LPW: an efficient data-aware cache replacement strategy for Apache Spark

Hui LI[1,2], Shuping JI[1], Hua ZHONG[1*], Wei WANG[1,2,3,4*], Lijie XU[1,2,3,4], Zhen TANG[1], Jun WEI[1,2] & Tao HUANG[1]

[1]State Key Lab of Computer Science, Institute of Software, Chinese Academy of Sciences, Beijing 100190, China;
[2]University of Chinese Academy of Science, Beijing 100049, China;
[3]Nanjing Institute of Software Technology, Nanjing 210000, China;
[4]University of Chinese Academy of Sciences, Nanjing 210008, China

Abstract   Caching is one of the most important techniques for the popular distributed big data processing framework Spark. For this big data parallel computing framework, which is designed to support various applications based on in-memory computing, it is not possible to cache every intermediate result due to the memory size limitation. The arbitrariness of cache application programming interface (API) usage, the diversity of application characteristics, and the variability of memory resources constitute challenges to achieving high system execution performance. Inefficient cache replacement strategies may cause different performance problems such as long application execution time, low memory utilization, high replacement frequency, and even program execution failure resulting from out of memory. The cache replacement strategy currently adopted by Spark is the least recently used (LRU) strategy. Although LRU is a classical algorithm and has been widely used, it lacks consideration for the environment and workloads. As a result, it cannot achieve good performance under many scenarios. In this paper, we propose a novel cache replacement algorithm, least partition weight (LPW). LPW takes comprehensive consideration of different factors affecting system performance, such as partition size, computational cost, and reference count. The LPW algorithm was implemented in Spark and compared against the LRU as well as other state-of-the-art mechanisms. Our detailed experiments indicate that LPW obviously outperforms its counterparts and can reduce the execution time by up to 75% under typical workloads. Furthermore, the decreasing eviction frequency also shows the LPW algorithm can generate more reasonable predictions.

Keywords   Spark, memory, cache replacement, least partition weight, data-aware

1   Introduction

In recent years, the entity, big data is widely used in various fields. Meanwhile, it brings true value to society. For example, it was reported that government departments could save more than $149 billion by using big data techniques in developed European economies, excluding the use of big data to prevent fraud and increase tax revenue[1]. To support various large-scale data processing applications, big data analysis frameworks play a fundamental role in such applications. Many big companies invest significant resources in developing and utilizing big data analysis frameworks, such as Twitter[2], Facebook[3], Google[4] and Alibaba[5][6]. To adapt to different big data processing application scenarios, many big data processing frameworks have been developed, including but not limited to, Apache-Hadoop [1],

---

* Corresponding author (email: zhonghua@iscas.ac.cn, wangwei@otcaix.iscas.ac.cn)

  1) https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/big-data-the-next-frontier-for-innovation.
  2) http://bahir.apache.org/docs/spark/current/spark-streaming-twitter/.
  3) https://databricks.com/blog/2016/08/31/apache-spark-scale-a-60-tb-production-use-case.html.
  4) https://www.zdnet.com/article/google-announces-kubernetes-operator-for-apache-spark/.
  5) https://github.com/alibaba/SparkCube.
  6) https://alibaba-cloud.medium.com/setting-up-spark-on-maxcompute-171b06c7049a.

Tachyon [2], ApacheTez [3], Storm[7], GridGain[8], and Spark. Among them, Spark is one of the most popular frameworks due to its high performance, good scalability, and user-friendly interface. Spark is a distributed in-memory computing solution. Compared with the disk-based solutions, such as Apache Hadoop, Spark is much faster. This difference is largely because Spark can cache the intermediate calculation results in memory and eliminate frequent disk accessing and network transmission [4]. It means the caching mechanism plays an important role in improving the performance of Spark.

Spark provides the resilient distributed dataset (RDD) [5] data structure, which is a fault-tolerant collection of elements that are partitioned across a set of machines on which operations can be performed in parallel. When a programmer uses Spark to process big datasets (stored as RDD), the programmer can use the persist and unpersist application programming interface (API) to cache data in memory or release cached data from memory. However, without good knowledge of the application workloads and the available resources of the Spark cluster, a programmer may invoke these APIs in an unreasonable way [6]. Due to the diversity of cluster resources and application characteristics, unreasonable utilization of caching would bring performance issues, such as insufficient memory resource allocation, data loss, job execution failure, frequent cache replacement, long execution time, and even worse, cluster crashing. These types of problems have often been reported in the open-source Spark community. For example, in SPARK-16440[9], in the Word2Vec.fit() function, excessive memory consumption on the driver will happen when hundreds of continuous training jobs are running. SPARK-14289[10] points that a single eviction strategy might not be enough and may cause expensive time cost. In SPARK-17503[11], caching all data easily causes memory extension and even can cause a memory leak problem. For these issues, specific solutions are proposed. For example, for the SPARK-14289 issue, one of the proposed solutions is called LCS [7]. The key idea of LCS is to use the RDD recovery time cost as the basis for replacement. This design means that when insufficient memory exists, the RDD with the lowest recovery cost will be eliminated from memory. LCS strategy can solve the SPARK-14289 issue; meanwhile, it shows better performance than the default least recently used (LRU) strategy under some scenarios. However, it does not consider many other important factors, such as reference count, memory space occupied of partitions. As a result, its generalization capacity is limited. These reported issues also suggest that even experienced programmers may not use caching in a suitable way. The need for memory varies widely during applications running, and unreasonable configurations often cause many unnecessary performance issues. Many developers are facing significant challenges to achieve high performance for critical business tasks when using Spark to process big data with given limited hardware resources. It means timely deletion of cached data from the memory could be helpful and only caching reasonable data in memory plays an important role in guaranteeing the efficiency of applications and memory utilization.

In a nutshell, Spark users usually try to cache the data in memory for re-use to speed up task execution. Under many situations, too much redundant data could be cached and the memory gradually exhausts. When the memory is not enough, the cache replacement mechanism of Spark, specifically the LRU strategy, starts to work. This process is similar to garbage collection in the Java virtual machine. However, the problem is more complex for Spark. Because Spark needs not only to collect the unused memory but also to collect the data that is potential to be reused in the future. The randomness of cache usage creates a chaotic situation for data analysis. A challenge is that we cannot easily know which cached data is most suitable to be replaced. The LRU strategy adopted by Spark is simple and has been widely used. However, it does not work well under lots of scenarios. Let us illustrate limitations of LRU by using a simple example. In Spark, a RDD may contain a number of partitions distributed in different hosts. To boost performance, the caching mechanism of Spark works at a RDD level. Suppose we have two partitions, p_{ai} and p_{bj}, belonging to two different RDDs. They are sequentially cached in the same node. The times for computing these two partitions are 1 hour and 5 min, respectively. Both partitions will be used in the future. When the cache memory of that node is used up while another partition p_{ck} needs to be cached, based on the LRU algorithm, the partition p_{ai} will be evicted from cache to release memory. However, this is obviously a hasty decision considering the big computing time differences of the p_{ai} and p_{bj} partition: replacing p_{bj} is more reasonable than replacing p_{ai}. There are also scenarios where we should focus on other factors, for example, when the computing times of the p_{ai} and p_{bj} are

7) http://storm.apache.org/.

8) https://www.gridgain.com/.

9) https://issues.apache.org/jira/browse/SPARK-16440.

10) https://issues.apache.org/jira/browse/SPARK-14289.

11) https://issues.apache.org/jira/browse/SPARK-17503.

Table 1   Popular cache replacement strategies

Cache replacement	Frequency	Recency	Reference count	Occupied space	Computation cost
Random	✗	✗	✗	✗	✗
FIFO	✗	✔	✗	✗	✗
LRU	✗	✔	✗	✗	✗
LFU	✔	✗	✗	✗	✗
LRC	✗	✗	✔	✗	✗
LPW	✔	✗	✔	✔	✔

Figure 1   Execution metrics of the SVD application.

Figure 2   Execution metrics of the SVM application.

similar. So, considering different factors in a unified way is very necessary.

The LRU algorithm presents limitations because it does not consider diverse factors affecting the caching performance. Based on our observation, at least four important factors need to be considered when deciding which partition should be replaced: (1) frequency, (2) reference count, (3) occupied space and (4) computation cost. Here, frequency means the period in which data would be accessed again. Reference count indicates the number of times data would be accessed. Occupied space is defined as the size of the block. Computation cost is defined as the processing time of the block. LRU only considers recency. A number of other widely used cache replacement algorithms, such as Random, first in first out (FIFO), least frequently used (LFU) and the recently proposed least reference count (LRC) [8] algorithm, only consider a single, at most two factors. As shown in Table 1, Random does not consider any factor, FIFO only considers recency, LFU only considers frequency, and LRC only considers the reference count. To overcome the limitations of these existing solutions, we propose a new cache replacement algorithm, called the least partition weight (LPW). The goal of the LPW is to minimize the total execution time of different applications. The key idea of LPW is to design a novel unified weight model to evaluate the necessity of caching for each partition. The weight model is calculated based on the analysis of various factors such as computation cost, reference count and partition dependency.

Considering the diversity of different applications is very important. Here we use the singular value decomposition (SVD) and support vector machine (SVM) application from HiBench as the examples to show the diversity of the application workloads and the necessity to consider different factors. Figures 1 and 2 show the execution metrics for SVD and SVM, respectively. The selected execution metrics include CPU, Load, MemoryUsed, and Network. In this situation, CPU indicates the percentage of CPU utilization that occurred while executing at the user level. Load represents the average workload per

minute (at one minute increments). MemoryUsed refers to the amount of used memory. The network represents the number of bytes transferred per second. For the SVD and SVM applications, 0.19 and 107.3 GB data were used as inputs, respectively. The Spark cluster contains four machines and is executed in the Standalone mode. The execution metrics are collected every 45 s. As shown in these two figures, different application workloads could exhibit very different execution behaviors. For example, compared with SVD, the CPU, Load, MemoryUsed, and Network cost are higher for SVM. For SVD, the MemoryUsed footprint fluctuated during execution.

Our proposed LPW algorithm tries to consider the diversity of different applications and the potential continuous variability of memory resource requirements for a single application. LPW dynamically calculates the weight to predict the most appropriate data for making the optimal caching decision. The online replacement decision could be made according to current memory resources to optimize the cache management mechanism of Spark. Specifically, in this paper, we made the following contributions.

• A number of factors that could profoundly affect the execution efficiency of different applications on Spark, such as the cache partition data size, the number of reference count and the calculation cost were analyzed.

• Based on the analysis, the partition weight model and the LPW algorithm were proposed targeting to reduce the overall program execution time by selecting reasonable data to be replaced.

• The LPW algorithm was implemented in Spark and comprehensive experiments were conducted to evaluate the performance. Detailed experiments indicate that LPW could help improve the performance for a variety of applications and obviously outperform its counterparts.

The rest of this paper is organized as follows. Section 2 reviews related background including the Spark programming model, memory management mechanism, and the LRU cache replacement algorithm in Spark. Section 3 illustrates our motivation to design a new cache replacement algorithm. Section 4 proposes the partition weight model after analyzing different factors affecting the partition calculation. The LPW cache replacement algorithm is also presented in this section. The implementation of our solution on the Spark platform is discussed in Section 5. Comprehensive evaluation results are reported in Section 6. Section 7 surveys related work and Section 8 concludes this paper.

2 Background

In this section, we briefly present the background for caching in the Spark framework. Specifically, the programming model of Spark is discussed. After that, Spark's memory management mechanisms are described. Finally, the LRU algorithm, which is adopted by Spark for cache replacement is reviewed.

2.1 Spark programming model

RDD is a core concept of Spark, which refers to an immutable data structure that can be divided into multiple partitions. These partitions can be computed in parallel and are stored in memory or on disk of different nodes of the cluster. Two types of operations can be executed on RDD: transformation (such as map, filter, and reduceByKey) and action (such as count, first, and saveAsTextFile). A transformation operation creates a new RDD from the existing RDDs, while an action operation returns the result after running some computation on a RDD. A Spark job is composed of a series of transformation and action operations. Transformation operations are executed in a lazy way, while action operations are executed in an instant way. Transformation means only when an action operation on a RDD is encountered, are the dependent transformation operations generating that RDD kicked off.

Let us illustrate this process through a simple example as shown in Figure 3. Figure 3(a) presents a segment of source code while Figure 3(b) portrays the corresponding RDDs in addition to their dependency relationship. Since two action operations exist, there are two execution workflows. One is formed by the action A_0 in Figure 3(b) that corresponds to line 4 in Figure 3(a), and the other one is formed by the action A_1 that is corresponding to line 9. In this example, each execution workflow contains a number of transformation operations such as flatMap, filter, reduceByKey, and join. The execution workflow triggered by A_0 can be represented as {input → file → words}. In Figure 3(b), the words RDD is marked as persisted since the persist operation is called at line 3 in Figure 3(a). The partitions of the words RDD are distributed at different nodes of the cluster. The execution workflow triggered by A_1 can be expressed as {words → pairs; pairs → results; pairs → filtered; · · · → counts}. The pairs RDD

```
 1:    val file = sc.textFile("hdfsPth…")
 2:    val words =file.flatMap(x=>x.split(" "))
 3:    words.persist(StorageLevel.MEMORY_ONLY)
 4:    words.count()
 5:    var pairs= words.map(word=>(word,1)).cache()
 6:    val filtered = pairs.filter(item =>
 7:        item._1.startsWith("a")).reduceByKey(_+_)
 8:    val results=pairs.reduceByKey(_+_)
 9:    val counts=results.join(filtered).first()
10:    println("counts"+counts.toString())
11:    words.unpersist()
12:    pairs.unpersist()
```
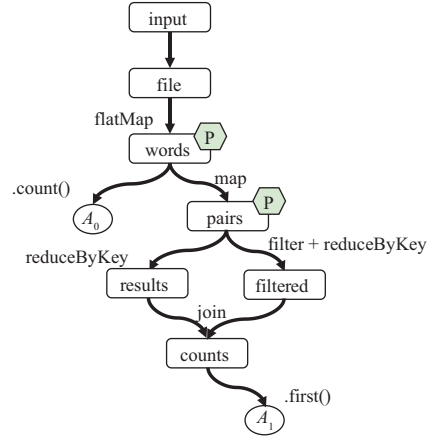
(a)   (b)

Figure 3   (Color online) Code segment (a) and execution graph example (b).

**Table 2**   Execution time of Bayes given different spark.memory.fraction and executor memory configurations

spark.memory.fraction	ExecutorMemory = 15 GB	ExecutorMemory = 30 GB
0.1	8000.934	**2232.523**
0.3	7834.907	7929.69
0.6	**2159.869**	7904.824
0.9	8980.962	7895.035

is also marked as persisted since the cache operation is invoked at line 5 in Figure 3(a). For the second execution workflow, the words RDD could be directly re-used. After the execution of the second workflow triggered by A_1 is completed, partitions of the pairs RDD could also be persisted in the memory.

2.2   Spark memory management

A typical Spark cluster contains a master node and several worker nodes. Each worker node has an executor, which is used to manage memory and execute tasks. The memory managed by an executor can be mainly divided into two types: storage memory and execution memory. Cache-related operations, such as persist, cache, and broadcast consume the storage memory, while the other operations, such as shuffle, join, and sort will take up the execution memory. The configuration spark.executor.memory is used to determine the total size of an executor's available memory, while the configuration spark.memoryFraction is used to decide the ratio of the storage memory. For example, if the value of spark.memoryFraction is set at 0.5, the executor can use a maximum of 50% of the memory for caching. The remaining 50% of the memory will be used to store the dynamically created temporal objects during the execution of different tasks. spark.memoryFraction is an important configuration that could significantly affect the performance of the whole system. However, due to the diversity of the workloads, it is often hard to provide an ideal configuration value. A different setting might cause severe performance problems such as longer delays in executing the task. Even in the same application with a different configuration, the performance can be different. We conduct an experiment using 70 GB data size as input for Bayes workload deployed in a cluster with four nodes. The result is shown in Table 2. The performance is good given a configuration spark.memoryFraction value of 0.6 under the executor memory value of 15 GB, while performance is poor under the executor memory value of 30 GB. The execution time increases by about 3 times, from 2159 to 7904 s. By running more experiments, we found that the best performance occurs when the executor memory value is set as 30 GB and the value of spark.memoryFraction is set as 0.1.

   In this experiment, we also analyze the cache replacement that happened on different work nodes. Specifically, we record the cache replacement frequency on each worker node. The result is shown in Figure 4. The red vertical line indicates that the replacement occurs. The blue line shows the free memory size after a block eviction. Compared with slave1, more frequent replacements occur on the nodes (slave2 and slave3). In real-world applications, since the storage memory is often not enough to cache all intermediate computing results, frequent cache replacement may happen.

Figure 4 (Color online) Free memory and eviction happened in slaves. (a) Slave1; (b) slave2; (c) slave3.

```
1: var g= sc.parallelize("hdfs://...")
2: for (i<-1 to N) do:
3:    var g=function1(g, i)
4:      g.persist()
5:      g.first()
6: end for
7: var result=g.first()
```

(a)

```
1: var g=sc.parallelize("hdfs://...")
2: for (i<-1 to N) do:
3:    var prevg=g
4:    var g=function1(g, i)
5:      g.persist()
6:      g.first()
7:      prevg.unpersist()
8:    end for
9: var result=g.first()
```

(b)

Figure 5 (Color online) Code segment for an example with cache API. (a) Native version; (b) improved version.

Inappropriate configuration settings can easily cause performance problems such as frequent cache replacement. To eliminate this type of problem, an effective cache replacement strategy becomes necessary and useful. In a nutshell, keeping the most useful data in memory and right evicting other data from the storage memory is one of the most important techniques for improving the performance of the whole system.
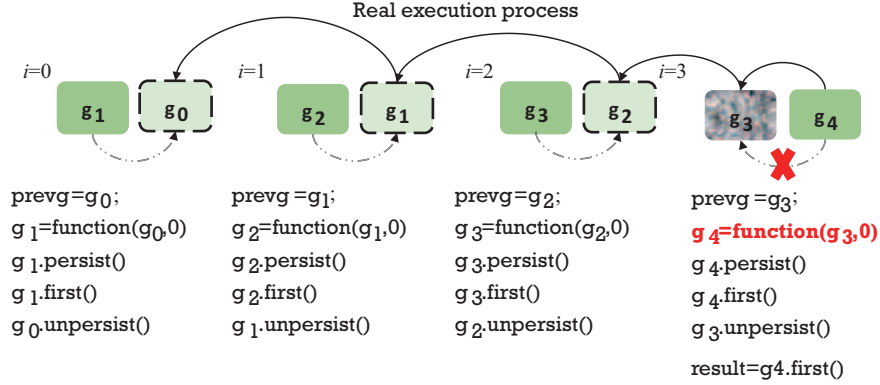
2.3 LRU in Spark

To manage the limited available memory, Apache Spark uses the LRU eviction strategy as default policy. The basic idea of LRU is to record each RDD's access times. When a new data is generated and the storage memory is not enough to cache the new data, the LRU algorithm will evict the cached partition belonging to the least recently accessed RDD. It is worth noting that the selected partition and the new partition cannot belong to the same RDD. This design is to prevent partitions of the same RDD from looping in and out. LRU's mechanism of evicting the least recently used data is based on the assumption that the recently used data is more likely to be reused again. This premise is not appropriate for many applications as will be shown in Section 3.

3 Motivation examples

Challenges of manual cache controlling. Figure 5 presents an iteration example, which shows the challenges of solely relying on programmers to manually control the caching and the necessity of adopting an effective cache replacement mechanism. In this example, we will analyze the limitations of two different cache-based implementations, as shown in Figure 5(a) and (b), respectively. In Figure 5(a), the application first loads the data from HDFS and initializes the execution graph (line 1). The iterative computation then starts, and many new execution graphs are created (lines 2–6). At last, the application outputs the final result (line 7). To improve the performance and reduce recovery time for fault tolerance, developers usually cache intermediate data into the memory by calling the persist API (line 4). For each iteration, a new execution graph would be created. As the number of iterations increases, the number of caches also increases ($g_1 \rightarrow g_2 \rightarrow \cdots \rightarrow g_N$). In this case, many memory resources are wasted.

A common practice for solving this problem is to add a variable in an ad hoc manner. As shown in Figure 5(b), Spark users can assign the result of a previous calculation to the variable (prevg in line 3) and cache the result of the current iteration calculation (line 5). After the current iteration calculation is finished, prevg will be released from memory by calling the unpersist API (line 7). The advantage of this approach is that only one intermediate calculation result is cached at any time, indicating that

Real execution process

$i=0$ g_1 g_0 $i=1$ g_2 g_1 $i=2$ g_3 g_2 $i=3$ g_3 g_4

prevg $=g_0$;
g $_1$=function(g_0,0)
g $_1$.persist()
g $_1$.first()
g $_0$.unpersist()

prevg $=g_1$;
g $_2$=function(g_1,0)
g $_2$.persist()
g $_2$.first()
g $_1$.unpersist()

prevg $=g_2$;
g $_3$=function(g_2,0)
g $_3$.persist()
g $_3$.first()
g $_2$.unpersist()

prevg $=g_3$;
g $_4$=function(g $_3$,0)
g $_4$.persist()
g $_4$.first()
g $_3$.unpersist()
result=g4.first()

**Figure 6** (Color online) Execution trace of the simple example when some partitions are lost suddenly.

when the $(N+1)$-th job is finished, the cached result from the N-th calculation iteration will be released. This seems to be a reasonable solution. However, serious performance deterioration could happen prior to data losses. Suppose that some data loss happens because of the crash of a work node, which is very common in large scale Spark clusters. In this case, data recalculation will be triggered according to Spark's fault tolerance mechanism. For example, suppose partial data of g_{N-1} are accidentally lost before g_N is calculated. Spark needs to recompute the data from g_0 to g_{N-1}, which could be a very time consuming process. The application may suffer from severe performance deterioration. Additional observation about recomputing the data from g_0 to g_{N-1} is shown in Figure 6. Here N is 3. When the program is running normally, the program can read g_3 directly in the memory and perform subsequent calculations while computing g_4. Suppose a buggy scenario that the program is executing the function $g_4 = \text{function}(g_3, 0)$ occurs, and some partitions of g_3 are lost suddenly. The program cannot read g_3 directly in the memory since g_3 has been lost from the memory. To perform subsequent calculations, the program needs to calculate g_3. When computing g_3 in the 2nd iteration, g_2 is evicted from the memory. So, we need to calculate g_2, continue to compute until g_0 is calculated. All in all, before computing g_4, calculation according to this order $\{g_4 \rightarrow g_3 \rightarrow g_2 \rightarrow \cdots \rightarrow g_0\}$ needs to be performed. So, this process means this common practice is still not an ideal solution, and a good cache replacement mechanism is desired for different scenarios. Although evicting cached data according to the LRU would happen, applications may still suffer inefficient performance issues.

Necessity of considering different factors. Figure 3 is used as an example to show the necessity of adopting a cache replacement strategy in addition to the necessity of considering different workload and environment related factors.

Before different features of operators in depth are observed, a simple typical test example is run, and the execution time is measured. Based on the example in Figure 3, this experiment is run in the local mode with 8 GB memory and 1 kB text file as input. It could be observed that the execution time of join operator in Stage$_3$ is nearly 3.7 times more than the time of map operator in Stage$_2$. The calculation time of these two operators is 70 and 19 ms, respectively. In this case, map belongs to the operator of the value data type, which merges the elements in each set to form a new set. Another operator join belongs to the key-value data type, which puts the same key data into a partition. Through our experiments, the features of operators are shown to be diverse. Thus, a good cache replacement should be fully aware of the features of operators and then make appropriate replacement decisions.

When the memory is insufficient and blocks need to be replaced, according to the LRC, words will be retained because of being accessed twice, which may not necessarily improve execution efficiency. This drawback occurs because the time to calculate words is short and the performance bottleneck lies in the execution time of the join operator. Another scenario can occur in which the amount of input data is gigantic, removing the cached words from the memory according to the LRU still could not effectively solve the performance issues. This issue occurs because the bottleneck lies in calculation of words. Obviously, either LRC or LRU is difficult to guarantee applications will have a better performance in complex environments. The advantage of the LRU mechanism to keep reasonable data in memory is not obvious. In short, application execution is affected by various factors, such as computing resources, and computation time of operating on blocks. In other words, the calculation time between RDDs is diverse and the free memory is changing. As a result, predicting future data access based on single historical

Figure 7 (Color online) Cache decisions considering different factors. (a) Reference count; (b) computation cost; (c) occupied space.

information (such as recency, frequency) to make a replacement decision for the big data computing framework is not effective. A good cache replacement strategy is supposed to be adaptive to various applications even though they have different characteristics.

4 Our solution

Optimizing memory usage is one of key techniques for improving the Spark application execution performance[12)13)]. Keeping the right data in memory for re-use is an important way to improve memory utilization and reduce recalculation cost. We design the LPW algorithm to guide Spark to choose the reasonable data for cache replacement. In this algorithm, the concept of weight to measure the value of caching a block is proposed. A larger weight indicates that the block is more valuable during the application execution and is more likely to be reused in the future. Blocks with smaller value are more likely to be evicted compared to blocks with higher weight value. Many factors could affect the execution time of Spark applications. In our design, the calculation of weight considers several important factors: reference count of the block, occupied space, the calculation cost, and the historical execution information (pastmod). These factors are comprehensively considered together. Compared with the current LRU algorithm adopted in Spark, which only considers the recency information of a block, our weight-based algorithm should be more feasible and effective for many different applications.

4.1 Considering factors

Once an application is submitted to a Spark cluster, one or several jobs will be created and executed. RDD is the base component for the execution of a job. A triggered job is composed of a series of transformation operations on a RDD. In this paper, r_i is used to represent a RDD, and τ is used to represent a job. The execution trace of a job can be defined as $\{r_0 \to r_1 \to r_j \to \cdots \to r_k\}$. Now suppose two jobs τ^p and τ^q for an application are present. If r_i is used in both τ^p and τ^q, it is better to keep r_i in memory. Figure 7 shows a sample lineage graph on which the vertices represent the RDDs. When a vertex is used by more than two jobs, it is marked in black. In this figure, the transformation between RDDs is connected by edges and transformation time is defined by the symbol T. In the following sections, Figure 7 is used as an example to explain the detailed meaning of different factors and how they can help determine which is the best cache to replace.

Reference count. When the execution time is the same for different transformation operations, the completion time of the two jobs will be shown in (1). The first job (left most in Figure 7(a)) Time_1 is counted by T_{00}, T_{10}, T_{11}, T_{12}, and T_{13}. The fourth job Time_4 is counted by T_{00}, T_{10}, T_{30}, T_{40}, T_{41}, T_{42} and T_{43}. It is not difficult to observe that one endpoint of T_{00} is used by seven jobs and the value of the reference count is 7. Meanwhile, one endpoint of T_{50} is used by three jobs, and the value of the reference count is 3. When not enough memory exists, the best caching replacement decision is to cache the endpoint of T_{00} because of the larger reference count. Figure 7(a) shows the optimal caching according to the reference count.

$$\begin{aligned}
\text{Time}_1 &= T_{00} + T_{10} + T_{11} + T_{12} + T_{13}, \\
\text{Time}_4 &= T_{00} + T_{10} + T_{30} + T_{40} + T_{41} + T_{42} + T_{43}.
\end{aligned} \tag{1}$$

12) https://unraveldata.com/common-reasons-spark-applications-slow-fail-part-1/.
13) https://unraveldata.com/common-failures-slowdowns-part-ii/.

In short, when the reference count of the RDD is larger, therefore, it is more necessary to cache for re-use.

Computation cost. The costs between different RDDs are often very different. The detailed execution time can be easily calculated from the task that completes the calculation of these RDDs. RDD_i and RDD_j are the different transformation operators. Suppose that it will take one hour to calculate RDD_i while the cost of RDD_j is only one ten minute. In this situation, RDD_i denotes the end vertex of T_{10}. RDD_j denotes the end vertex of T_{40}. Two jobs consisted of RDD_i and RDD_j. The completion time on the partial path of two jobs is shown as

$$
\begin{aligned}
\mathrm{Time}_m &= \mathrm{RDD}_i.\mathrm{cache}() + \mathrm{RDD}_j, \\
\mathrm{Time}_n &= \mathrm{RDD}_i + \mathrm{RDD}_j.\mathrm{cache}(),
\end{aligned}
\tag{2}
$$

where $\mathrm{RDD}_i.\mathrm{cache}()$ and $\mathrm{RDD}_j.\mathrm{cache}()$ denote caching their respective data into the memory. As shown in (2), Time_m would be longer than Time_n when the calculation time of the RDD_i is larger. It is obvious that, caching RDD_i rather than caching RDD_j in memory is beneficial for improving execution efficiency. A larger computation cost implies the need for a longer calculation time. The best caching replacement decision for this example is to keep the end vertex of T_{10} because of the larger computation cost. Figure 7(b) shows the optimal caching according to the computation cost. In short, the computation cost of the RDD is smaller and more likely to be evicted for releasing memory.

Occupied space. RDDs are divided into partitions that are distributed on each node as blocks occupying memory. Based on the example in Figure 3, experiments with caching all data are performed. With an input data size of 50 GB, we find that different cached partitions belonging to the pairs occupy different memory sizes. One partition is 72.1 MB and another one is 1281.1 MB. The larger the space occupied by the partition is, the less the necessity to cache the block is. Excessively abundant data occupy a significant portion of storage space, which may easily waste considerable memory resources and even slow down the execution efficiency. Evicting the partitions occupying a large space in memory can help decrease the application execution time. Suppose the block size of the end vertex of T_{50} is much larger than other blocks. Figure 7(c) shows the optimal caching according to the occupied space.

Pastmod. Blocks frequently used in past jobs are more likely to be accessed during subsequent job calculations. If data have been accessed more frequently during the previous period of time, this data is more likely to be used again in the future. Pastmod indicates the reference count of the partition which is computed in the completed job. For example, 10 jobs have been executed before the 11th job is executed, and P_{13} is used in 5 jobs that have been executed, and then the value of pastmod is calculated by $5/10$ to get 0.5.

Overall, evicting the suitable data block in a timely manner when the memory is not enough is a key technique for optimizing memory usage and improving performance. As shown in the example in Figure 7, we should comprehensively consider many factors to make a good cache replacement decision.

4.2 Partition weight model

All factors mentioned in the above section were comprehensively considered in this study. To identify the valuable partition that needs to be kept in memory, the weight of a partition is calculated. For a partition $\mathrm{partition}_i$, the formula to calculate its weight is shown as

$$
\mathrm{weight}_i = \frac{\mathrm{cost}_i \times \mathrm{ref}_i \times (1 + \mathrm{pastmod}_i)}{\mathrm{size}_i},
\tag{3}
$$

where cost_i is the execution time to calculate a $\mathrm{partition}_i$; ref_i is the number of jobs that depend on the current partition; $\mathrm{pastmod}_i$ is the reference count of the partition which is computed in already finished jobs; size_i is the memory space size to store the current partition.

4.3 LPW algorithm

Our proposed LPW algorithm is designed to adapt to different Spark applications, which are executed based on in time data perception and dynamic adjustment. Algorithm 1 illustrates the pseudocode of the algorithm to identify reasonable partitions for the cached blocks in memory. The variable partition denotes the new block needed to be cached in memory. The variable cachedParts is a list of data blocks that are already cached in memory. The variable freeMem represents the remaining memory of the node.
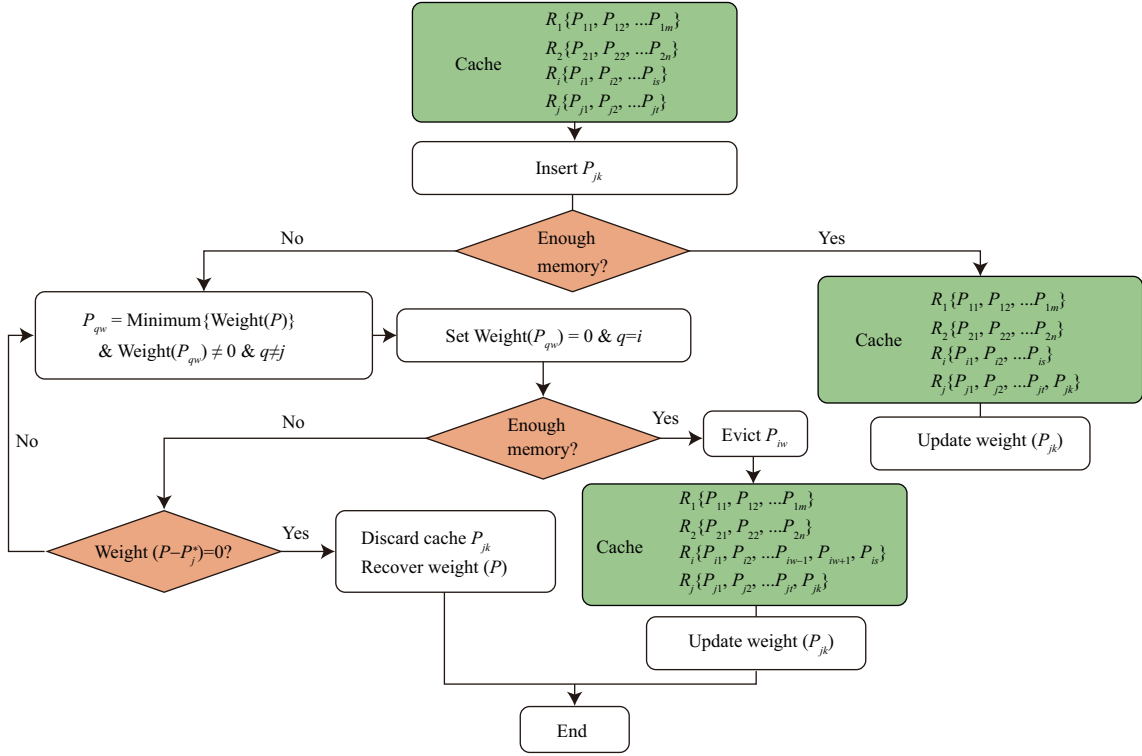
Figure 8   (Color online) The replacement process of cache replacement strategy with LPW.

The variable pWeight denotes the weight value of each data block cached in memory. First, if the partition is found in cachedParts, the data do not need to be cached since the block is hit (lines 1–4). The free memory and the weight in the pWeight (lines 13–15) were updated. Second, the weight value of each partition is calculated (line 5). The calculation method is shown in (3). During the replacement process, the weight value of partition is sorted in ascending order and added the sorted weight to the pQueue (lines 6–8). Before making a decision to cache, LPW continues to check whether the space is enough when currPart is removed from memory (lines 9–12). In the beginning, the partition from memory with the smallest weight in pWeight is removed and free memory space is tracked to cache partition. If enough space is available to cache partition, LPW would stop to evict block from memory and update the value of pWeight and cachedParts (lines 13–15).

Algorithm 1 can be used to traverse the cached partition to find a most suitable partition to be replaced, based on the prediction of which data blocks have high probabilities to be reused. To make the algorithm more defined, we use Figure 8 as a flow chart to elaborate the execution details of the LPW algorithm.

Algorithm 1 lpwRep(cachedParts, partition, freeMem)

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

A weight table named Weight(P_{jk}) is constructed and maintained the weight of each partition. Note that, all initial weights in the table are zero at the beginning. Suppose a new block P_{jk} that belongs to
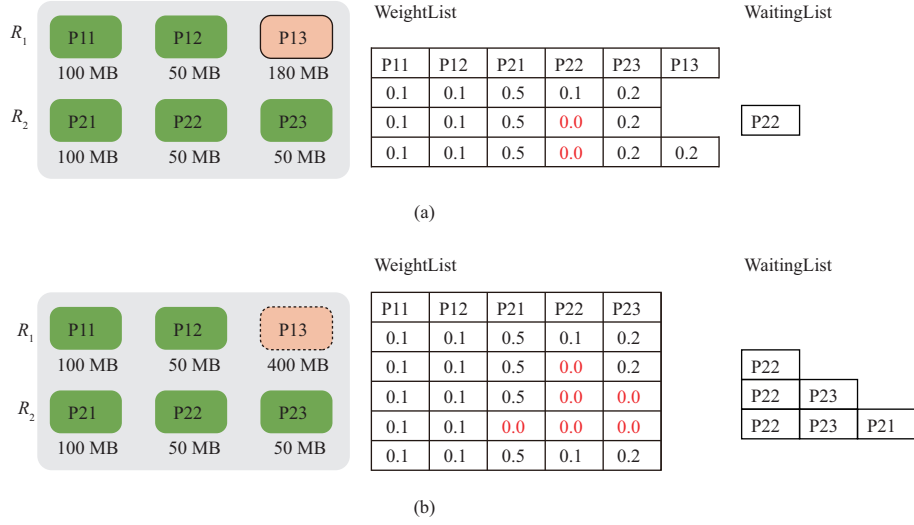
RDD$_j$ needs to be cached. LPW would first check free memory space after which it would directly add the new partition into memory and update the weight value of the weight table if enough space is available. A new cachedParts would then be generated. A list of cached data blocks $R_j\{P_{j1}, P_{j2}, \ldots, P_{jt}\}$ would also be updated to become $R_j\{P_{j1}, P_{j2}, \ldots, P_{jt}, P_{jk}\}$. If enough space is not available to put P_{jk}, the block with the smallest weight value (not zero) would be fetched from the table and they do not belong to the same RDD. The formula can thus be expressed as shown below as $P_{qw} = \text{Minimum}\{\text{Weight}(P)\}$ & Weight(P_{qw}) $\neq 0$ & $q \neq j$. The weight value of the partition is set with its minimum at zero and moves to waitingList. Spark continued to check whether enough space is available to cache P_{jk}. LPW continues to search for the smallest weight block and the value is not zero. More and more space is obtained by releasing partitions in the waitingList. In detail, after the partition P_{iw} is chosen for eviction, enough space can be finally obtained to cache P_{jk}, and then partitions could be deleted from waitingList while concurrently, evicting them from the memory and updating the value of P_{qw} in the weight table. A new cachedParts is generated and a list of cached blocks $R_j\{P_{j1}, P_{j2}, \ldots, P_{jt}\}$ is updated to $R_j\{P_{j1}, P_{j2}, \ldots, P_{jt}, P_{jk}\}$. Meanwhile, because of the eviction of P_{iw}, $R_i\{P_{i1}, P_{i2}, \ldots, P_{is}\}$ changes to $R_i\{P_{i1}, P_{i2}, \ldots, P_{iw-1}, P_{iw+1}, P_{is}\}$. Note that, when P_{iw} is replaced from the memory, it is also directly removed from the weight table. If all partitions which do not belong to the same RDD with P_{qw} are released, enough space is still not available to cache P_{jk}, then give up caching P_{jk} and recover the value of the weight table. In Figure 8, P_j^* indicates the weight value of all partitions except P_{jk}.
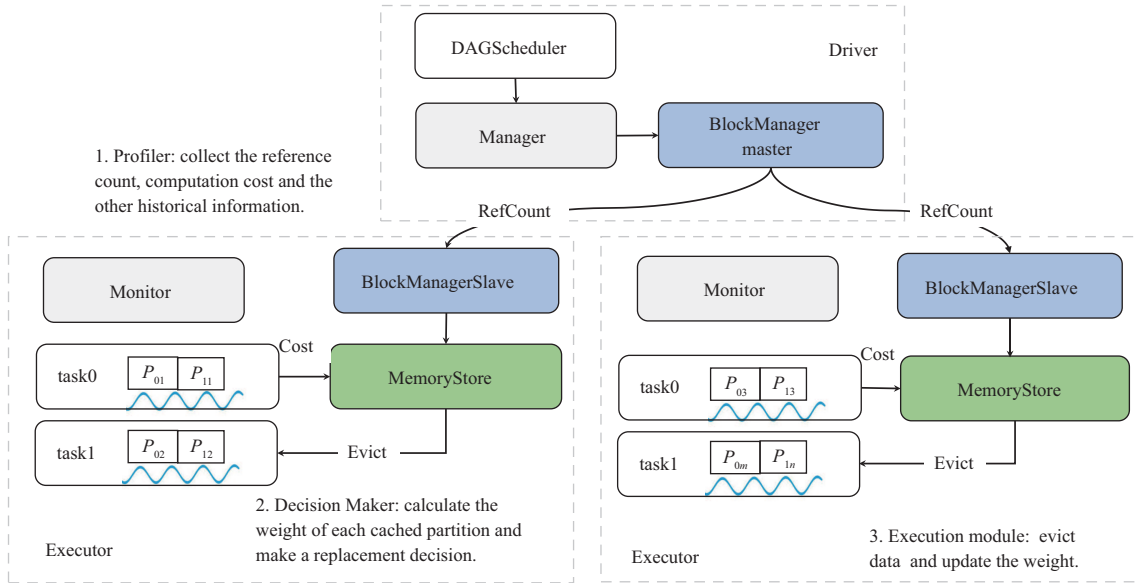
4.4 A simple example for LPW

A simplified example of a self-adaptive scenario in pseudocode is given below. The initial value of the weight table is zero. In Figure 9, five partitions $\{P_{11}, P_{12}, P_{21}, P_{22}, P_{23}\}$ are cached and located on the same node. Partitions P_{11} and P_{12} belong to RDD$_1$, and other partitions such as P_{21}, P_{22} and P_{23} belong to RDD$_2$. New data P_{13} is to be cached. The total memory size of the worker Total$_{\text{mem}}$ is initialized to 500 MB.

In Figure 9(a), the memory space of P_{13} is 180 MB ($S_{p13} = 180$ MB) and the cached sequence pQueue is $\{P_{11}, P_{12}, P_{21}, P_{22}, P_{23}\}$. The remaining memory could be easily calculated according to the formula Free$_{\text{mem}} = $ Total$_{\text{mem}} - S_{R1} - S_{R2}$ to obtain enough free memory space (150 MB). We compare free memory size with P_{13} space and find that not enough space for caching P_{13} data exist (Free$_{\text{mem}} < S_{p13}$). Thus, replacement process starts. First, Spark traverses the partition weight table and selects partitions where P_{13} does not belong to the same RDD. The pre-selection list $\{P_{21}, P_{22}, P_{23}\}$ is obtained. Second, Spark searches for the minimum weight value among the list and finally selects P_{22} as a block to be replaced. Third, the weight value of P_{22} is set to zero and moves to the waitingList. The remaining memory is updated to 200 MB according to formula Free$_{\text{mem}} = $ Total$_{\text{mem}} - S_{R1} - S_{p21} - S_{p23} + S_{p22}$. Free memory size is continuously updated with P_{13}, and it is found that enough space to cache P_{13} is available. Finally, the replacement decision according to LPW is made: P_{22} is selected for eviction from memory to cache P_{13}. We get a new cached sequence pQueue $\{P_{11}, P_{12}, P_{13}, P_{21}, P_{23}\}$ and the remaining memory of the node Free$_{\text{mem}}$ is changed to 20 MB.

We continue to observe another situation. As shown in Figure 9(b), P_{13} still represents a new data to be cached. The memory space occupied by P_{13} is 400 MB and the cache sequence pQueue is $\{P_{11}, P_{12}, P_{21}, P_{22}, P_{23}\}$. First, Spark checks if enough memory to store P_{13} is available, and then begins to traverse the partition weight table and find the pre-selection list $\{P_{21}, P_{22}, P_{23}\}$. Second, the weight value of the minimum P_{22} is set to zero and moves to the waitingList. The remaining memory is calculated using the equation Free$_{\text{mem}} = $ Total$_{\text{mem}} - S_{R1} - S_{p21} - S_{p23} + S_{p22}$ to yield 200 MB. In this situation, not enough memory is available to cache P_{13}. We continue to choose the minimum weight value as 0.2. Partition P_{23} is chosen for eviction after which the free memory size is changed to 250 MB. The weight value for each partition is also updated to $\{0.1, 0.1, 0.5, 0.0, 0.0\}$. Third, P_{21} is selected sequentially from the pre-selection list $\{P_{21}\}$. Fourth, the remaining memory is computed using the equation Free$_{\text{mem}} = $ Total$_{\text{mem}} - S_{R1} + S_{p21} + S_{p23} + S_{p22}$, which increases to 350 MB from 250 MB. Although all partitions consisting of P_{21}, P_{22} and P_{23} are all released, there is still not enough memory to cache P_{13}. Finally, the replacement decision according to LPW is made: giving up caching P_{13} and resetting the weight value to the initial value $\{0.1, 0.1, 0.5, 0.1, 0.2\}$.

WeightList

P11	P12	P21	P22	P23	P13
0.1	0.1	0.5	0.1	0.2	
0.1	0.1	0.5	0.0	0.2	
0.1	0.1	0.5	0.0	0.2	0.2

WaitingList

P22

(a)

WeightList

P11	P12	P21	P22	P23
0.1	0.1	0.5	0.1	0.2
0.1	0.1	0.5	0.0	0.2
0.1	0.1	0.5	0.0	0.0
0.1	0.1	0.0	0.0	0.0
0.1	0.1	0.5	0.1	0.2

WaitingList

P22		
P22	P23	
P22	P23	P21

(b)

Figure 9 (Color online) Cache replacement example of LPW. (a) 180 MB; (b) 200 MB.

Figure 10 (Color online) System architecture of the LPW solution.

5 Implementation

LPW first collects information by modifying Spark implementation and then calculates the weight value online on executors. After applications are submitted to Spark Cluster, DAGScheduler would divide directed acyclic graph (DAG) of applications into different stages after which each stage will be divided into multiple tasks, which will be assigned to executors for execution by TaskScheduler. It is not difficult to obtain a lineage graph from the DAGScheduler, which provides an opportunity to design a data-aware cache replacement strategy that considers the characteristics of applications. By default, RDDs are divided into multiple partitions on executors of the cluster by hash partitions and then handed over to Tasks on Worker nodes. For example, partitions such as P_{01}, P_{02} of RDD_1 are in Worker1 and the rest partitions are located in Worker2. Partitions such as P_{01}, P_{12} of RDD_2 are located in Worker1 and the remaining partitions are located in Worker2. We first present system architecture that enables the implementation of the LPW in Spark and then elaborate on the implementation details. For comparison, LRC in Spark is also prototyped. Figure 10 shows an architecture overview of LPW. Our extended modules are shown in the shaded box. The Manager module is located on the Driver node, and the Monitor modules are distributed on Executor nodes.

Table 3 Key APIs for LRW implementation

Key API	Location	Description
putOrUpdateWholeBlock	MemoryStore.scala	Update the weight value of the partition
updateRef	MemoryStore.scala	Update local RefCount
updateCost	MemoryStore.scala	Update local cost of the partition
calculateWeight	MemoryStore.scala	Calculate the weight for each partition
broadcastRef	BlockManagerMasterEndpoint.scala	Send RefCount to the BlockManagerSlave on the worker
profileRefCountStageByStage	DAGScheduler.scala	Calculate RefCount that needs to be cached in a single job
profileRefCountOneStage	DAGScheduler.scala	Calculate RefCount that needs to be cached in a single stage
handleJobSubmitted	DAGScheduler.scala	Calculate RefCount
calculateTrueTime	ShuffleMapTask.scala ResultTask.scala	Get the calculation time of a partition

Table 3 summarizes the key APIs for LPW. Here, the implementation details are introduced, including system architecture for the overall design and key APIS for information collection. All modules in this experiment are based on existing modules of Spark. The DAGScheduler module on the Driver node learns the dependency relationship between RDDs, allowing us to obtain the reference count (RefCount) which defines the number of times on which each RDD is relied during the calculation. RefCount is sent to WorkerManager BlockManagerSlave through BlockManagerMasterEndpoint. The BlockManagerSlave on the Worker is to keep sending heartbeats to the BlockManagerMaster to update the block (partition) information. Once the BlockManager object is created, it will also create a MemoryStore to manage the Block information. By instrumenting MemoryStore.scala, the size of the partition and the time spent in calculating could be obtained. The number of times on which each RDD is relied on in the completed job calculation is obtained and used to calculate the partition weight according to the model. The weight is calculated and saved in the worker in which the corresponding task is executed. After each job is completed and when the dependency count information of the new job is passed to the MemoryStore, the MemoryStore will record its current dependency information such as the number of times RDD is referenced in the previous job and receives the new dependency information. Whenever insufficient memory on the Worker is found, the MemoryStore will send the data of the current node and select the appropriate partition to complete the replacement by using LPW decision information. An RDD with more than one partition is distributed on each node in the cluster. The life cycle of the RDD is the period from the first job to the last job that relies on the partition to complete the calculation. The computation cost of a partition contains calculation time and communication time. The ShuffleMapTask and ResultTask on the workers in the cluster record the running time of each task and the calculation cost of each partition can be obtained through code instrumentation. The BlockManager is used to manage data in Spark. We can obtain the size of memory space occupied by partitions by monitoring the API. The BlockManager on each worker makes the replacement decision based on weight and sends information to Driver.

6 Evaluation

In this section, we evaluate the performance of our proposed LPW algorithm by using the Hibench [9][14] benchmark, which contains a number of typical application workloads. Through several experiments, we try to answer the following three questions.

RQ1. Can LPW effectively address the performance issues faced by applications under the environment with limited memory resources?

RQ2. Compared to existing solutions, can LPW obviously speed up the execution of Spark applications?

RQ3. Does LPW introduce extra overhead during the execution of different Spark jobs?

In the following sections, we review the hardware and software environment used to run experiments. Second, the distribution and characteristics of our adopted workloads are presented. Third, the detailed evaluation results are presented. Finally, a brief analysis of the characteristics of our proposed LPW algorithm is provided.

---

14) https://github.com/Intel-bigdata/HiBench.

Table 4 Hardware and software environment

Hardware/Software configuration	Detail
CPU Model	16 * Intel(R) Xeon(R) Bronze 3106 CPU
CPU FREQUENCY	1.70 GHz
RAM	64 GB
DISK	4 TB (1.9 TB for HDFS)
OPERATING SYSTEM	CentOS Linux release 7.3.1611 (Core)
SPARK VERSION	2.2.3
SCALA VERSION	2.11.12
JDK VERSION	1.8.0_131
HADOOP VERSION	2.7.7

6.1 Environment

Our experiments are running on a Spark cluster with four machines. Each machine has 64 GB of memory and an Intel(R) Xeon(R) Bronze 3106 processor with 16 cores. The CPU frequency is 1.70 GHz. The disk size is 4 TB, among which 1.9 TB are configured as HDFS. The installed operating system is CentOS. The version of the installed Spark is 2.2.3, on which our LPW algorithm is implemented. In our experiments, the Spark's built-in LRU algorithm and the latest LRC algorithm are used as a baseline to compare with our proposed LPW. In our Spark cluster, one machine is set as the master node, and the left three machines are set as slave nodes. During experiments, the tested applications are submitted to the cluster. The applications are run on the Standalone deploy mode. The hardware and software configuration details of our experiment environment are also shown in Table 4.

6.2 Workloads

To observe system performance with different cache replacement policies, our experiments are conducted with benchmarks having different characteristic workloads. In this subsection, the popular Hibench benchmark is chosen and system performance is evaluated. Hibench is an open-source big data benchmark tool including various application categories, such as micro, ml (Machine Learning), SQL, GraphX, and Websearch. It can be used to evaluate the speed, throughput, and system resource utilization of different big data frameworks. The sparkbench framework is chosen as our experimental framework. In total, there are 19 workloads in HiBench (V7.0). Specifically, we use two rules to select these 7 workloads for our evaluation. First, we select the workloads that are highly relied on the cache mechanism. We use the command grep -rn '.cache(' and grep -rn '.persist(' to search each workload which call cache or persist API. Second, for each algorithm, we also keep one. For example, /ml/DenseKmeans.scala and /dal/Densekmeans.scala both implement the KMeans algorithm. We only choose the first one in our experiments. To differentiate from these applications, seven typical applications are finally obtained from HiBench Benchmark as our experimental objects.

HiBench supports the generation of input data with different scales, such as tiny, huge, large, and gigantic-scale. To simulate big data scenarios as much as possible, larger and above data scale as input is selected. We calculate the input data size for each workload, as shown in the fourth column of Table 5. More details about Hibench are elaborated in Table 5. For the seven workloads (C1–C7), we find cache or persist API is called. The column 3 in Table 5 (category) shows the application category. It can be seen that cache or persist API is widely used in iterative scenarios, such as Machine Learning and GraphX. The last column shows some other important input parameter configurations required for applications to run.

6.3 Evaluation results

For each workload, these workloads in cluster configured with different executor memory and driver memory are deployed. Applications are run many times in the Standalone mode, and then use the average execution time. The system performance of each workload is observed and tested under different cache replacement strategies. The other configuration parameters are the default value, such as MemoryFraction, which is set to 0.6. Execution time, speed up and eviction frequency under different cache replacement strategies are compared.

Table 5 Summary of workloads for Hibench

| ID	Category	Workload	Input data size (GB)	Parameter configurations
C1	Websearch	Pagerank	2.79	pages = 5000000; num_iterations = 3
C2	GraphX	Nweight	0.29	edges = 10000000; degree = 3; max_out_edges = 30
C3	Machine Learning	Kmeans	37.41	dimensions = 20; num_of_samples = 200000000
C4	Machine Learning	LogisticRegression	37.25	examples=10000; features = 100000
C5	Machine Learning	PCA	0.22	examples = 1000; features = 1000
C6	Machine Learning	SVM	107.3	examples = 120000; features = 300000
C7	Machine Learning	Bayes	70.14	pages = 20000000; classed = 20000; ngrams = 2

Figure 11 (Color online) Execution time for different workloads with 15 GB of executor memory and driver memory.

Application execution time. A set of experimental results under the configuration of 15 GB executor and driver memory are summarized. First, the execution time of each application under different cache replacement policies, i.e., LRU, LRC, and LPW is observed. The overview result is shown in Figure 11.

As we all know, the configuration of cluster parameters has a great impact on the performance of the application. The execution time of workloads against different spark.memoryFraction (from 0.1 to 0.9) is measured, and results are depicted in Figure 12. To answer RQ1, LPW effectively addresses the performance issues faced by applications, especially in the limited resource. Compared with LRC, application Nweight have poor performance especially in smaller spark.memoryFraction, such as 0.1 and 0.3 as shown in Figure 12(a). This is because the nweight has only 1 job. When the only one job is executed, the advantage of reusing cached data is not obvious. However, when the LogisticRegression application with 44 jobs and Bayes with 9 jobs are executed, the cached data is relied on by many jobs. The performance has been greatly improved.

SpeedUp. Next, speedup and average execution time of all workloads having cache or persist API in HiBench test suites in-depth were observed. The experimental results are shown in Table 6. It can be seen that Kmeans, LogisticRegression, and Bayes show good performance under the configuration of 15 GB executor memory and driver memory. The other parameter configuration of the experimental results is default. This technique saves more than 45.80% time over LRU. The performance improvement of Bayes is the most noticeable, which can be up to 74.87% higher than LRU. Caching reasonable data reduces the execution time for some typical applications. It should be noted that some applications in HiBench have good improvement because of the memory-intense feature. It is common for iterative scenarios such as Kmeans and LogisticRegression to cache input data in memory before model training. These distributed applications repeatedly perform the same calculation on different data. For answering RQ2, LPW can
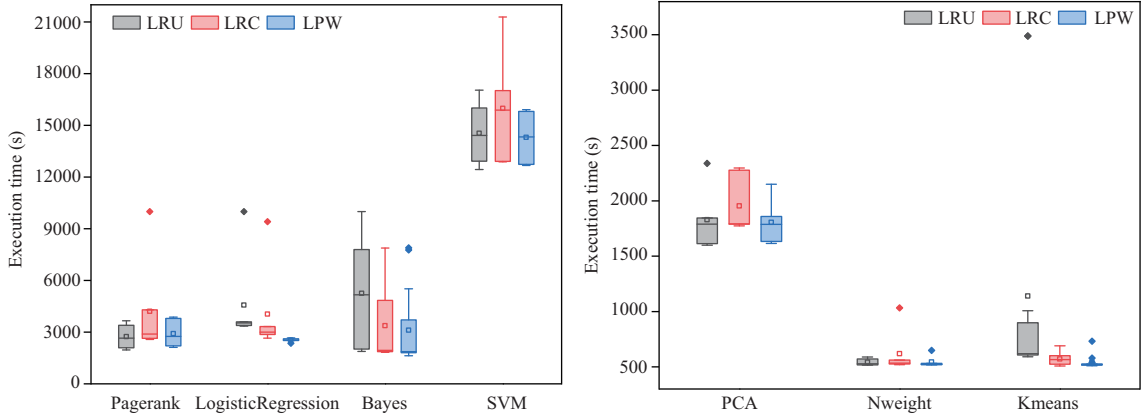
Figure 12 (Color online) Execution time under the LPW, LRU and LRC. (a) Nweight; (b) LogisticRegression; (c) Bayes.

Table 6 Execution time and speedup with different cache replacement strategies under 15 GB executor memory

Workload	LRU (s)	LRC (s)	LPW (s)	SpeedUp-LRC (%)	SpeedUp-LRU (%)
Pagerank	7921.324	5760.366	2047.133	27.28	74.16
Nweight	541.5135	618.2463	543.3262	−14.17	−0.33
Kmeans	1385.608	570.9917	541.4555	58.79	60.92
LogisticRegression	4677.865	4044.374	2535.377	13.54	45.80
PCA	1829.108	1953.538	1805.344	−6.80	1.30
SVM	14540.98	15981.48	14299.84	−9.91	1.66
Bayes	2154.762	570.9917	541.4555	73.50	74.87

Figure 13 (Color online) Execution time for workloads with different cache replacement policies.

effectively speed up the execution of Spark applications compared to LRU and LRC.

To verify the effectiveness of our cache replacement algorithm in-depth, we focus on selecting workload having cache API. The seven workloads of Hibench are chosen. Figure 13 sketches out the results. To answer RQ3, LPW is lightweight without bringing extra overhead for the performance of Spark applications without cache or persist API. For these seven different characteristics of workloads, we can see that when compared with LRU, these applications have better performance. Especially, in the case of LogisticRegression (LR) and Bayes, they have excellent performance. To answer RQ2, LPW can speed up the execution of Spark applications compared to the LRU and the other state of solutions such as LRC.

Eviction frequency. The efficiency of LPW can also be illustrated from evicting frequency. Less eviction frequency for the same application under the same configuration exists. For example, the number of evictions for Bayes workload under different cache replacement policies is further summarized. As shown in Table 7, during the entire application execution process, the total number of replacements that occur decreases. LRU has 3958 evictions and LRC has 3718 evictions, while LPW had 3685 replacements. Additionally, the frequency of replacement on each node is also reduced. For Node_1, the number of evictions is 1376 under LRU, while the replacement number under LRU is 1185, indicating that our method LPW can effectively predict hot spot data to cache without causing frequent replacement.

Compared with LRU and LRC, LRW presents good advantages in workloads such as graph computing

Table 7 Eviction numbers of Bayes with different cache replacement policies

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

```
1.        ......
2.  val pairs=inputRDD.cache()
3.  val rawFiltered=pairs.filter(item=> item._1.startsWith("i")).cache()
4.  val fineFiltered = rawFiltered.filter(item =>
5.     item._1.startsWith("id")).cache()
6.  val fineResults = fineFiltered.reduceByKey(_+_)
7.  val rawResults=rawFiltered.reduceByKey(_+_)
8.  val counts=rawResults.union(fineResults)
9.  println("counts===="+counts.first().toString());
10. pairs.unpersist()
```
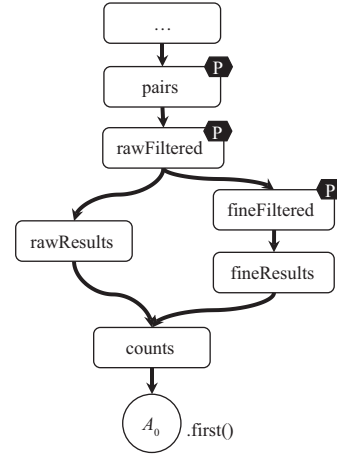
(a)   (b)

Figure 14 (Color online) Code segment (a) and corresponding lineage graph (b).

and machine learning. The main reason is that these applications call the cache or persist API, when the optimal partitions are selected for replacement.

6.4 Analysis and discussion

Compared with existing cache replacement algorithms, the extra overhead introduced by our proposed LPW algorithm is quite small. Only the newly introduced factors are collected and maintained. The calculation time for these factors is negligible. For example, when a job is executed, the time to compute a partition and the space needed to store that partition can be gotten easily. For the reference count factor, some network overhead is needed, since the reference count of a RDD needs to be synced from the master to the slaves. However, this overhead is also very small since this reference count information only has several bytes of data. These factors for different partitions are maintained in an in-memory table. The formula to calculate weight of different partitions is simple, and the consumed time is negligible.

We select typical and popular open-source Spark applications in Github. Experiments show that LPW is promising in identifying hot data for Spark applications. A challenge in efficient cache replacement is to select important factors affecting performance. Our studied factors, also known as, computation cost, reference count, and block size, are representative and important. The arbitrariness of cache API usage, the diversity of application characteristics, and the variability of memory resources construct challenges for efficient system performance. LPW can dynamically be aware of hot data and provide a wise decision to improve performance for various workloads.

An independent evaluation to evaluate LPW is provided as shown in Figure 14. There are three RDDs that would be cached. Given the same input data and the same workload, parameters such as executor memory are configured the same. The workload is run in local [9] mode and the memory with $-X_{mx} = 512$ MB. We then observe the execution time and detailed replacement process under the LPW and LRU. In our experiments, the execution time of the example is 79 and 96 s under LPW and LRU, respectively.

Logs are collected and analyzed under different cache replacement strategies. Our results are shown in Figure 15. Under the guidance of the LRU strategy, the cache blocks belonging to paris (RDD_3) and rawFiltered (RDD_4) are preferentially eliminated. On the contrary, the decision made by LPW is to give priority to removing cache blocks belonging to fineFiltered (RDD_5). In calculation, when compared with RDD_5, RDD_3 and RDD_4 are accessed first. According to recency in LRU, blocks of RDD_3 and

Figure 15 (Color online) Cache replacement process under LRU and LPW. (a) Replacement under LRU; (b) replacement process under LPW.

RDD_4 are firstly evicted as shown in Figure 15(a). According to LPW, blocks in RDD5 with the smallest weight are first eliminated under the LPW as shown in Figure 15(b). When the application is executed, evicting RDD_3 or RDD_4 would not be the best choice. Because of the calculation time and other factors, the weight value of these two RDDs calculated under LPW would be greater than the weight of RDD_5. The optimal choice would be to evict partitions of RDD_5. Our experiments show that LPW can guide Spark to keep the reasonable data to be reused and evict other data for replacement to enhance execution efficiency.

7 Related work

Cache replacement strategies. Remarkable advancements in identifying factors affecting caching are ongoing. Many cache replacement strategies, including traditional strategies, such as LFU, FIFO, and the newer work LRC [8] can be found. GD-Wheel [10], a cost-aware replacement policy can reduce the total recomputation cost and average latency for web applications. An imitation learning approach PARROT [11] has been proposed, which automatically learns data access patterns by leveraging Belady and computing the optimal eviction decision. The LRFU [12] algorithm takes the computational cost and size of the data block into account but does not consider the number of uses of the data block. The AWRP [13] algorithm captures the frequency of data usage, but it misses the size of the data block. LERC [14] stores the relevant blocks of the computing task as a whole and keep them in the memory thus, omitting the data dependency. The LRC algorithm puts the concept of the number of references forward. The system monitors the number of references of each data block in the memory of each slave node. The principle of cache replacement is to remove the data block from the memory with the lowest number of references. The performance of LRC is better than LRU because it considers the problem of

data dependency when the program is running. LRC evicts the data that will not be used in subsequent calculations in time and retains the data that will be used multiple times in the future not considering the data calculation. For the Spark framework, these algorithms ignore some important factors such as the data dependency or the computational cost of data. It is often difficult to have better performance in various scenarios just based on frequency or recency.

Memory optimization solutions. In the data parallel computing framework, caching data is one of the key technologies for achieving performance improvement. Dache [15] proposed caching intermediate result to reuse in MapReduce framework. For the Spark framework, Yang et al. [16] showed that the technology backs up intermediate data in different storage levels for re-use without recalculation to achieve higher execution efficiency. Many studies have elaborated performance issues caused by persisted data in memory during data processing. As developers understand the difference in professional knowledge such as caching mechanism, it brings uncertainty in caching API usage, which further causes performance problems for the application. In real-world cases, Spark's typical and representative built-in libraries such as GraphX [17] and MLlib [18] currently generally have improper use of cache API codes that cause application performance problems. Memory cache optimization aims to improve the utilization of data cached in memory during data processing and optimize application execution. Xu et al. [19] proposed one of the most common reasons for an executor out of memory (OOM) is that developers always want to cache or load too much information into memory. Ousterhout et al. [20] pointed out that the increase in bandwidth and the use of SSDs or disk arrays can effectively solve network or disk IO problems. Memory-based computing frameworks such as Spark have performance bottlenecks mainly in CPU and memory rather than network overhead or Disk IO and other issues. It is of great significance to improve memory utilization in big data parallel data processing systems. The MemTune [21] considers the partial data dependency of the currently running task. In the multi-tenancy scenario, MemTune traverses all submitted DAG graphs after each computing task is completed, for which the complexity is relatively large. Stark [22] optimized memory calculations on dynamic data set collections and enforced data locality to avoid unnecessary data duplication and shuffling, ignoring whether the data is needed for subsequent calculations. PACMan [23] designed two cache replacement strategies to coordinate access to distributed caches. The first one cleared a large amount of incomplete inputs to minimize the average completion time. Another one is to clear infrequently accessed inputs to maximize cluster efficiency. Data access dependencies about applications are not considered. Ref. [24] was conducted on the analysis of the temporal and spatial distribution of information such as data references in the entire workflow. Neutrino [25] captured the data access dependencies between RDDs in different stages of the Spark application, and adaptively converted or discarded RDDs from different cache storage levels. If the cost of data calculation is added, it is more meaningful to make appropriate caching decisions for the entire cluster.

8 Conclusion

Spark is a widely used distributed computing framework for processing big data. Intermediate result caching is an important mechanism for Spark to achieve high performance. Due to the limitation of storage memory, the cached data often need to be replaced. As a result, cache replacement strategies can largely affect the performance of Spark. However, the LRU cache replacement algorithm, currently adopted by Spark, only considers the least recently accessed and simply evicts the least recently used data. This mechanism presents limitations for many applications. To overcome the limitations of LRU and further improve the performance of Spark, we designed a data-aware cache replacement algorithm, the LPW. In view of the diversity of application characteristics and the continuous variability of memory usage requirements, the LPW establishes a unified weight model based on factors and calculates the tracked statistics for the block to identify hot data to be swapped into the memory. The weight is used as the basis for a replacement decision. Our comprehensive experiments show the effectiveness of LPW especially for iterative applications. LPW outperformed related algorithms, such as LRU and LRC, with respect to the execution time and eviction frequency.

References

1 Shvachko K, Kuang H, Radia S, et al. The Hadoop distributed file system. In: Proceedings of Symposium on Mass Storage Systems and Technologies, 2010. 1–10

2 Li H Y, Ghodsi A, Zaharia M, et al. Tachyon: reliable, memory speed storage for cluster computing frameworks. In: Proceedings of the ACM Symposium on Cloud Computing, Seattle, 2014. 1–15

3 Saha B, Shah H, Seth S, et al. Apache Tez: a unifying framework for modeling and building data processing applications. In: Proceedings of International Conference on Management of Data, Melbourne, 2015. 1357–1369

4 Zaharia M, Chowdhury M, Das T, et al. Fast and interactive analytics over Hadoop data with spark. Adv Comput Syst Assoc, 2012, 37: 45–51

5 Zaharia M, Chowdhury M, Das T, et al. Resilient distributed datasets: a fault-tolerant abstraction for in-memory cluster computing. In: Proceedings of USENIX Conference on Networked Systems Design and Implementation, 2012

6 Li H, Wang D, Huang T Z, et al. Detecting cache-related bugs in Spark applications. In: Proceedings of International Symposium on Software Testing and Analysis, Virtual Event, 2020. 363–375

7 Geng Y Z, Shi X H, Pei C, et al. LCS: an efficient data eviction strategy for Spark. Int J Parallel Prog, 2017, 45: 1285–1297

8 Yu Y H, Wang W, Zhang J, et al. LRC: dependency-aware cache management for data analytics clusters. In: Proceedings of Conference on Computer Communications, Atlanta, 2017. 1–9

9 Huang S S, Huang J, Dai J Q, et al. The Hibench benchmark suite: characterization of the MapReduce-based data analysis. In: Proceedings of International Conference on Data Engineering Workshop, Long Beach, 2010. 41–51

10 Li C, Cox A L. GD-Wheel: a cost-aware replacement policy for key-value stores. In: Proceedings of the European Conference on Computer Systems, Bordeaux, 2015. 1–15

11 Liu E, Hashemi M, Swersky K, et al. An imitation learning approach for cache replacement. In: Proceedings of the International Conference on Machine Learning, 2020. 6237–6247

12 Lee D H, Choi J, Kim J H, et al. On the existence of a spectrum of policies that subsumes the least recently used (LRU) and least frequently used (LFU) policies. SIGMETRICS Perform Eval Rev, 1999, 27: 134–143

13 Swain D, Paikaray B, Swain D. AWRP: adaptive weight ranking policy for improving cache performance. Comput Sci, 2011, 3: 2151–9617

14 Yu Y H, Wang W, Zhang J, et al. LERC: coordinated cache management for data-parallel systems. In: Proceedings of Global Communications Conference, 2017. 1–6

15 Zhao Y X, Wu J. Dache: a data aware caching for big-data applications using the MapReduce framework. In: Proceedings of International Conference on Computer Communications, 2013. 35–39

16 Yang Z Y, Jia D L, Ioannidis S, et al. Intermediate data caching optimization for multi-stage and parallel big data frameworks. In: Proceedings of International Conference on Cloud Computing, San Francisco, 2018. 277–284

17 Gonzalez J E, Xin R S, Dave A, et al. GraphX: graph processing in a distributed dataflow framework. In: Proceedings of USENIX Symposium on Operating Systems Design and Implementation, Broomfield, 2014. 599–613

18 Meng X R, Bradley J, Yavuz B, et al. MLlib: machine learning in Apache Spark. J Mach Learn Res, 2016, 17: 1235–1241

19 Xu L J, Dou W S, Zhu F, et al. Characterizing and diagnosing out of memory errors in MapReduce applications. J Syst Softw, 2018, 137: 399–414

20 Ousterhout K, Rasti R, Ratnasamy S, et al. Making sense of performance in data analytics frameworks. In: Proceedings of USENIX Symposium on Networked Systems Design and implementation, Oakland, 2015. 293–307

21 Xu L, Li M, Zhang L, et al. MemTune: dynamic memory management for in-memory data analytic platforms. In: Proceedings of International Parallel and Distributed Processing Symposium, Chicago, 2016. 383–392

22 Li S, Amin M T, Ganti R, et al. Stark: optimizing in-memory computing for dynamic dataset collections. In: Proceedings of International Conference on Distributed Computing System, Atlanta, 2017. 103–114

23 Ananthanarayanan G, Ghodsi A, Warfield A, et al. PACMan: coordinated memory caching for parallel jobs. In: Proceedings of Symposium on Networked Systems Design and Implementation, San Jose, 2012. 267–280

24 Perez T B, Zhou X B, Chen D Z. Reference-distance eviction and prefetching for cache management in Spark. In: Proceedings of International Conference on Parallel Processing, Eugene, 2018. 1–10

25 Xu E, Saxena M, Chiu L. Neutrino: revisiting memory caching for iterative data analytics. In: Proceedings of USENIX Workshop on Hot Topics in Storage and File Systems, Denver, 2016. 16–20