Correlation-Aware Flow Consolidation for Load Balancing and Beyond

Shiva Ketabi[*], Matthew Buckley[*], Parsa Pazhooheshy[*], Faraz Farahvash[+], Yashar Ganjali[*]

[*]University of Toronto, [+]Sharif University of Technology

November 12, 2021

UNIVERSITY OF TORONTO

SHARIF UNIVERSITY OF TECHNOLOGY

Motivation: Load Balancing Challenge

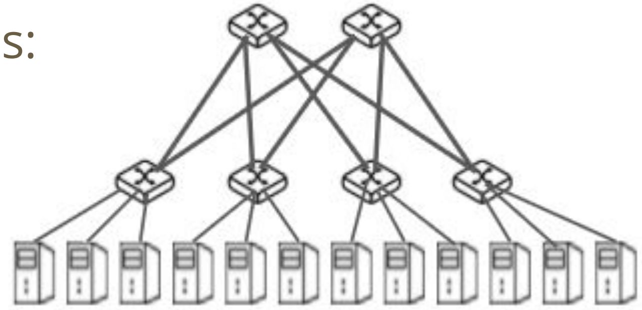

- Modern topologies offer several equal-cost paths:
  - **Load balancing is a critical component**
- Existing schemes:
  - Random and congestion-oblivious:
    - Packet-based: ECMP. Flowlet: Presto, Let it flow, ...
  - Congestion-aware:
    - Local: Drill, Global: Conga, Hula, Clove, ...
- Challenge:
  - **Volatile nature of flows:** Fluctuations in rates lead to under/overshooting of link capacities.
  - **Existing schemes are random or rely on average measurements.**
  - **None of them consider second degree variations in flow rates**!

Motivation: Correlation-Aware Flow Consolidation

- Goal: **Minimizing collective rate variations.**
- **Aggregating inversely correlated or independent flows into "*superflows*".**
- Feeding load balancing schemes with "superflows".
- Enhancing existing load balancing schemes:
  - **Reducing peak requirements.**
  - **Estimating future group demands with higher confidence.**
  - **Eliminating congestion in time-slots smaller than control intervals.**

Example: Flow Consolidation

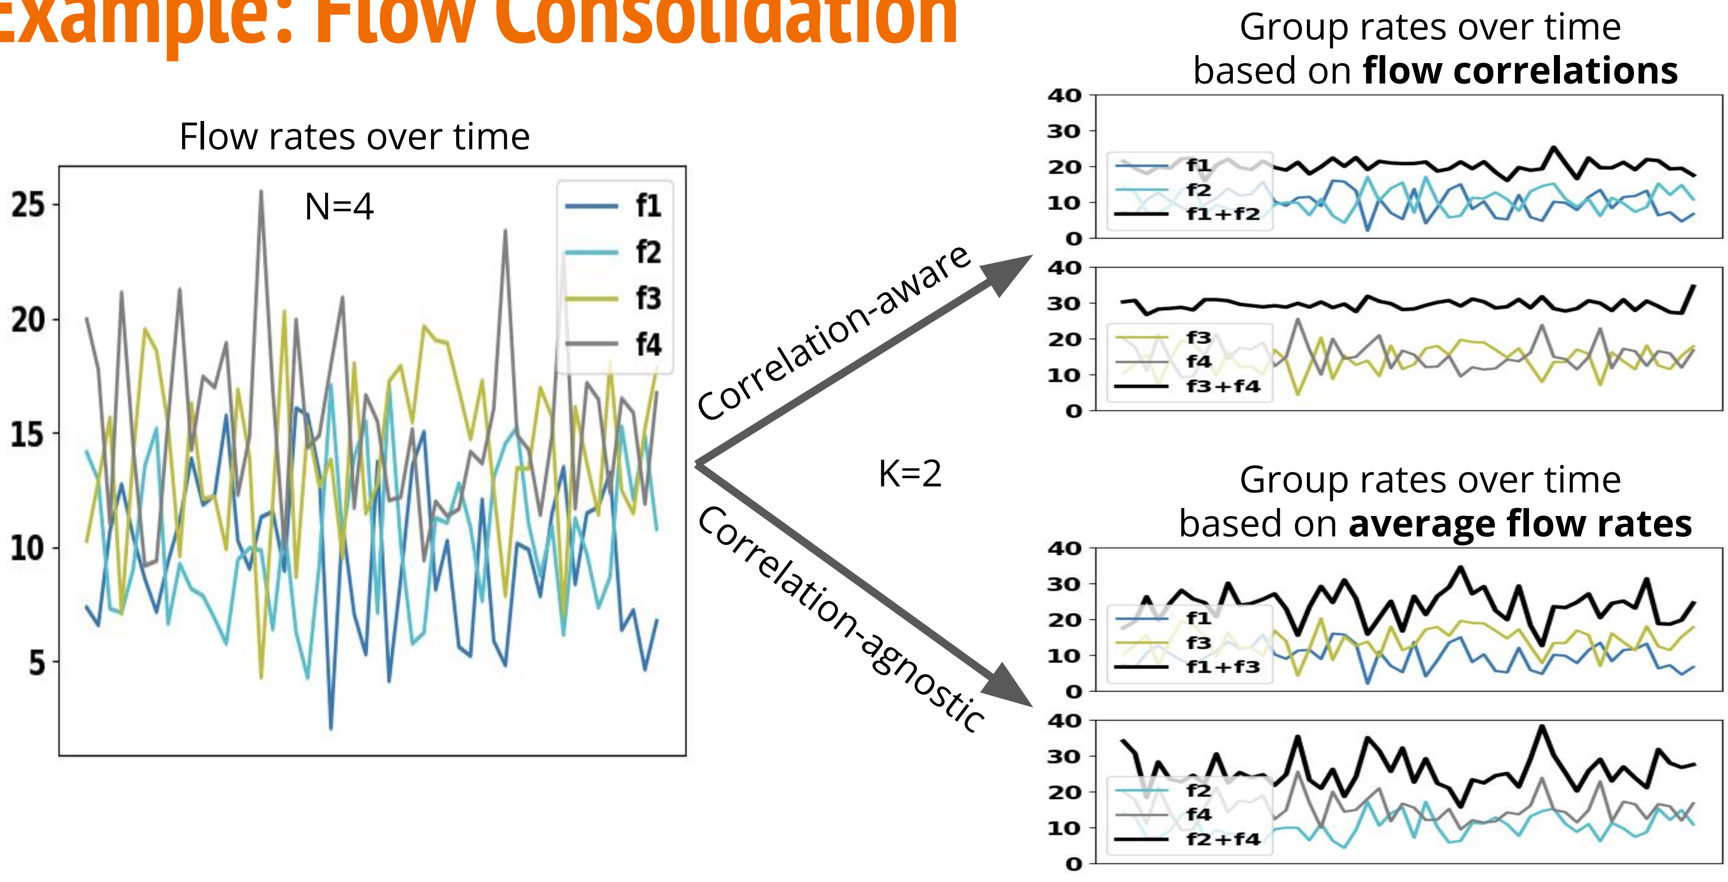

Formulation: Flow Consolidation Problem

N flows $\mathcal{F} = \{f_1, \ldots, f_N\}$

$\mathcal{G} = \{g_1, \ldots, g_K\}$

$\text{Group} : \mathcal{F} \to \{1, \ldots, K\}$

$\text{Group}(f_i) = k$ implies that $f_i \in g_k$

We formulate an optimization problem that minimizes the maximum of the aggregate rates among all the possible grouping functions.

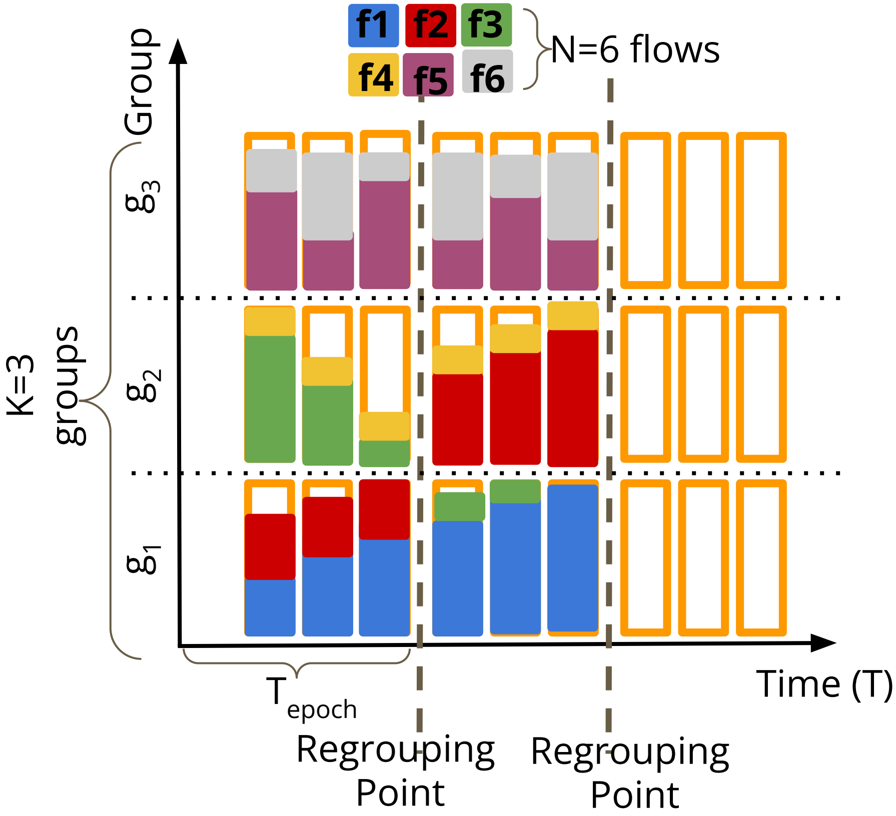

Formulation: Correlation-Aware Flow Consolidation

- The well-known *multi-processor scheduling* problem reduces to it.
- Computational complexity: **NP-hard**.
  - Solution: we use ***variance*** of aggregated rate instead of maximum.
  - To control the variance of group rates, we focus on ***flow correlations***.
  - Intuition: aggregation of independent (inversely-correlated) flows shows less variance over time.
- Practical challenge: **individual flow rates are not available a priori**.
  - Solution: to estimate flow correlations, we use the flow rates of the previous epoch to predict the future flow rates.

Solution: System

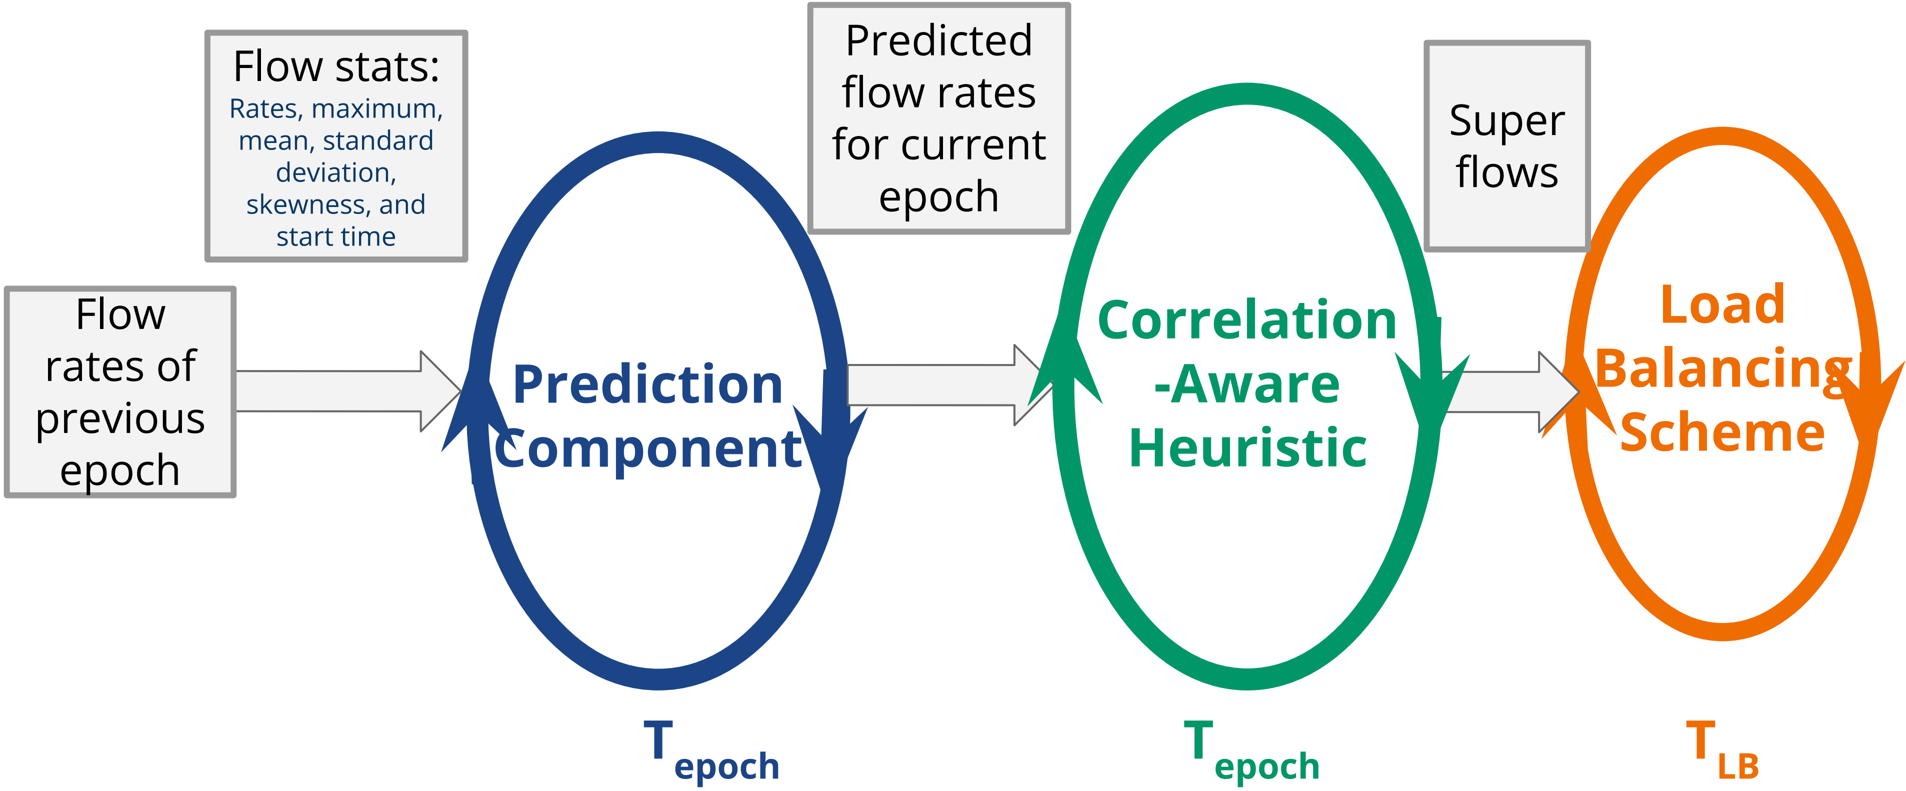

Solution: Correlation-Aware heuristic

- **Lowest Correlation Grouping (LCG):**

1. Initializing K empty groups g_1, g_2, ..., g_K.
2. Maintaining a total aggregate flow, S_k.
3. For each flow f:

4. Compute the correlation between f and S_k for all non-empty groups.

5. Find the minimum correlation, c*, and the corresponding group, g*.

6. Comparing c* to a threshold, if smaller assign f to g*.

7. Otherwise, assign to the next empty group.

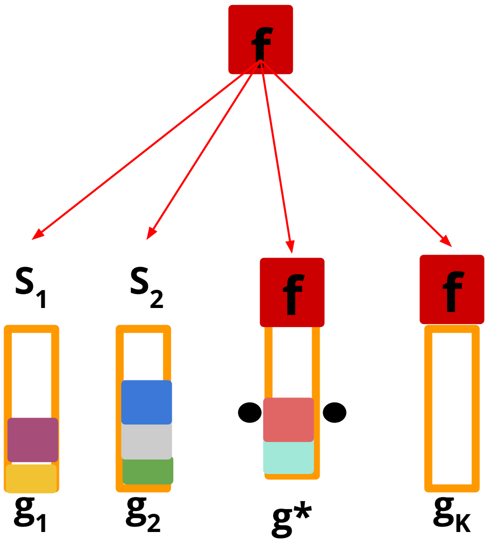

Solution: Prediction Component

- Tested Models:
  - Random Forest Regression
  - XGBoost
  - Ridge Regression
  - Stochastic Gradient Descent Regression
- Learning phases:
  - 30% epochs for training
  - 20% epochs for validation
  - 50% epochs for testing
- Prediction metric:
  - R-squared

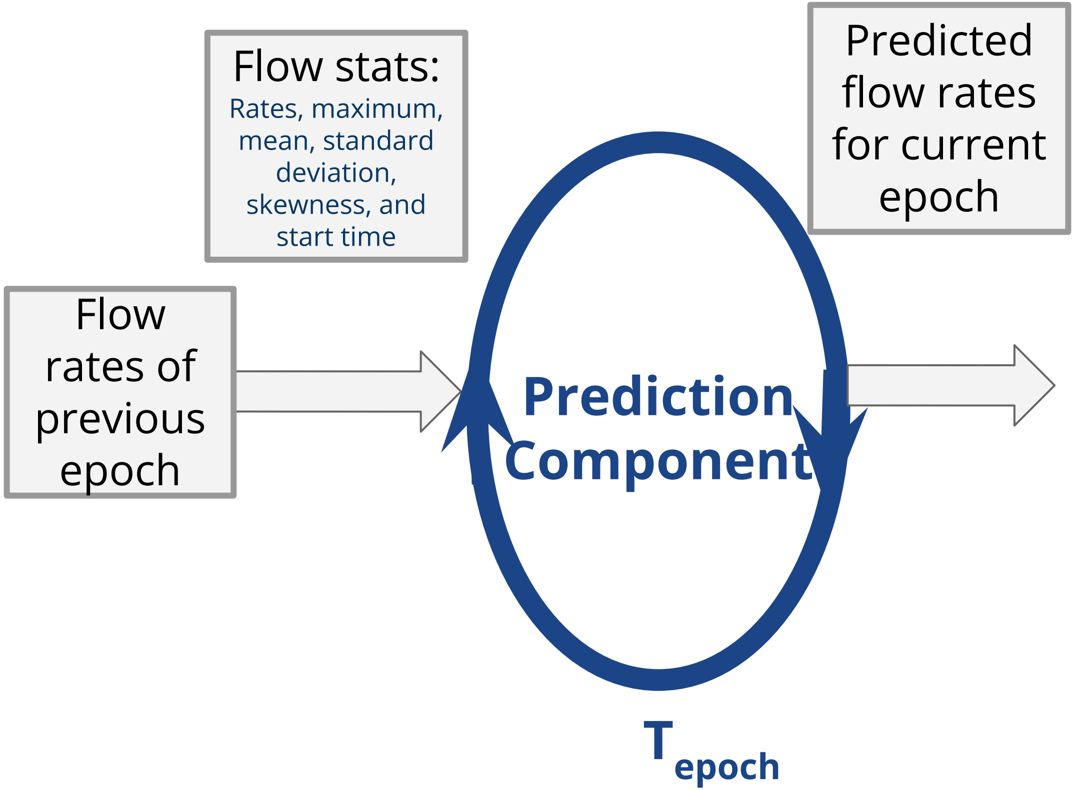

Evaluation: Setup

- Traffic:
    - **A real ISP data** (W. N. R. Group. WAND, ISPDSL II dataset. https://wand.net.nz/wits/ispdsl/2/.)
    - **500 flows for half an hour**
    - Flows are randomly passed through 30 paths
    - T_{epoch} is 5 seconds
    - Long-lived flows: active in at least in 1 time slot of previous epoch
    - Short-lived flows: only active in current epoch, randomly grouped
- Comparables:
    - **LCG**: Lowest Correlation Grouping
    - **HRF**: Highest Rate First
    - **Random**
- Comparables:
    - **Mean and max group standard deviations**

Evaluation: Predictive Models Importance

- LCG using current epoch outperforms random and HRF.
- LCG using previous epoch does not perform well at the tail.
- LCG with Ridge Regression reduces standard deviation by 33% at the 50th percentile, and 12.5% at the 99th percentile.

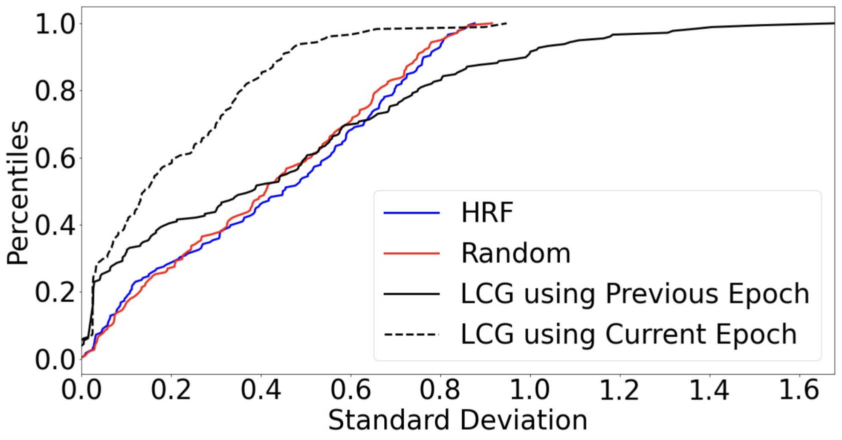

Without using predictive models

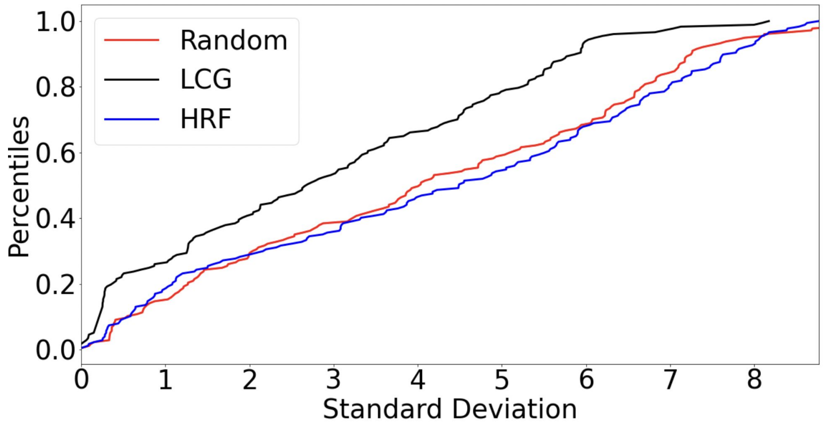

With Ridge Regression

Evaluation: Number of Groups

- LCG outperforms random more apparently as the number of groups increases.
- LCG with ridge regression outperforms oracle-based HRF.

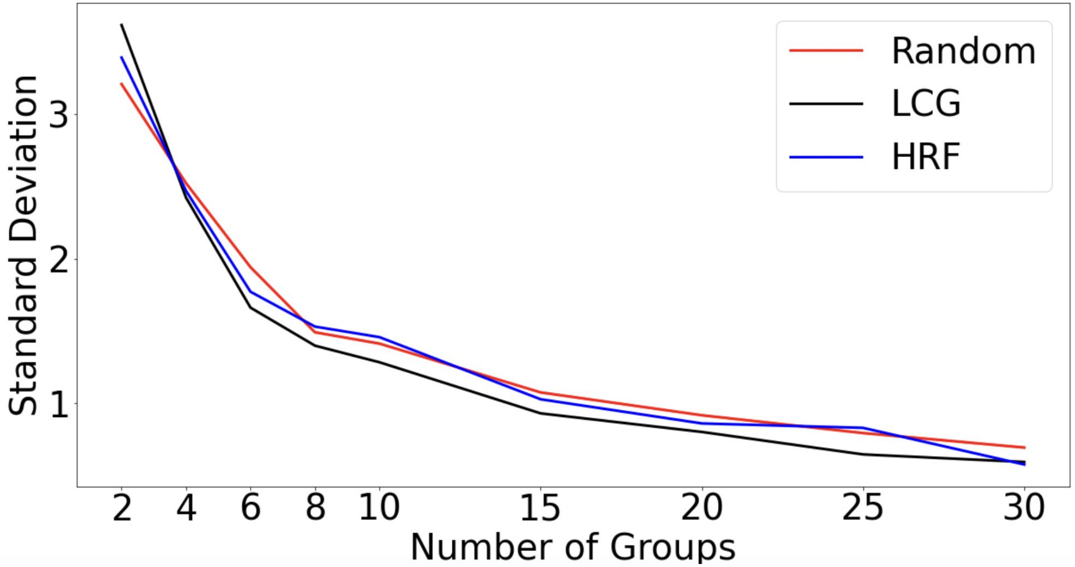

Conclusion and Future Work

- The first solution for load balancing that:
  - Considers **flow rate variations** and **correlations**.
  - Consolidates inversely correlated or independent flows.
  - Reduces rate fluctuations in superflows.
- Our design includes:
  - A correlation-aware heuristic for the NP-hard problem.
  - A prediction component to resolve the issue of unknown future rates.
- Our experiments show:
  - Significant reduction in standard deviation of group rates.
  - The importance of the prediction component.
- Future work: evaluation of load schemes+our solution, other applications.

Thank you!