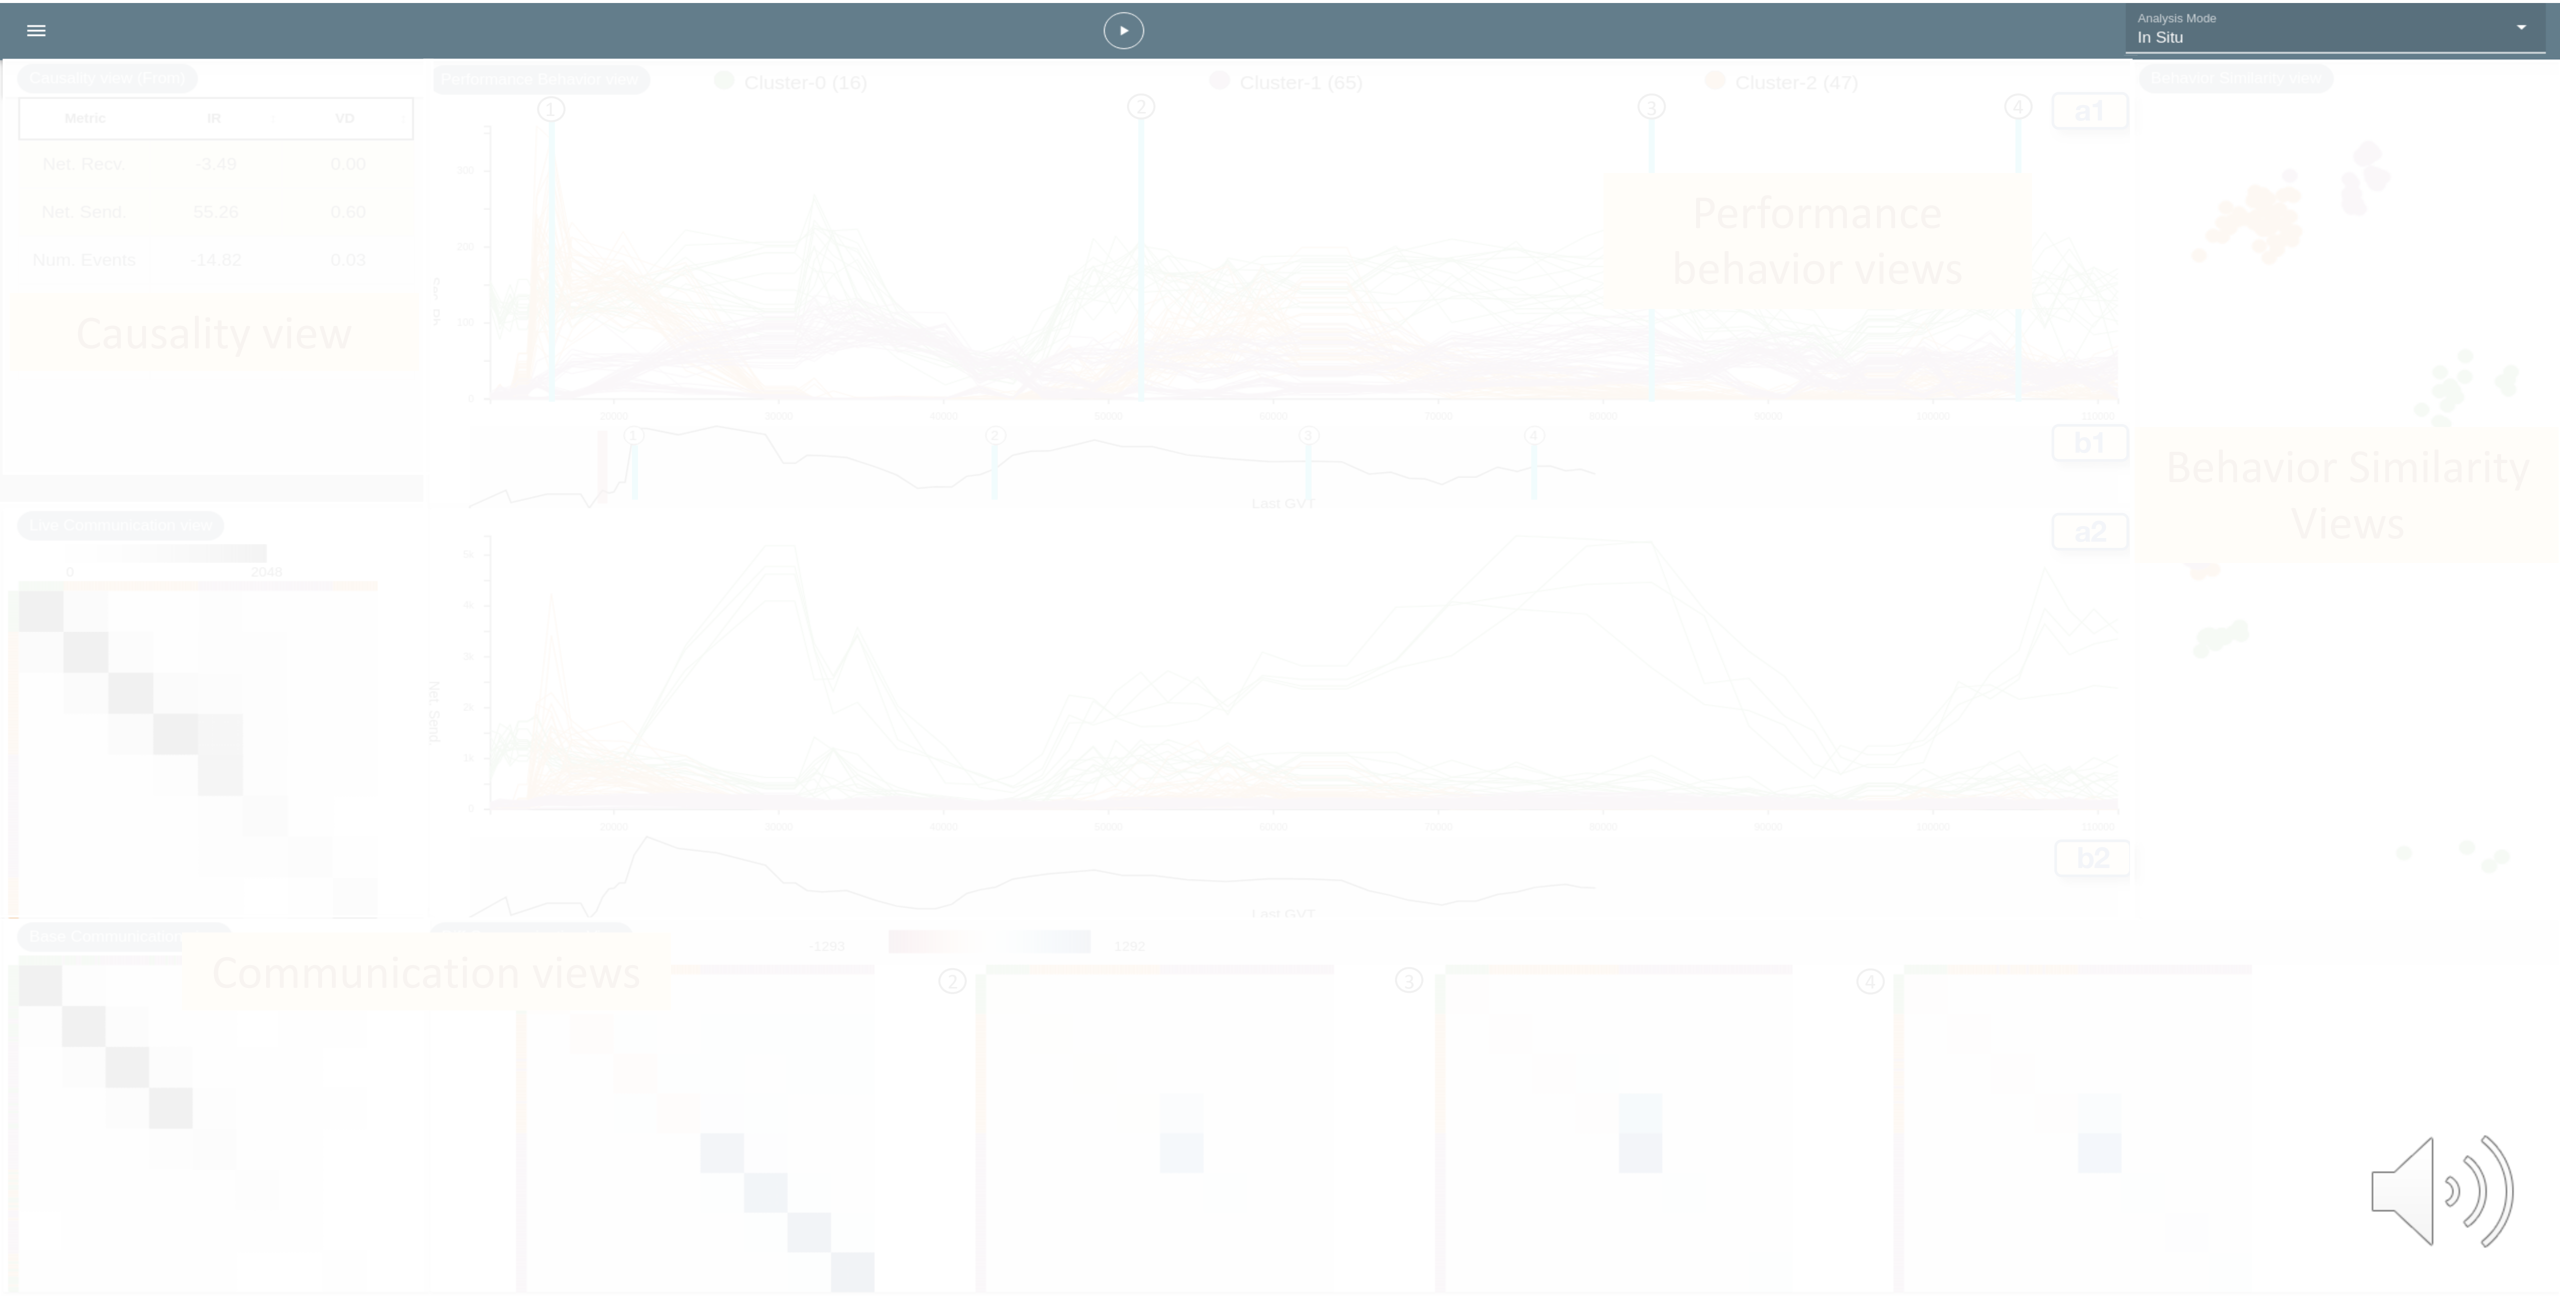

Performance behavior view

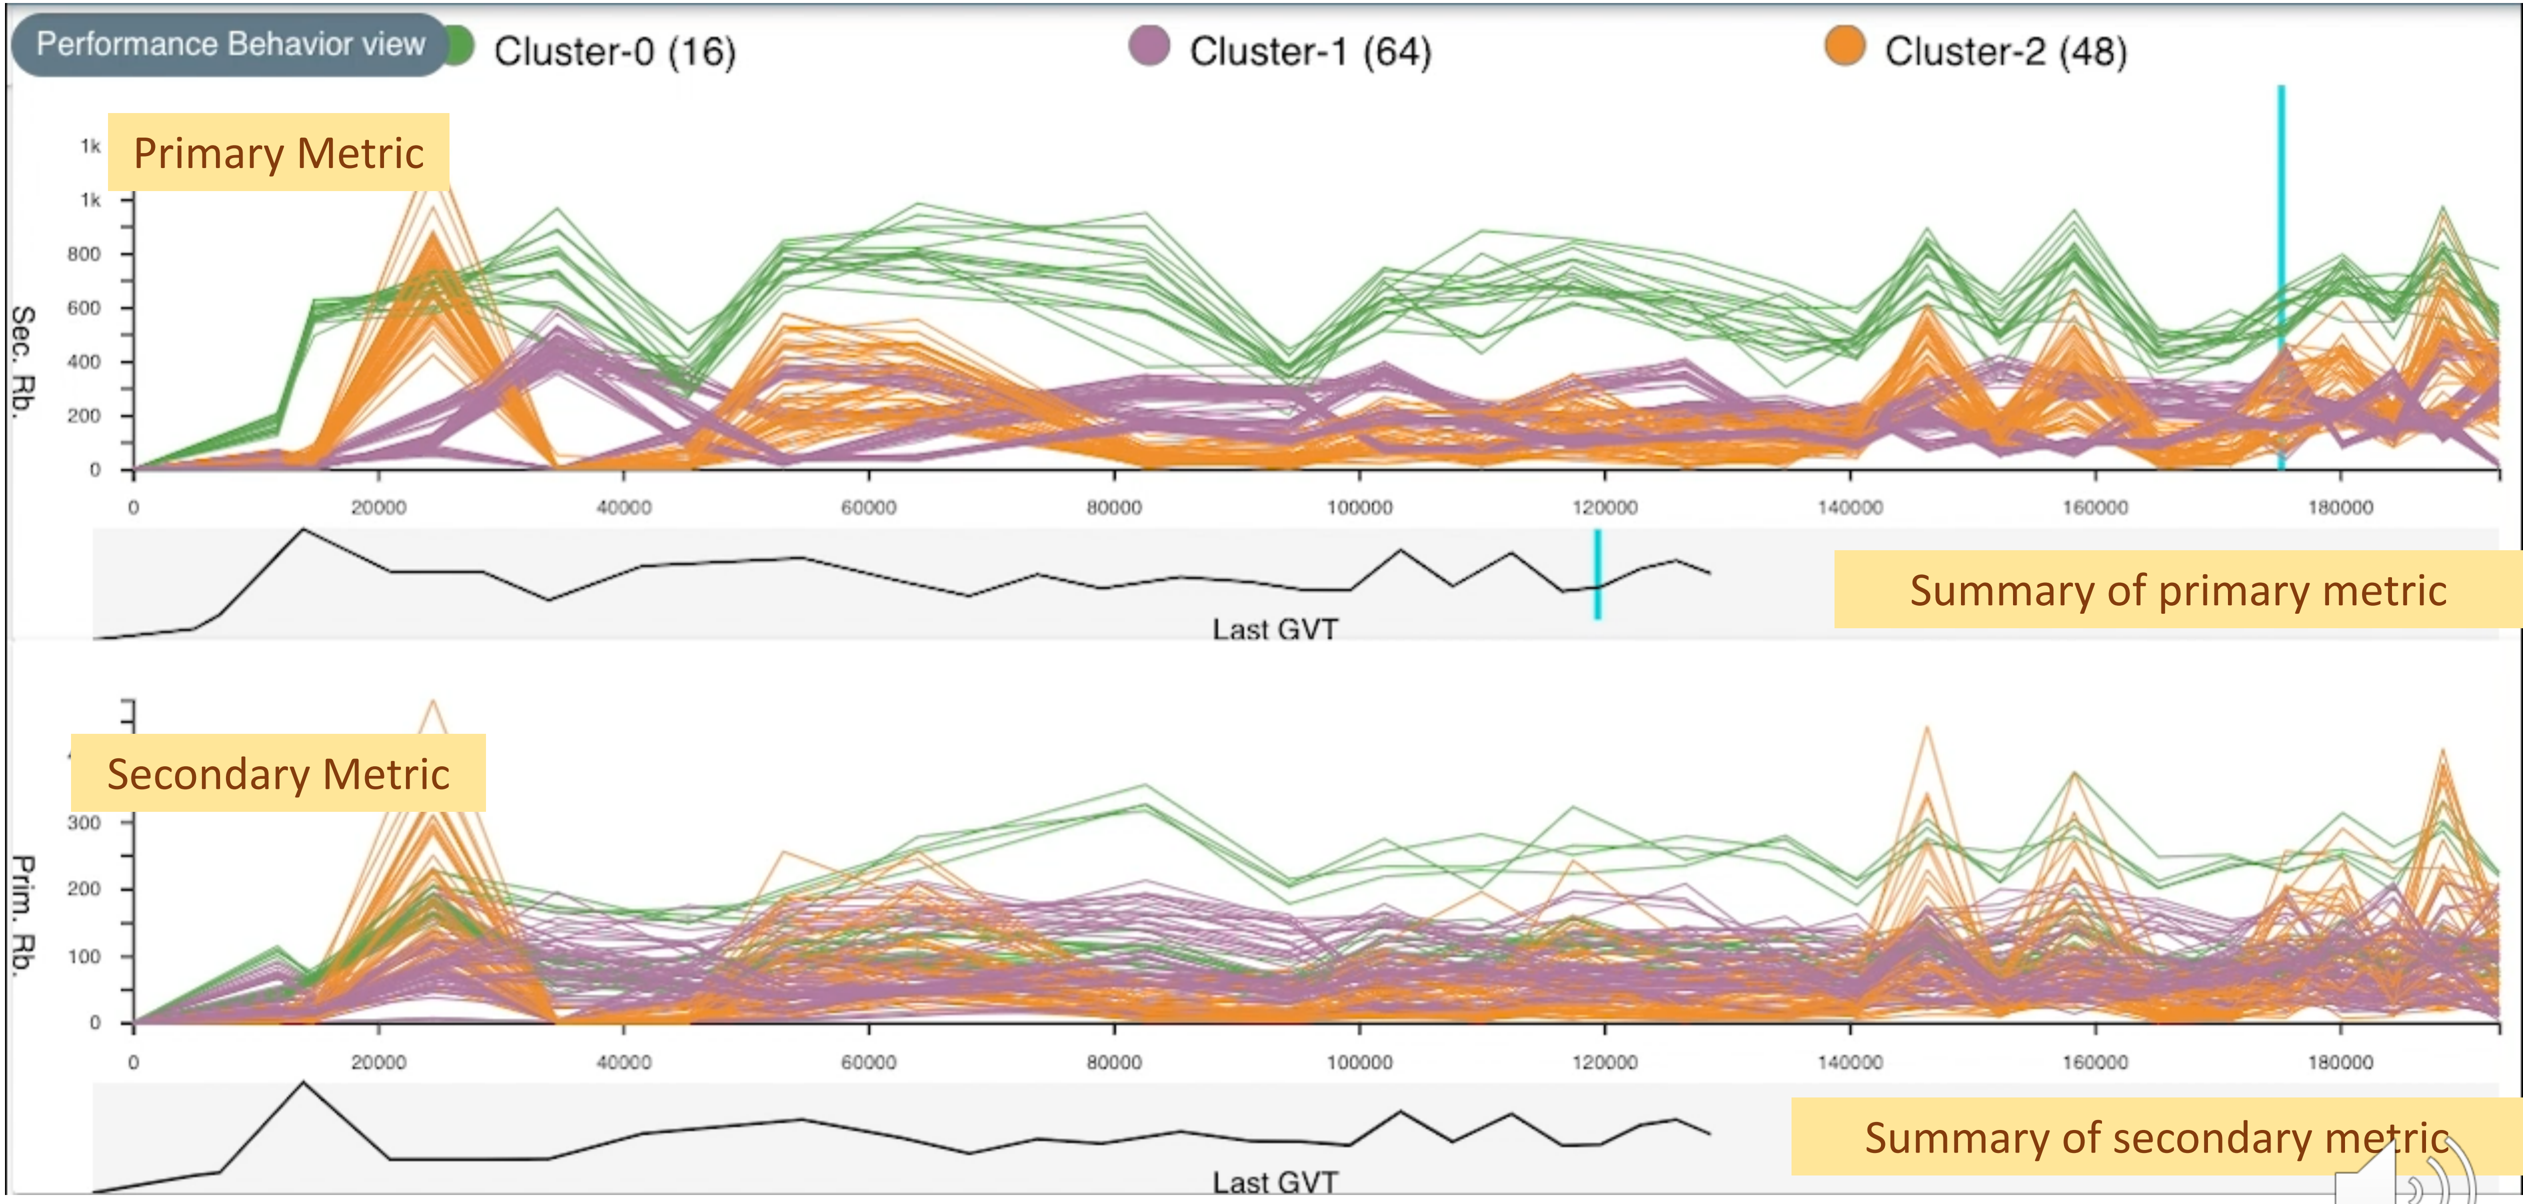

Behavior Similarity view

Progressive DR results are projected onto a 2D scatterplot.

Points are positioned based on the results from their PC1 and PC2.

Colored by the cluster IDs calculated by Progressive time-series clustering.

Visualizing primary and secondary metric side—by-side allows the HPC expert to compare two metrics at a time.

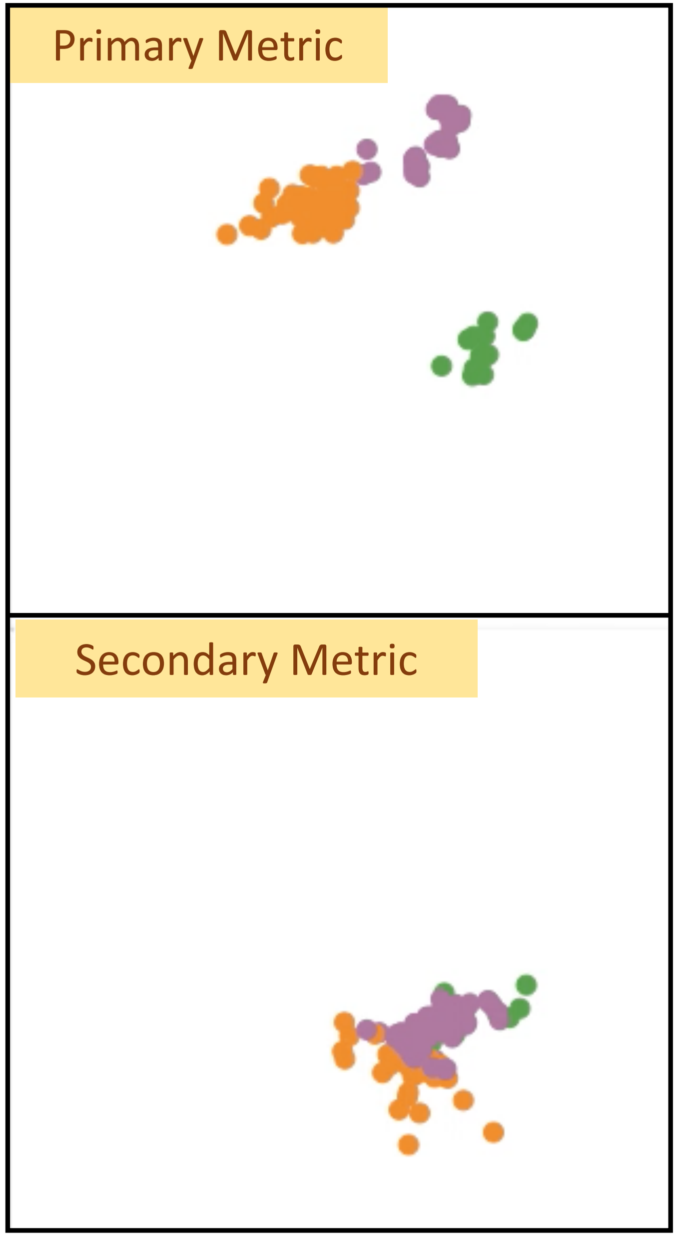

Communication views

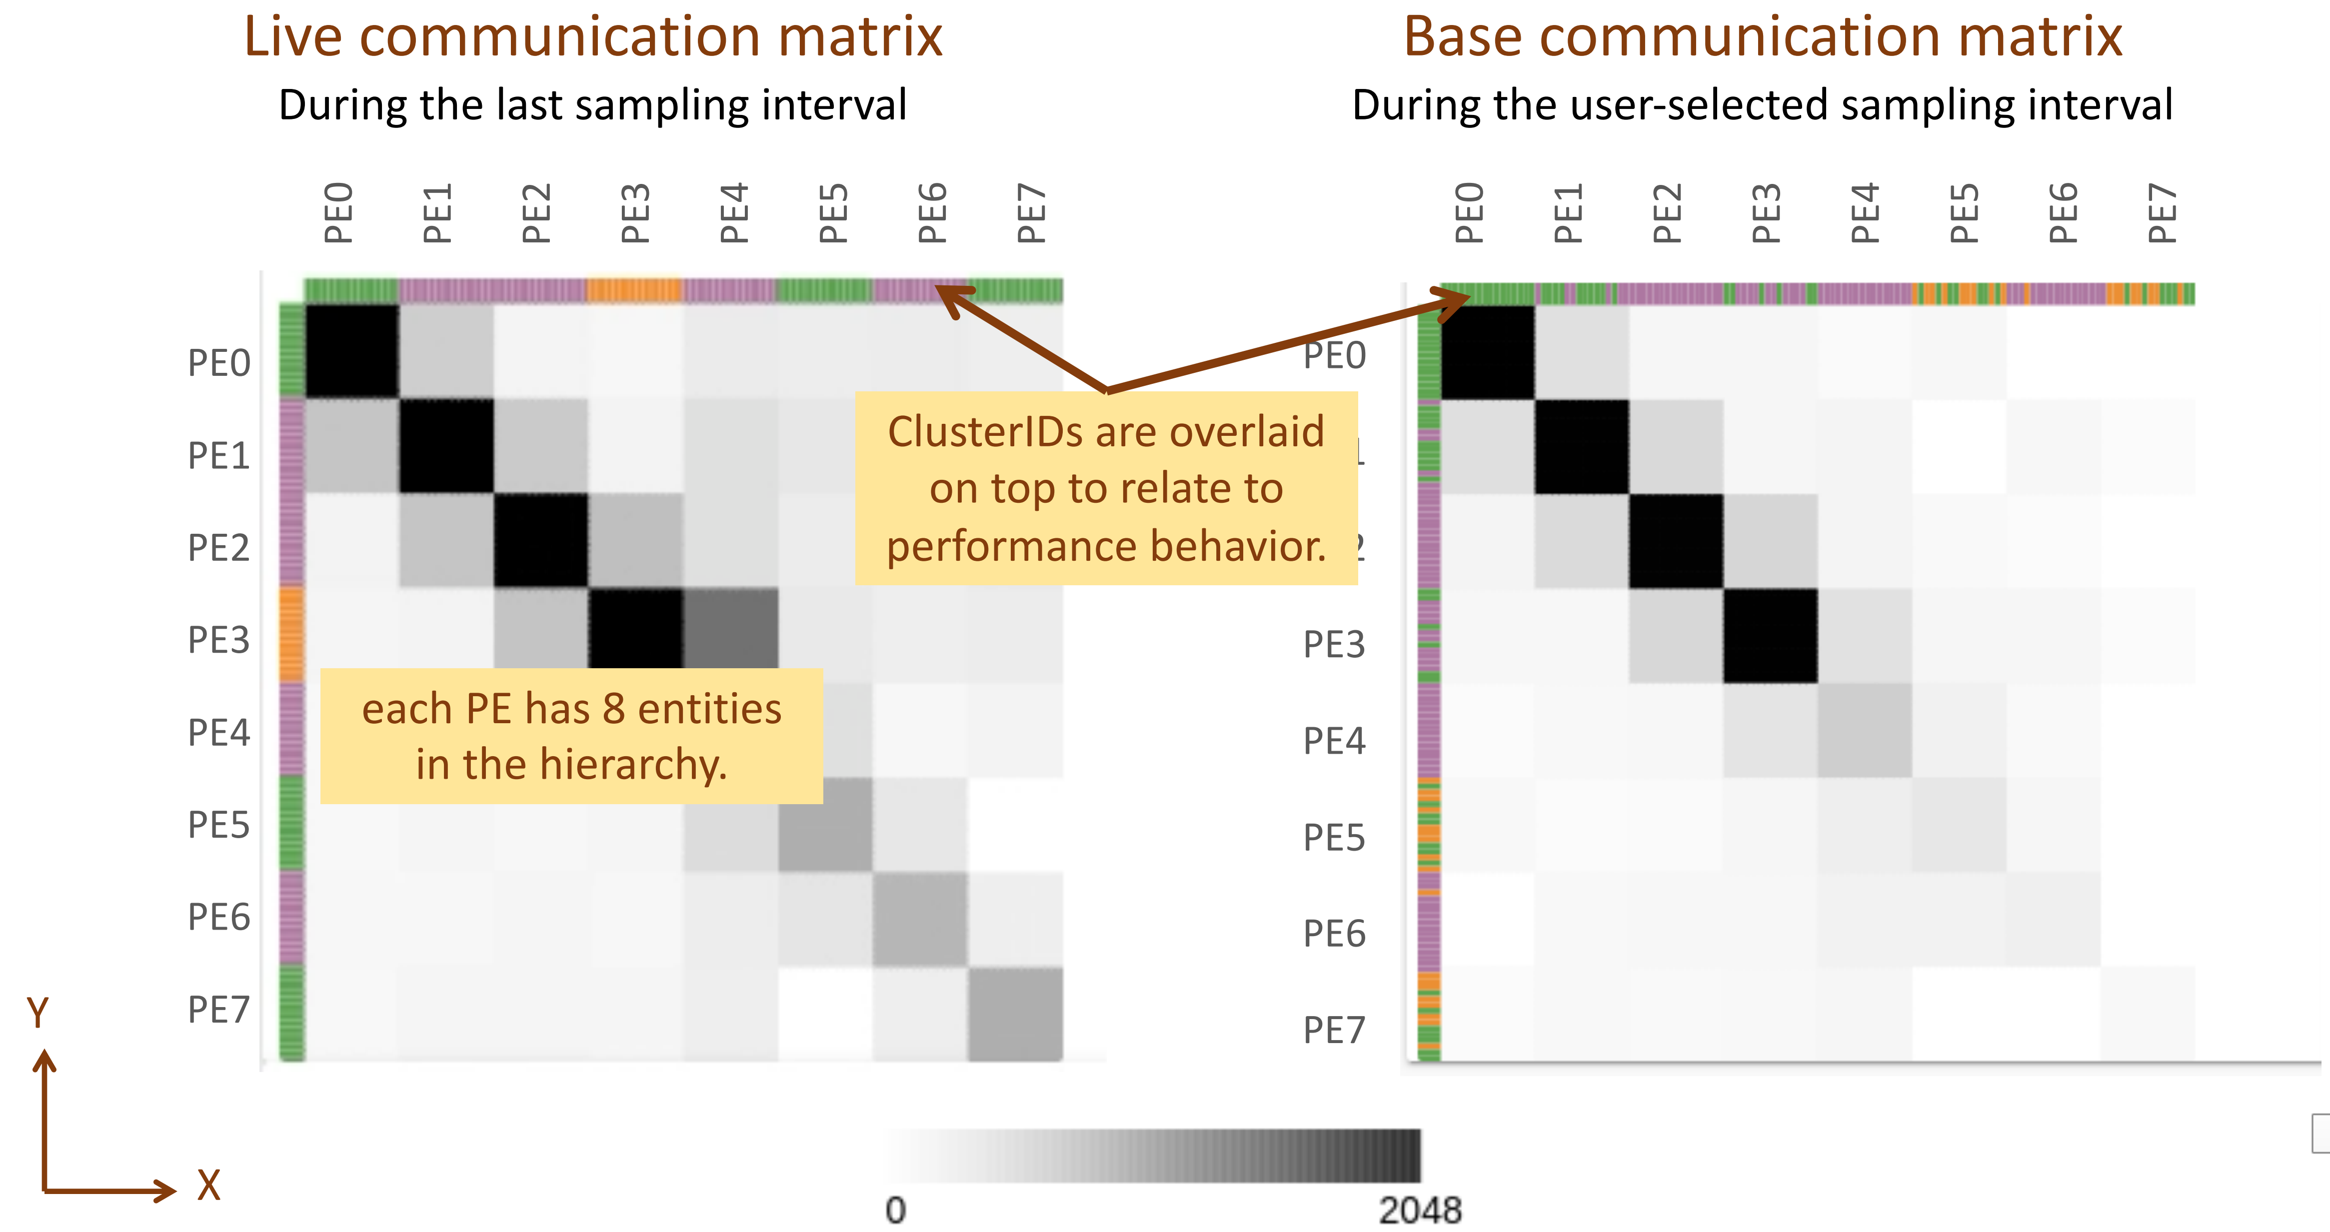

Communication views

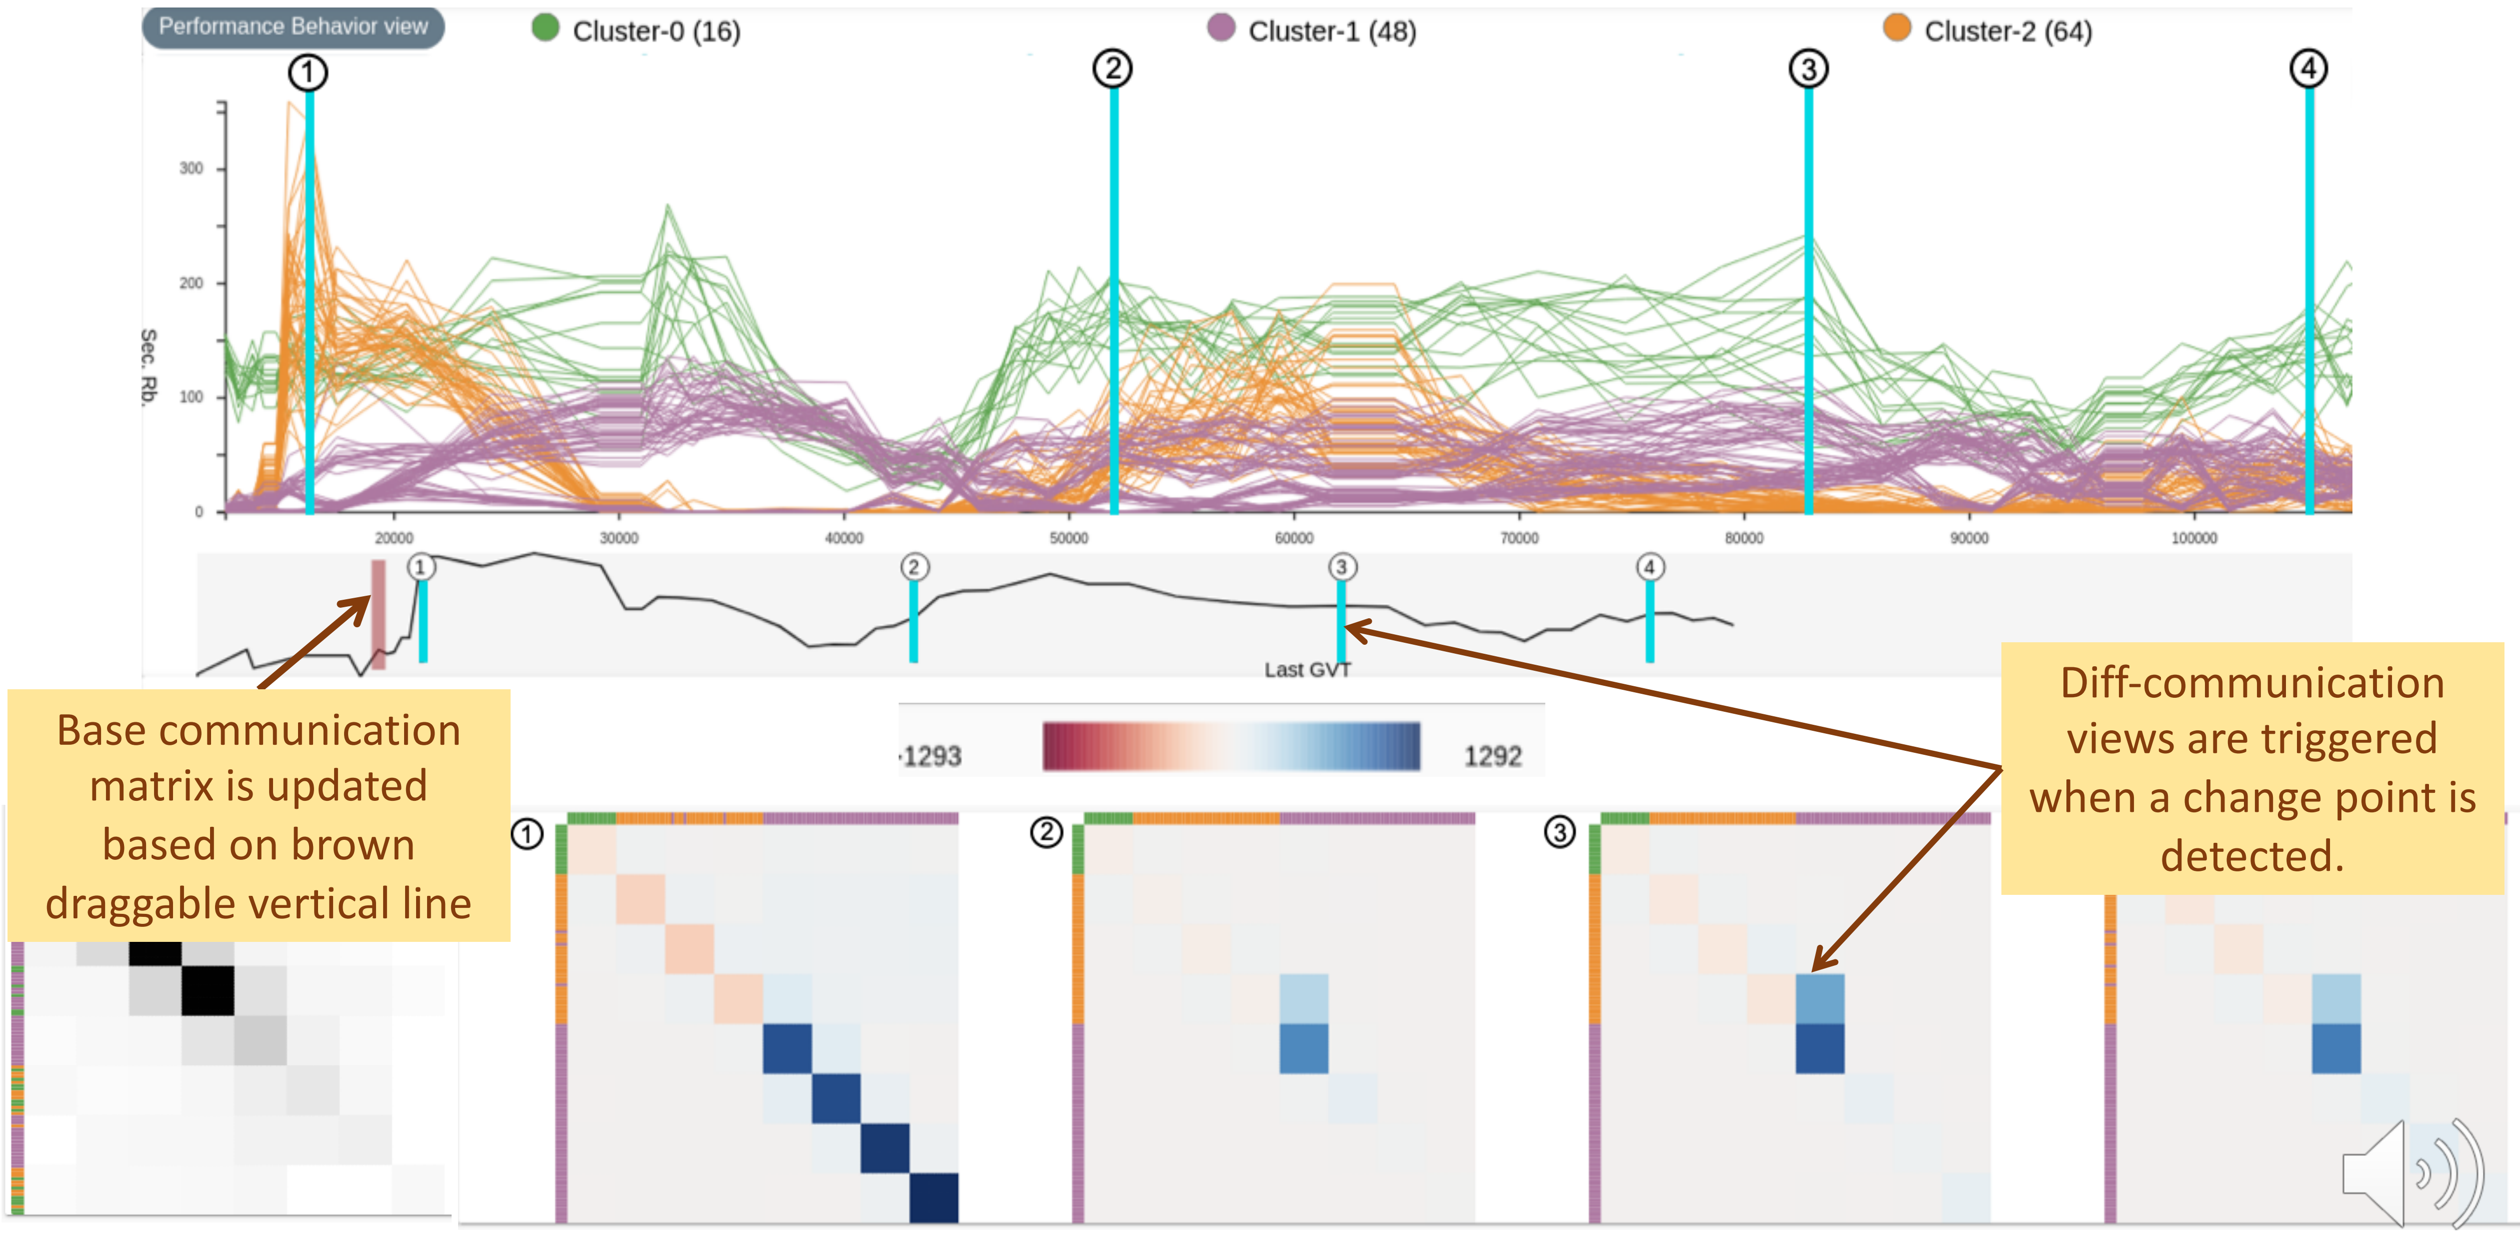

Causality view

From Causality

Effects from other metrics on the metric of interest

Causality view (From)

Metric	IR	VD
Net. Send.	7.90	0.09
Prim. Rb.	-46.15	0.04
Net. Recv.	20.04	0.02
Num. Events	-85.05	0.19
Sec. Rb.	116.22	0.64

p-value < 0.5

To Causality

Effect of metric of interest on other metrics

Causality view (To)

Metric	IR	VD
Net. Recv.	-237.75	0.00
Net. Send.	453.22	0.02
Num. Events	159.21	0.00
Sec. Rb.	86.05	0.72
Prim. Rb.	626.25	0.10

Case Study: Parallel Discrete-Event Simulation (PDES)

System

Theta at Argonne National Laboratory with the CODES network simulation toolkit [Cope et al., 2011] run with 864 routers.

Rensselaer's Optimistic Simulation System (ROSS)

Number of PEs: 8
Number of KPs: 128
Number of LPs (entities): 16384

Application

AMG solver application [Yang et al., 2002]

Metrics

Secondary rollback: The number of rollbacks on a KP caused by an cancellation message.
Network Sends (Net. Send.): The number of events sent by LPs over the network.
Last Global Virtual Time (Last GVT.): sampling interval in virtual time

Case Study: Monitoring Key Changes in PDES Performance

Metric	IR	VD
Net. Recv.	-3.49	0.00
Net. Send. 1	55.26	0.60
Num. Events	-14.82	0.03
Sec. Rb.	61.83	0.37
Prm. Rb.	-9.43	0.00

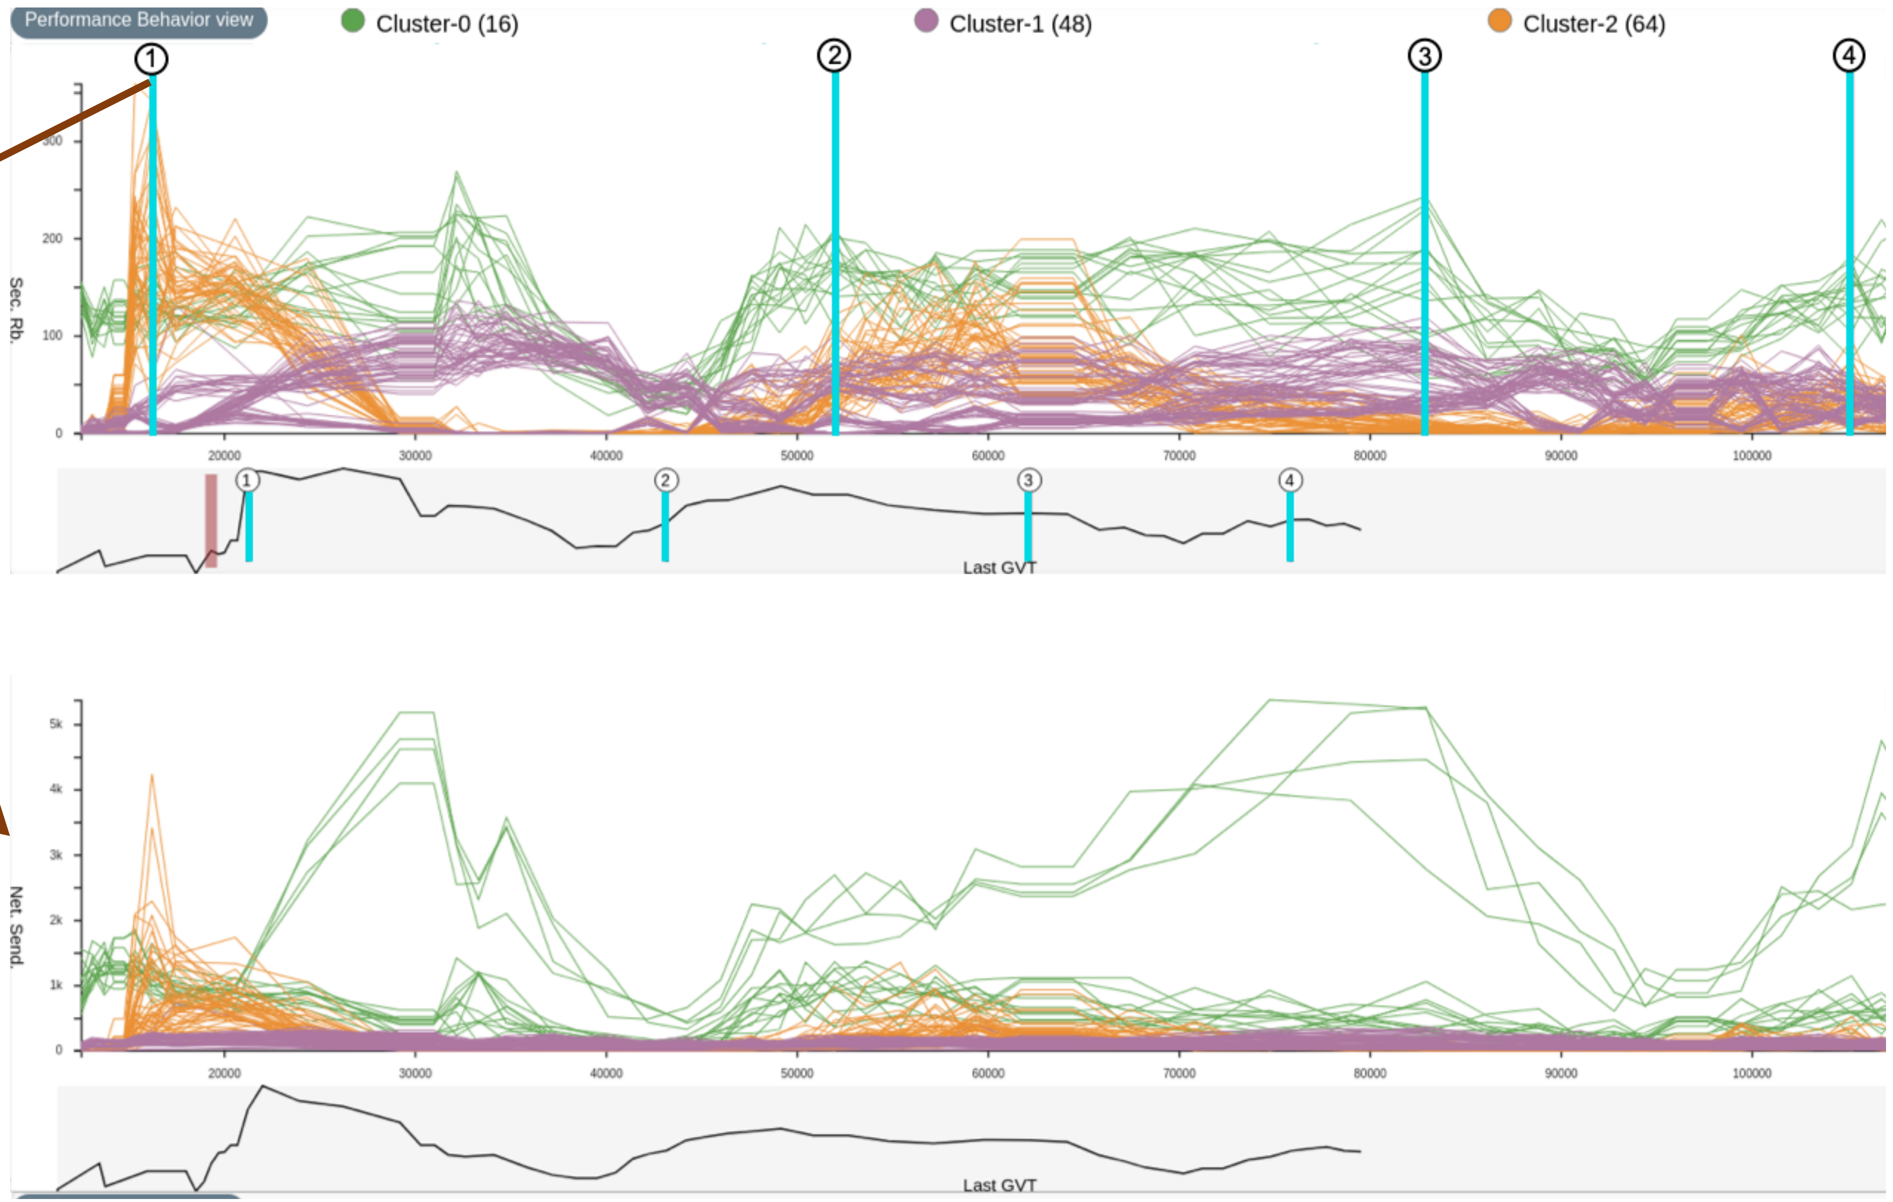

Case Study: Tracing Performance Bottlenecks

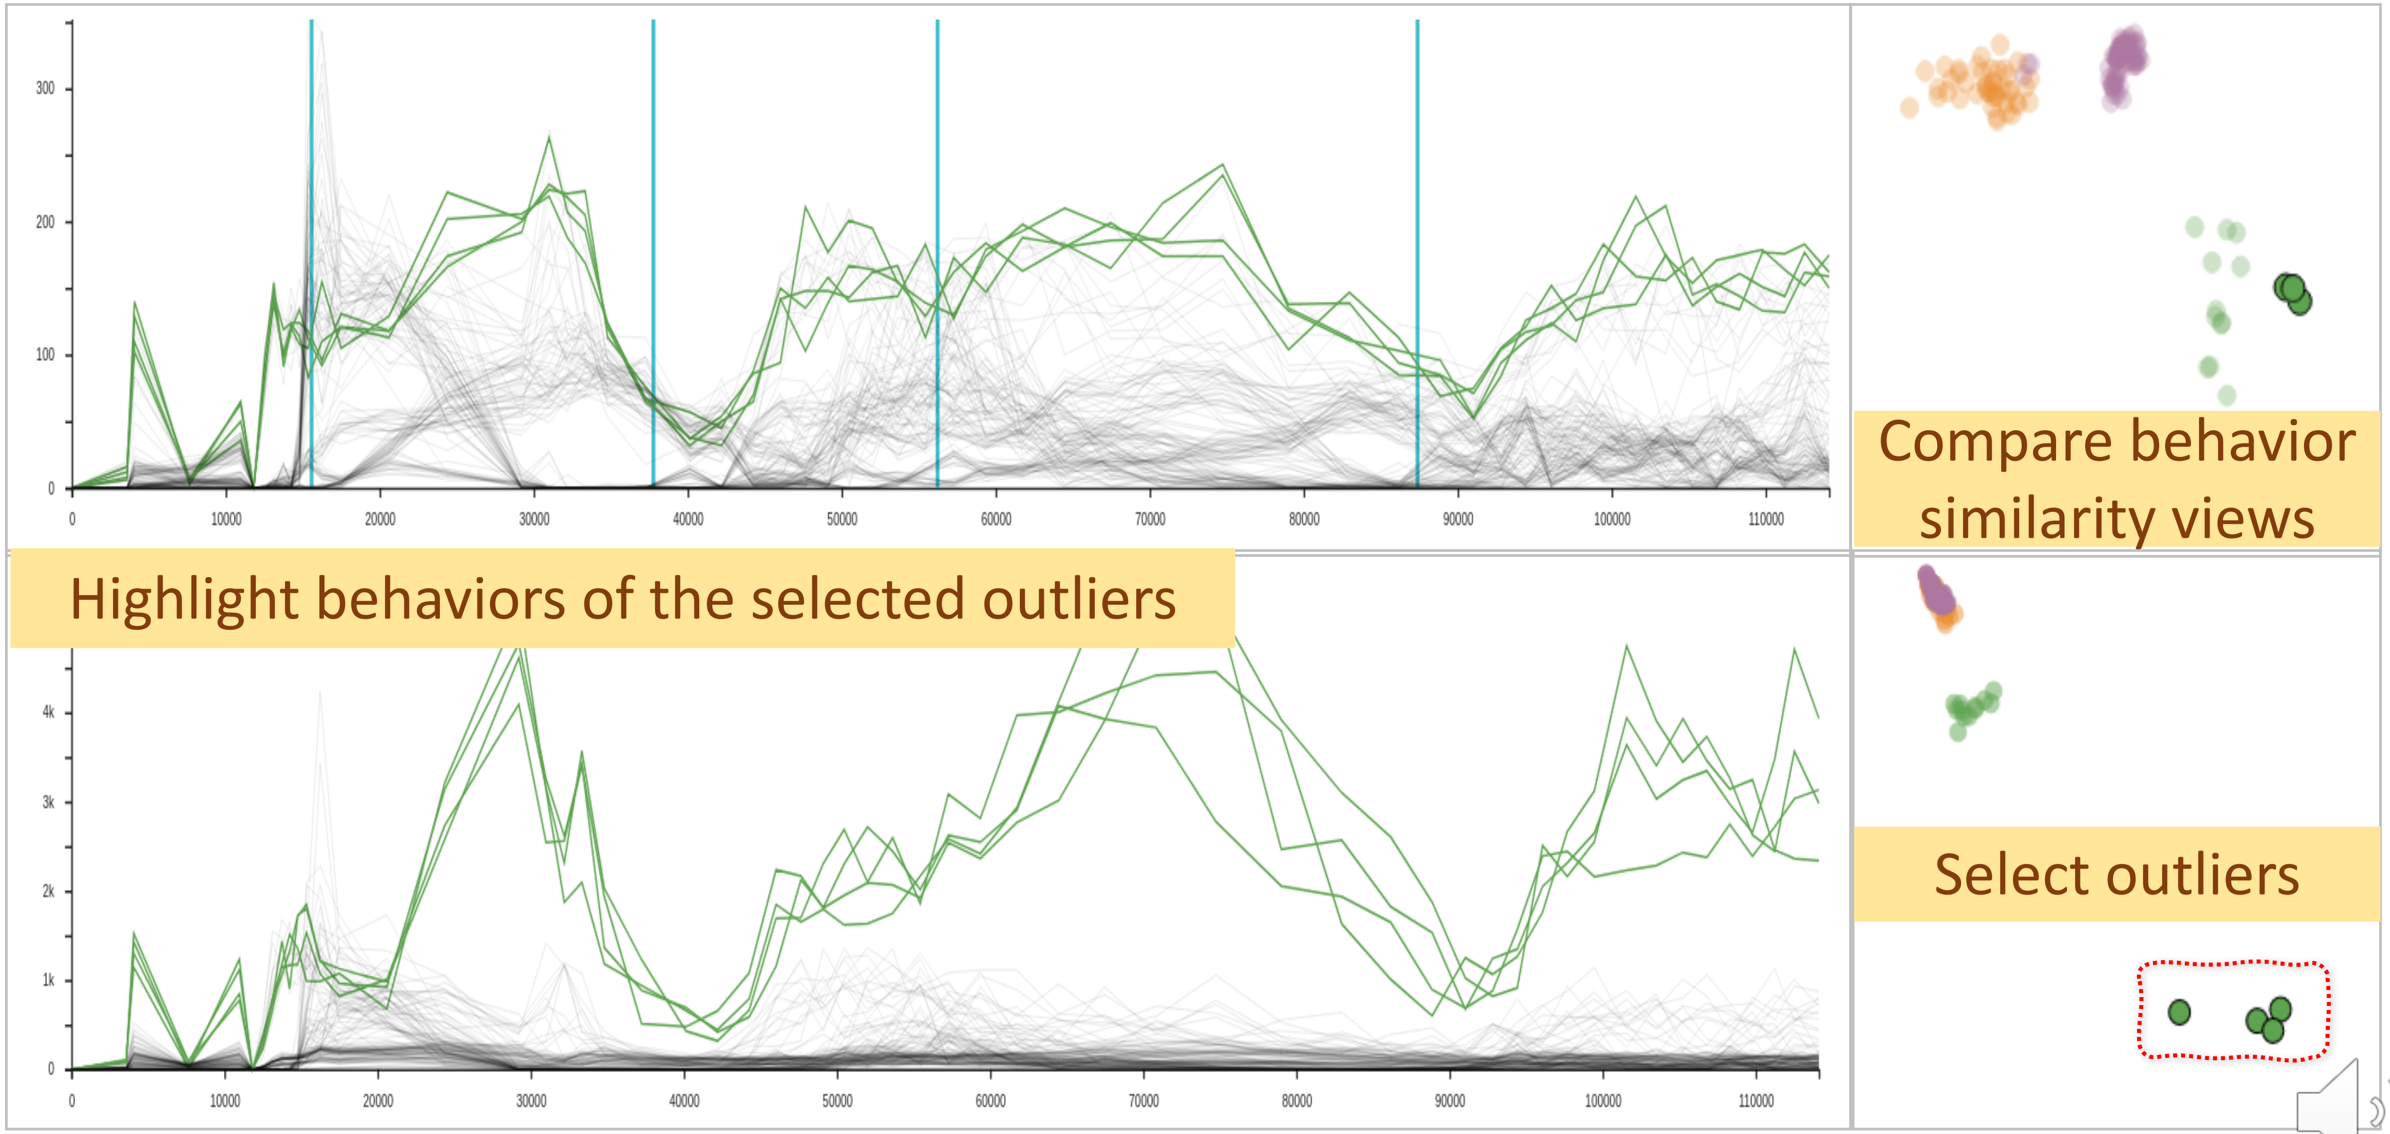

Case Study: Analyzing Communication Patterns

Visualize PE level communications

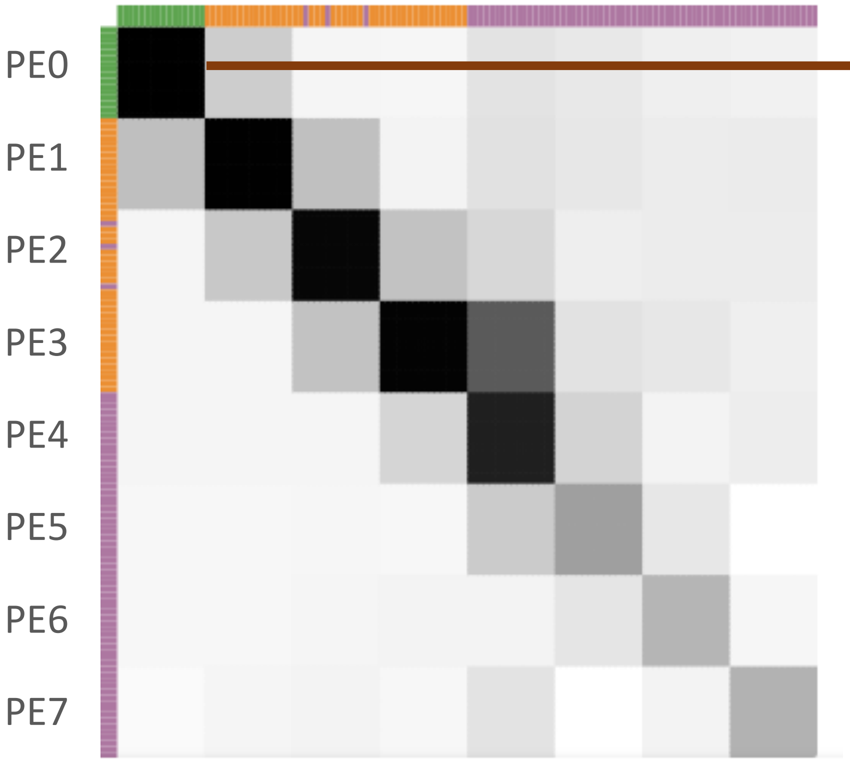

Show KPs belonging to green cluster

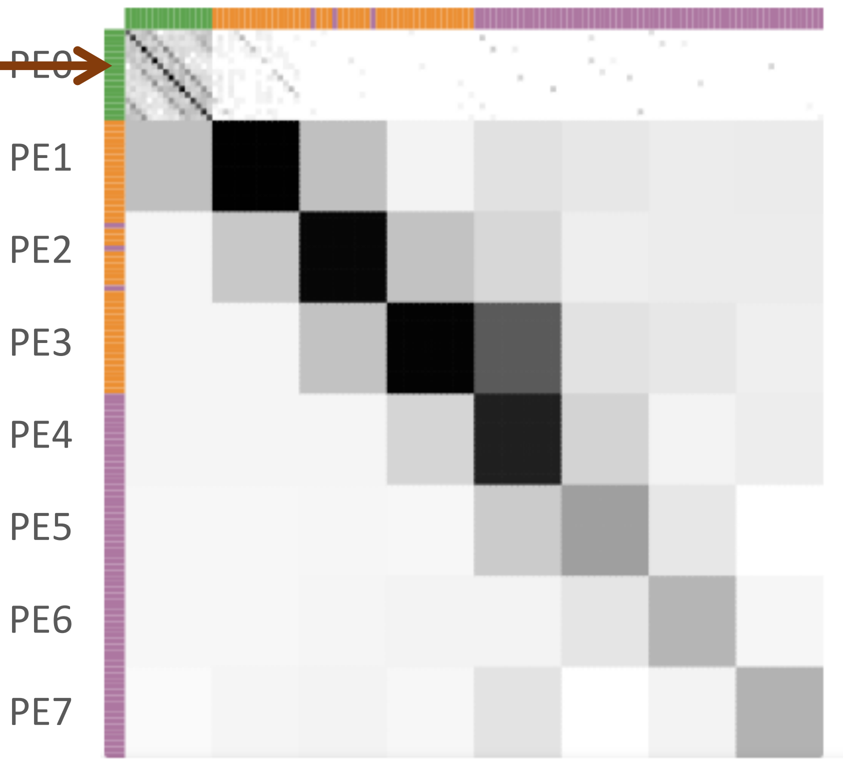

Limitations and Future Work

Latency

- For **extremely short sampling rate**, the progressive algorithms could not provide useful intermediate results because of **limited number of entities.**
- Visualizations can keep updating too frequently, whereby following the patterns becomes more challenging.

Controlling the frequency of updates in the data management module.

Scalability

- Current implementation supports **limited number of metrics** (< 20 metrics) and **number of entities** (< 10,000 entities) that can be processed.

Aggregating multiple metrics based on similarity

Comparison

- Our framework can only allow **comparison of 2 selected metrics (side-by-side views)** and **communication matrices (diff-communication views)**.

Tracking the performance metrics through animated overviews.

Questions?

Mail: spkesavan@ucdavis.edu
Website: CODES-VIS http://vis.cs.ucdavis.edu/codes-vis/

This research has been supported by U.S. Department of Energy through grant DE-SC0014917.