The Power of Both Choices: Practical Load Balancing for Distributed Stream Processing Engines

Muhammad Anis Uddin Nasir[1], Gianmarco de Francisci Morales[2], David Garcia Soriano[2], Nicolas Kourtellis[2], Marco Serafini[3]

[1] KTH Royal Institute of Technology, [2] Yahoo Labs Barcelona, [3] Qatar Computing Research Institute

1. Abstract

- **Load Balancing** in Distributed Stream Processing Engines directly affects the hardware utilization, latency, and throughput of the system
- We introduce **Partial Key Grouping** that leverages the classical "**power of two choices**"
- It achieves **better load balancing** than Key Grouping while being **more scalable** than Shuffle Grouping

2. Problem

- Streaming applications are represented by directed acyclic graphs (DAG)
- Key Grouping and Shuffle Grouping are two common partitioning strategies
- Highly skewed distribution of workloads creates imbalance among workers for Key Grouping
- Shuffle Grouping incurs additional memory and aggregation cost
- A streaming and distributed solution is required for the problem

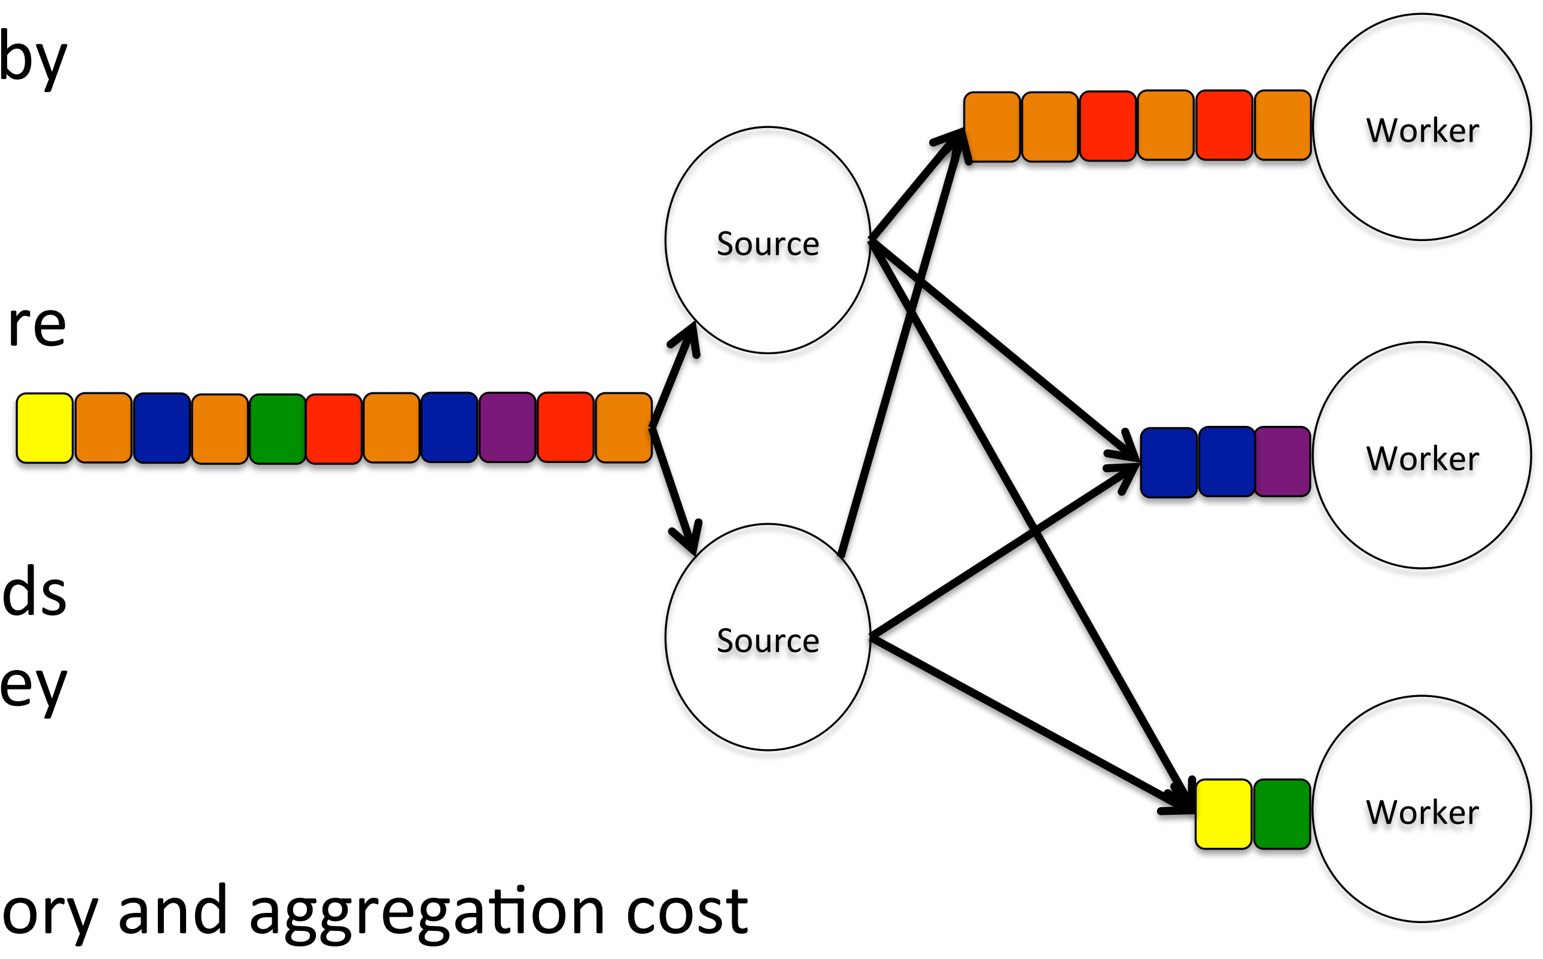

3. Partial Key Grouping

- PKG makes "power of two choices" applicable to our problem by using two novel techniques: ***Key Splitting*** and ***Local Load Estimation***
- ***Key Splitting*** relaxes the atomicity constraint of key grouping and assigns each key to exactly two servers
- ***Local Load Estimation*** solves the problem of gauging the load of downstream servers without any communication overhead

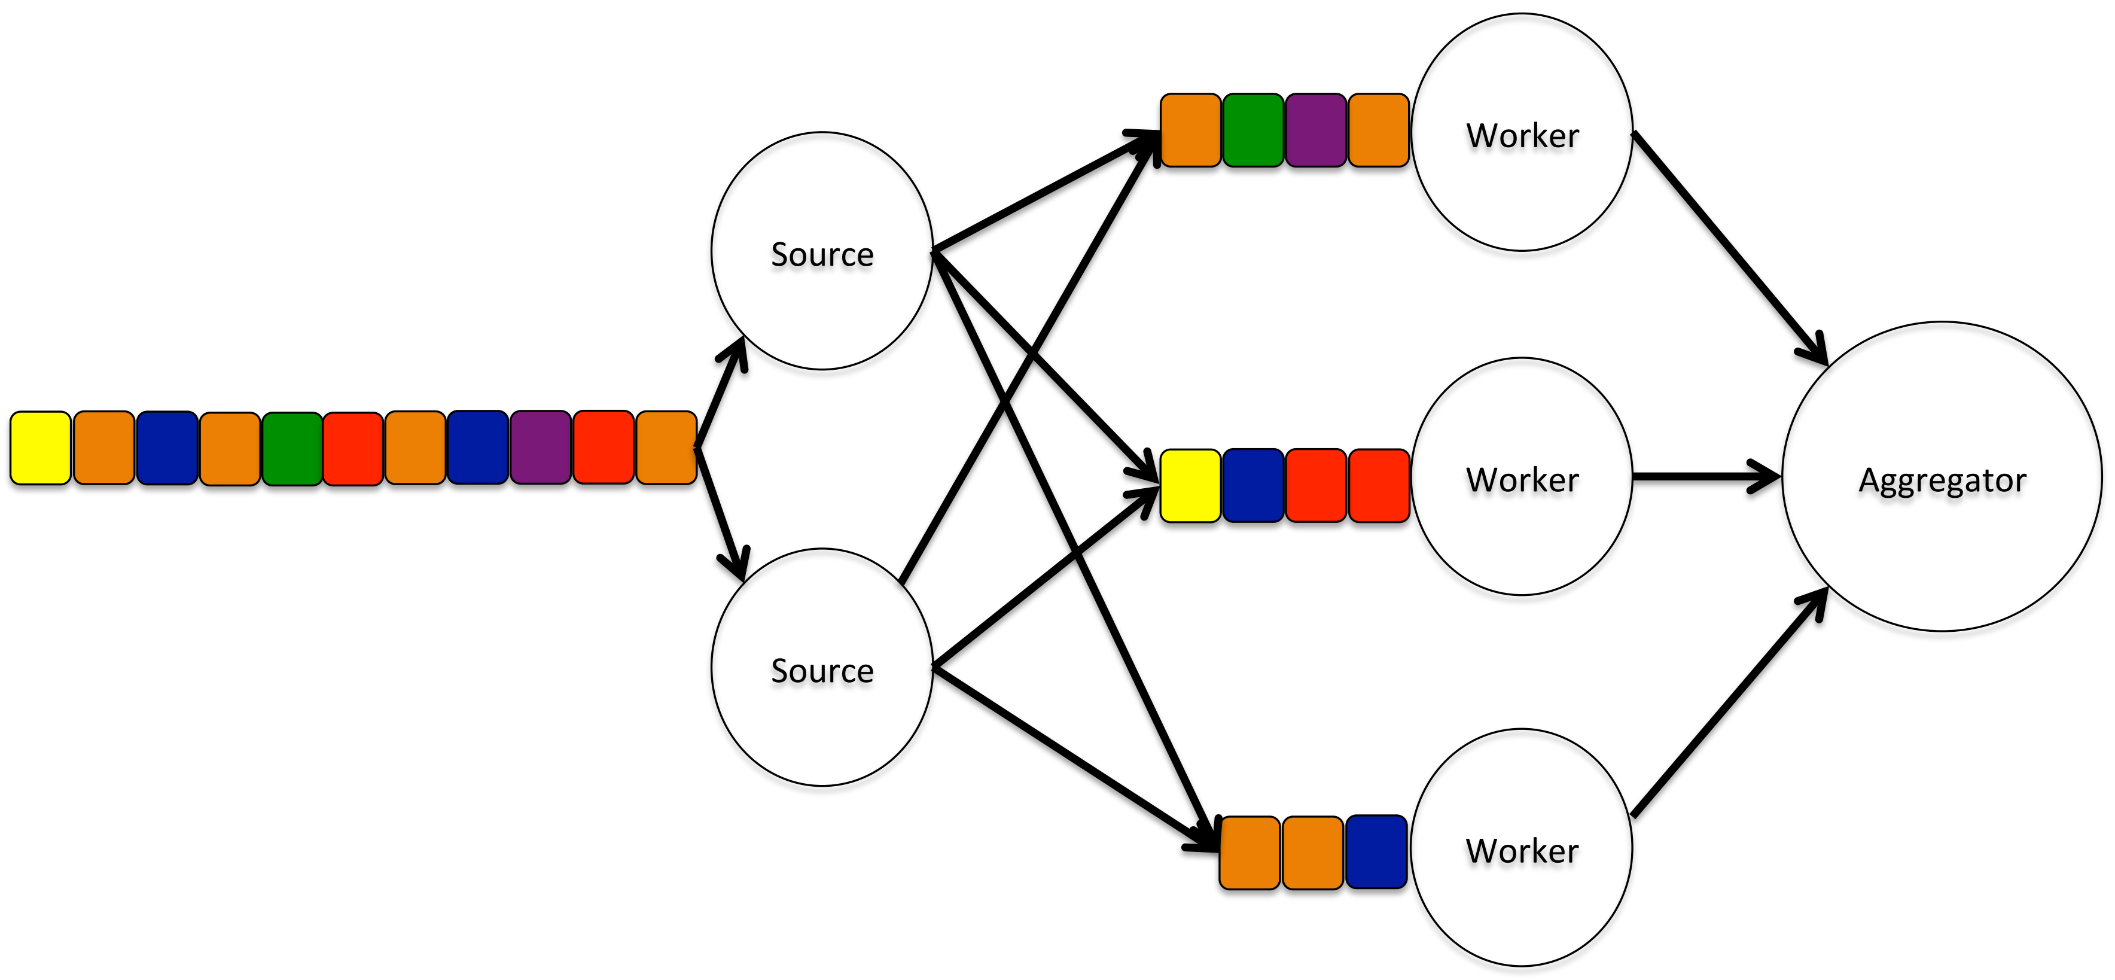

4. Theoretical Analysis

- Suppose we use m messages and n bins and let $m \geq n^2$. Assume a key distribution *D* with maximum probability $p_1 \leq 1/5n$. Then the imbalance after m steps of Greedy-d process satisfies, with probability at least $1 - 1/n$

$$I(m) = \begin{cases} O\left(\frac{m}{n}\cdot\frac{\ln n}{\ln\ln n}\right), & \text{if } d = 1 \\ O\left(\frac{m}{n}\right), & \text{if } d \geq 2 \end{cases}$$

5. Experimental Setup

Questions

- Effect of Key Splitting
- Local vs Global Estimation
- Robustness
- Real Implementation

Datasets

- Twitter: 1.2G tweets
- Twitter: 690k cashtags
- Wikipedia: 22M page visits
- Live Journal: 69M edges

6. Real Implementation

Storm Implementation

- 60% improvement in throughput
- 45% improvement in latency

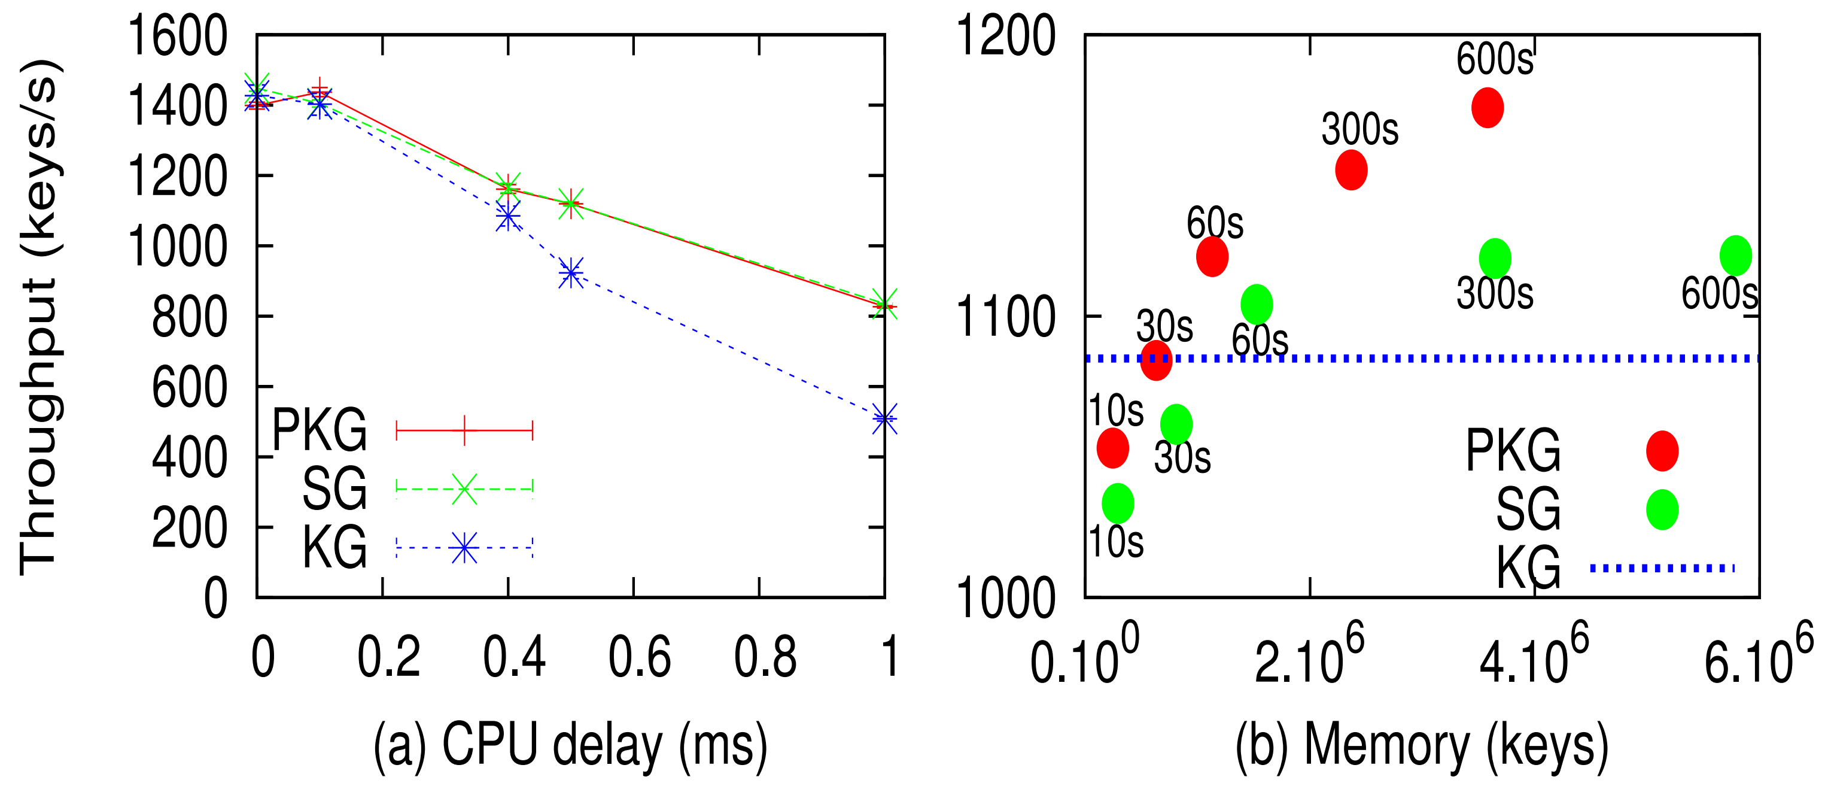

7. Simulations

- PKG provides better load balance compared to other strategies, e.g., greedy and power of two choices
- PKG with local load estimation works as well as with global oracle
- PKG reduces the load imbalance by up to seven orders of magnitude compared to Key Grouping

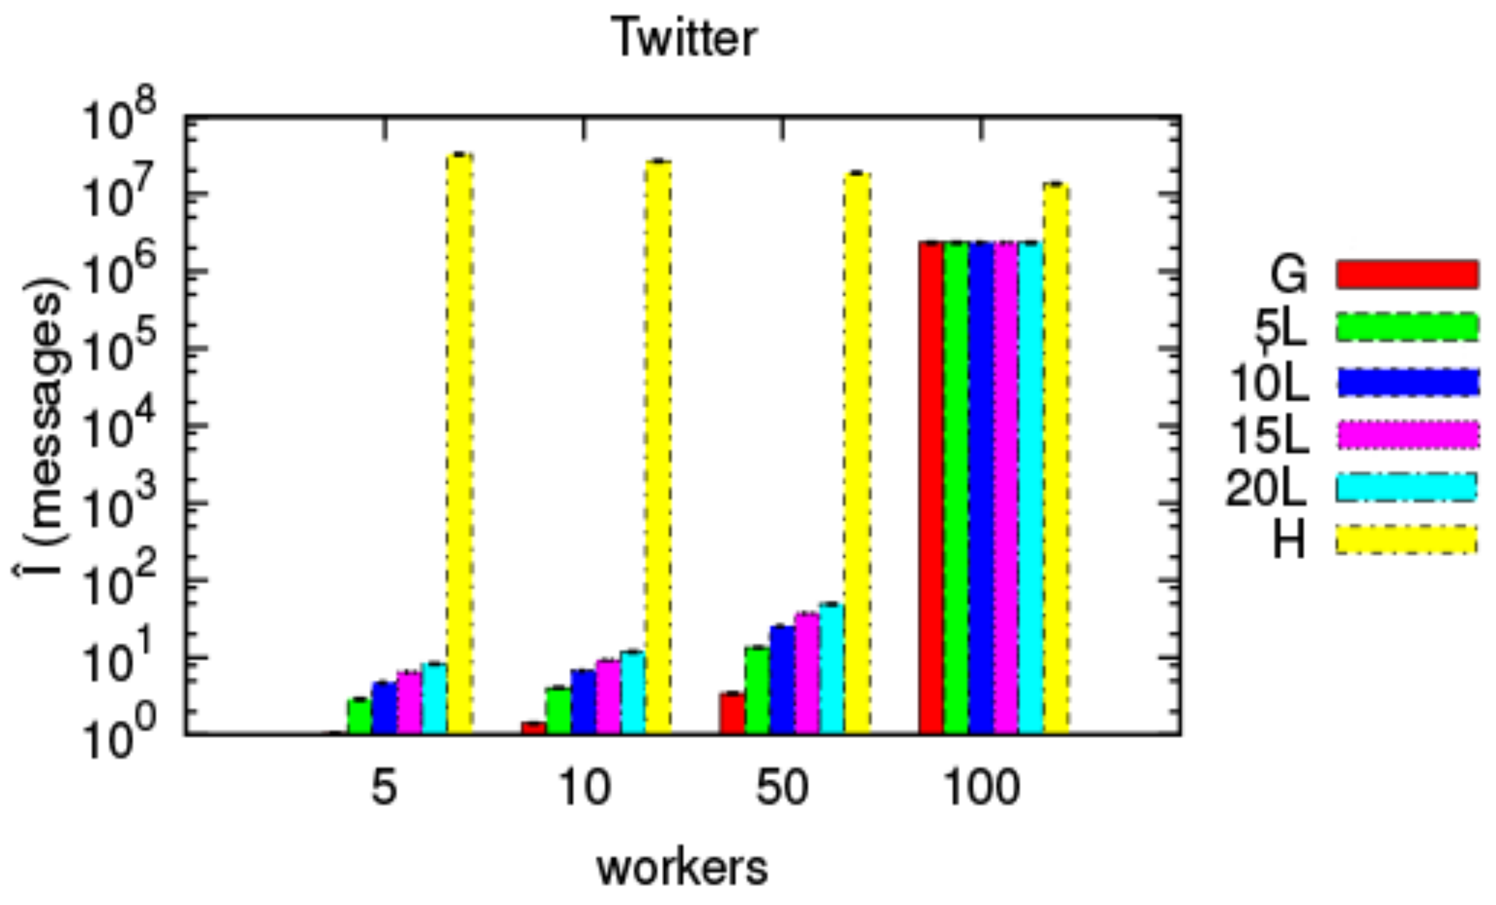

7. Future Work

- Load Balancing with dynamic key migration
- Handling worker churn with Partial Key Grouping
- Applying queuing theory with Partial Key Grouping for load balancing

YAHOO! LABS

Contact: anisu@kth.se, {gdfm, davidgs, kourtel}@yahoo-inc.com, mserafini@qf.org.qa

معهد قطر لبحوث الحوسبة
Qatar Computing Research Institute
Member of Qatar Foundation عضو في مؤسسة قطر

KTH VETENSKAP OCH KONST