

LPW: An Efficient Data-Aware Cache Replacement Strategy for Apache Spark

LPW: 一种面向Spark的高效缓存替换优化策略

Hui Li, Shuping Ji, Hua Zhong, Wei Wang, Lijie Xu,
Zhen Tang, Jun Wei, Tao Huang

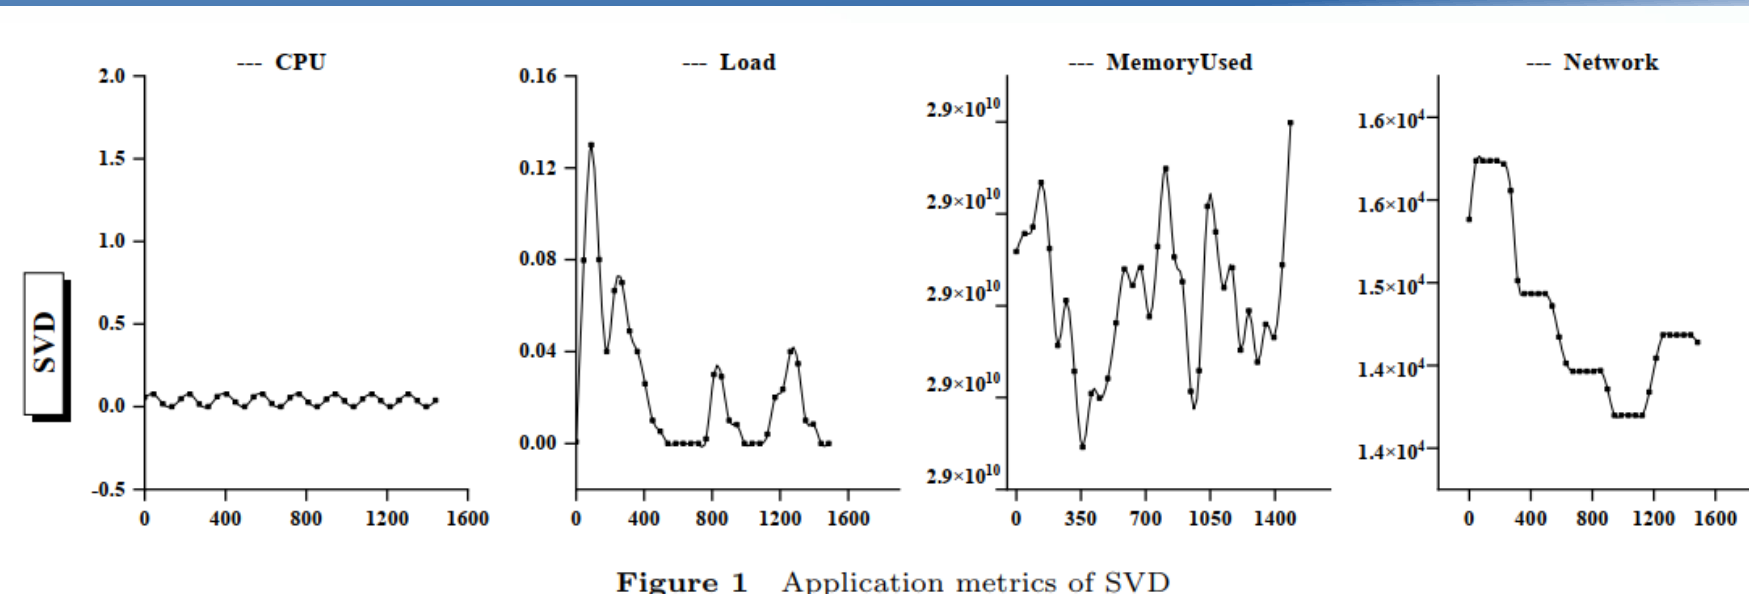

SCIENCE CHINA Information Sciences (SCIS 2022)

DOI:10.1007/s11432-021-3406-5

联系方式: 李慧 (15210848139, lihui2012@otcaix.iscas.ac.cn)

Background & Motivation

Spark users usually try to cache data in memory for re-use to speed up application execution. In real-world, since the storage memory is often not enough to cache all intermediate computing results, frequent cache replacement may happen according to LRU.



- Diversity of applications characteristics.
- Variability of memory resource requirement.
- Uncertainty of cache API usage.

LPW establishes a weight model based on factors to achieve effective use of cached data.

Algorithm

We take comprehensive consideration of different factors affecting performance, such as partition size, computation cost and reference count to find a most suitable partition to be replaced.

Algorithm 1 lpwRep(cachedParts, partition, freeMem)

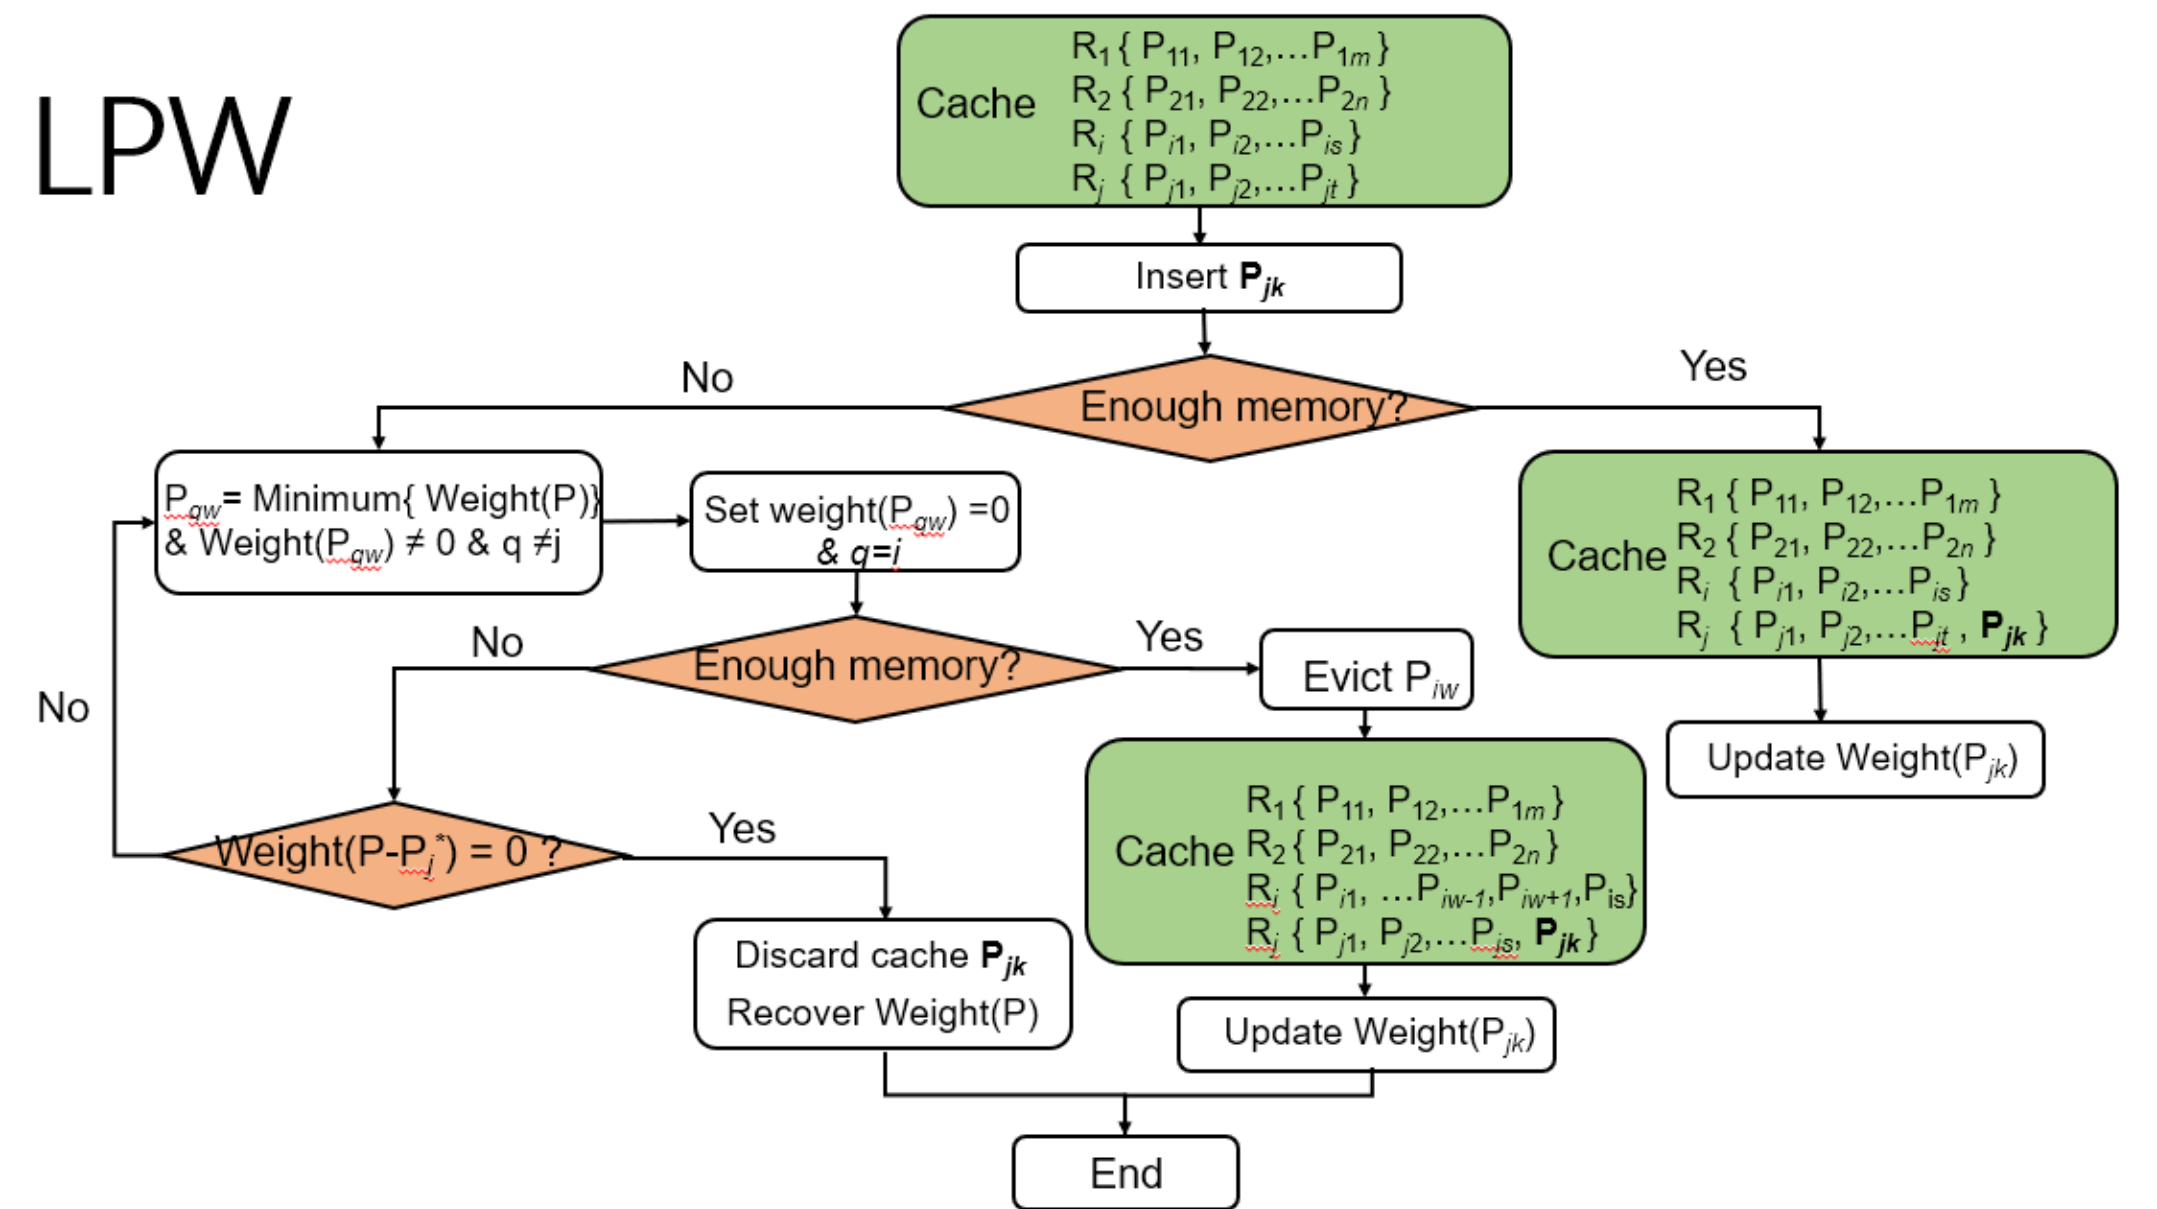
```

1: if partition ∈ cachedParts then
2:   return cachedParts[partition]
3:   break
4: end if
5: weight ← Compute(partition)
6: if freeMem < partition.size then
7:   pQueue ← SortByWeight(cachedParts)
8: end if
9: while freeMem < partition.size do
10:  currPart ← pQueue.pop()
11:  freeMem += currPart.size
12: end while
13: cachedParts.add(partition)
14: freeMem -= partition.size
15: pWeight.add(partition)

```

Methodology

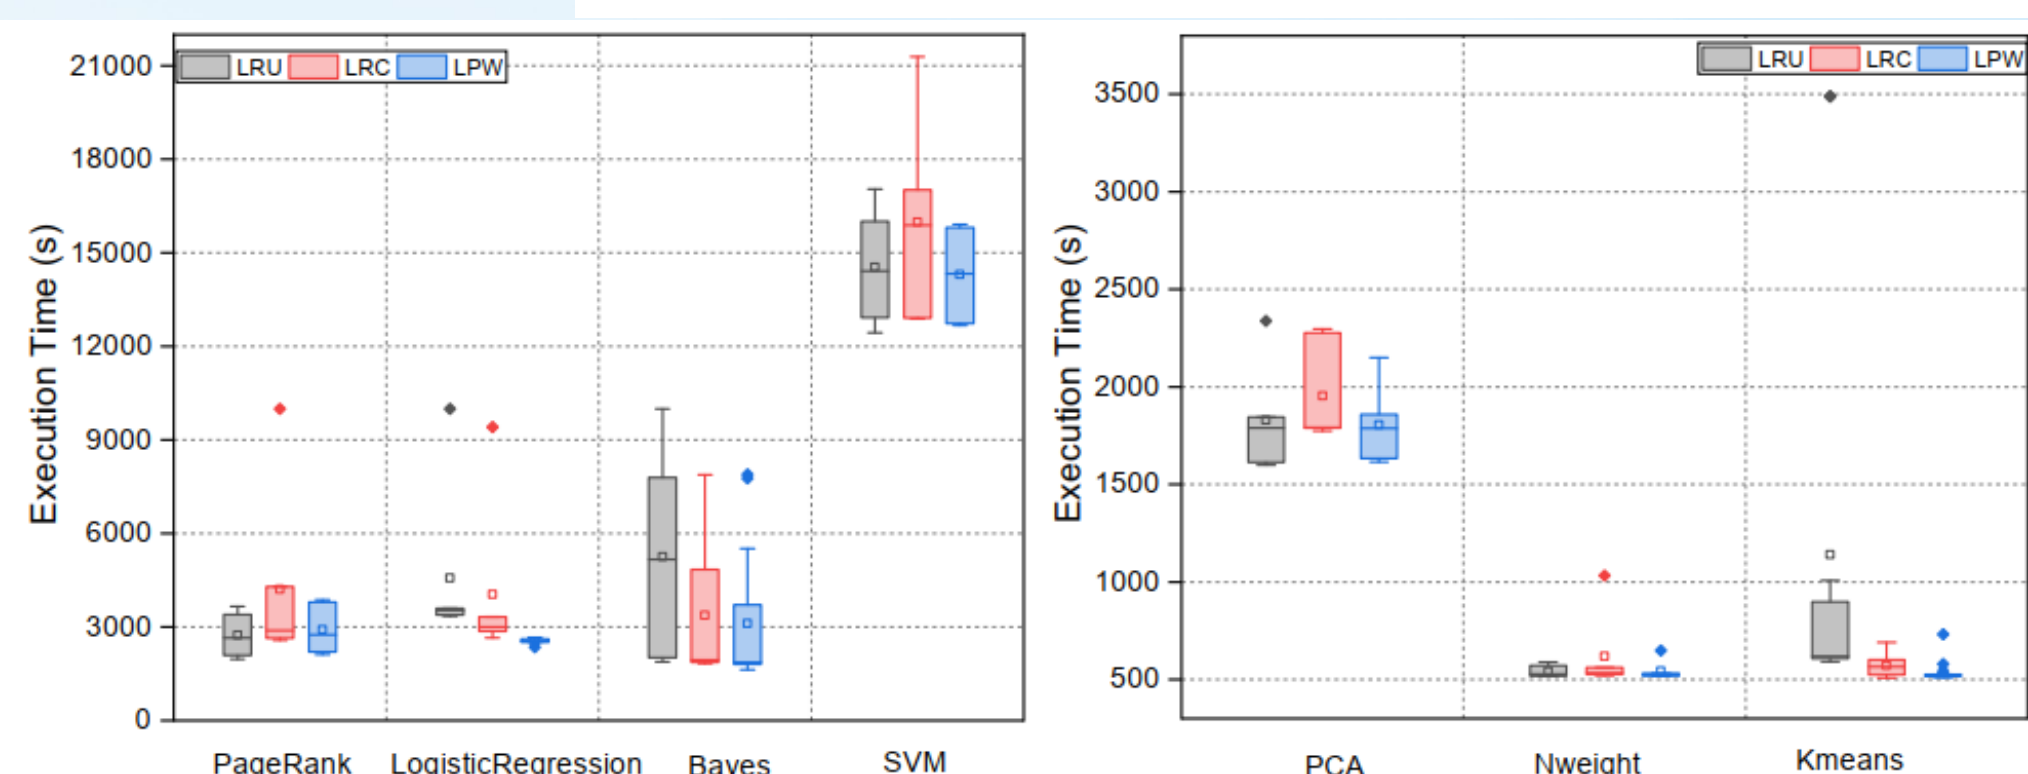
A new block P_{jk} that belongs to RDD_j need to be cached. LPW dynamically check the weight of partitions and make reasonable replacement decision.



Evaluation

We select workloads having *cache* API in HiBench^[1]. LPW can speed up the execution of applications compared to LRU and LRC. Also, LPW could find hot data to keep in memory without causing frequent replacement.

Cluster-NodeID	LRU	LRC	LPW
Node_1	1376	1227	1185
Node_2	1391	1346	1376
Node_3	1191	1145	1124
Total	3958	3718	3685



Conclusion

- We deeply analyze multiple factors of cached partitions that affect application performance.
- We build a weight model to comprehensively evaluate the necessity based on various factors to achieve efficient use of cached data.
- We implement the cache replacement strategy LPW. Our comprehensive experiments show the effectiveness of LPW especially for iterative applications.

[1] "The HiBench benchmark suite: Characterization of the MapReduce-based data analysis" (ICDEW 2010)