

Efficient B-tree Based Indexing for Cloud Data Processing

Sai Wu ^{#1}, Dawei Jiang ^{#2}, Beng Chin Ooi ^{#3}, * Kun-Lung Wu ^{§4}

[#]*School of Computing, National University of Singapore, Singapore*

^{1,2,3}{wusai, jiangdw, ooibc}@comp.nus.edu.sg

[§]*IBM T. J. Watson Research Center*

⁴klwu@us.ibm.com

ABSTRACT

A Cloud may be seen as a type of flexible computing infrastructure consisting of many compute nodes, where resizable computing capacities can be provided to different customers. To fully harness the power of the Cloud, efficient data management is needed to handle huge volumes of data and support a large number of concurrent end users. To achieve that, a scalable and high-throughput indexing scheme is generally required. Such an indexing scheme must not only incur a low maintenance cost but also support parallel search to improve scalability. In this paper, we present a novel, scalable B⁺-tree based indexing scheme for efficient data processing in the Cloud. Our approach can be summarized as follows. First, we build a local B⁺-tree index for each compute node which only indexes data residing on the node. Second, we organize the compute nodes as a structured overlay and publish a portion of the local B⁺-tree nodes to the overlay for efficient query processing. Finally, we propose an adaptive algorithm to select the published B⁺-tree nodes according to query patterns. We conduct extensive experiments on Amazon's EC2, and the results demonstrate that our indexing scheme is dynamic, efficient and scalable.

1. INTRODUCTION

There has been an increasing interest in deploying a storage system on Cloud to support applications that require massive scalability and high throughput in storage layer. Examples of such systems include Amazon's Dynamo [15] and Google's BigTable [13]. Cloud storage systems are designed to meet several essential requirements of data-intensive applications: manageability, scalability, availability, and low latency. Computer nodes that are allocated from Cloud are maintained as a resource pool and can be dynamically added/removed from the pool as resource demands change over time. Datasets are automatically partitioned and replicated among available nodes for scalability and availability. Query efficiency is achieved by either employing a pure key-value data model, where both key and value are arbitrary byte strings (e.g. Dy-

namo), or its variant, where key is an arbitrary byte string and value is a structured record consisting of a number of named columns (e.g. BigTable, which supports efficient retrieval of values via a given key or key range).

However, existing solutions lack of built-in support for secondary index, a useful feature for many applications. In real world, users tend to query data with more than one keys. For example, in an online video system, such as Youtube, each video could be stored in a key-value store with a unique video id as the key and video information, including title, upload time and number of views as the value. Although the video can be efficiently retrieved via video id, a common scenario is that the end user wants to find videos with given titles or within a date range. Current practice to solve this problem is to run a MapReduce job that scans the whole dataset and produces the necessary second indices in an offline batch manner. Problems of this approach is that the secondary index is not up-to-date and newly inserted tuples cannot be queried until they are indexed. For instance, when a new item is inserted into Google Base, that item could be delayed for one day to be seen by users.

This paper presents CG-index (Cloud Global index), a secondary indexing scheme for Cloud storage systems. CG-index is designed for Cloud platform and built from scratch. It is tailored for online queries and maintained in an incremental way. It shares many implementation strategies with shared-nothing databases [16], peer-to-peer computing [14, 20], and existing Cloud storage systems [18, 15]. CG-index supports usual dictionary operations (*insert*, *delete* and *lookup*), as well as range search with a given key range.

CG-index software consists of two components: a client library which is linked with user application and a set of index servers which store the index. The CG-index servers operate in a shared pool of compute nodes allocated from Cloud and the index server process can reside in the same physical machine with the storage server process. Besides high scalability and availability, CG-index can be easily integrated into various storage systems to support high throughput and high concurrency. These features are achieved by adopting three techniques: 1) a generic key-pointer representation, 2) partition aware indexing, and 3) eventual consistency.

CG-index stores each index entry as an s_k -handle pair, where s_k is the secondary key that will be indexed and handle is an arbitrary byte string which could be used to fetch the corresponding value in the Cloud storage system. Throughout this paper, the term *primary key* is referred to the key stored in key-value store and the term *secondary key* is referred to the key stored in CG-index. This design facilitates the integration of CG-index with various storage systems. CG-index treats a handle as an uninterpreted string. Users can serialize arbitrary information into a handle. For example, users can directly store the primary keys in handles or serialize the primary keys along with timestamps into handles. The latter case is a typi-

*The work of the three NUS authors was in part funded by Singapore Ministry of Education Grant No. R-252-000-394-112 under the project name of Utab.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Articles from this volume were presented at The 36th International Conference on Very Large Data Bases, September 13-17, 2010, Singapore.

Proceedings of the VLDB Endowment, Vol. 3, No. 1

Copyright 2010 VLDB Endowment 2150-8097/10/09... \$ 10.00.

cal usage of indexing data in BigTable since each value in BigTable is timestamped.

All existing storage systems employ some horizontal partitioning scheme to store large datasets in a cluster. The idea is to partition the dataset into a number of small pieces called data shards. Each data shard is a distribution unit and is stored on a unique computer node. CG-index is designed to be aware of and optimized for such form of partitioning. Instead of building an index for the whole dataset, CG-index builds a local B^+ -tree index for each data shard called an index shard. The index shard is a distribution unit in CG-index, which is stored and maintained on a unique index server. CG-index relies on this index distribution technique for desired scalability. Queries are served by searching all qualified index shards. The returned results is a stream of s_k -handle pairs. We can group the handles by their data shard IDs. An optimization is to retrieve a group of handles from the same data shard in a batch mode.

The index server is responsible for serving requests that involve data shards indexed by the server. To route queries among the servers, all index servers are organized as a structured peer-to-peer network, BATON [20]. Each index server maintains connections to its neighbors in the network. It collects some B^+ -tree nodes of its neighbors and thus knows data indexed by other servers. A query routing algorithm traverses the network with neighbor links and returns all s_k -handle pairs. Since each index server only sends a portion of its local B^+ -tree nodes to neighbors, only the updates involving the published B^+ -tree nodes trigger the synchronization process. Therefore, in most cases, index servers update index entries locally, achieving high throughput and concurrency. Finally, to obtain required availability and resilience to network partition, we replicate the data of an index server to multiple servers. Eventual consistency is adopted to maintain consistency between replicas.

The above techniques complete the design of CG-index. Although some implementation techniques applied by CG-index have been studied in the literature, adapting them in a Cloud system makes CG-index design unique and challenging. The rest of the paper is organized as follows: Section 2 reviews related work. Section 3 outlines our system architecture. Section 4 and Section 5 present the proposed indexing and tuning algorithms. Section 6 empirically validates the effectiveness and efficiency of our proposed indexing scheme. We conclude in Section 7. The algorithms and theorems are listed in the appendix (Section A). We also discuss some other optimizations in the appendix.

2. RELATED WORK

Building scalable data storage systems is the first step towards Cloud computing. These storage systems are always tailored for specific workload. The most important and fundamental storage system is the distributed file system. Google's GFS [18] and its open source implementation, HDFS [3], are designed to support large-scale data analytic jobs, where datasets are split into equal-size chunks. The chunks are randomly distributed over the computer nodes. Amazon's Simple Storage Service (S3) [2] is a data storage service that allows users to store and retrieve objects on Amazon's Cloud infrastructure. S3 can be used to support high-frequency access over the internet. OceanStore [21], Farsite [9] and Ceph [26] provide peta-bytes of highly reliable storage. They can support thousands of online users and simultaneous accesses.

Based on above systems, some more sophisticated systems are proposed to support various applications. Most of them are key-value based storage systems. Users can efficiently retrieve the data via the primary key. BigTable [13] is a distributed storage system for managing large-scale structured datasets. BigTable maintains

its SSTable File in GFS and combines the techniques of row-based and column-based databases. To reduce the overheads, it applies the eventually consistent model, which is also adopted in this paper. HyperTable [4] is an open source implementation of BigTable on HDFS. Amazon's Dynamo [15] is a key-value store for many Amazon's services. It applies the consistent hash function to distribute data among the computer nodes. Similar to Dynamo, Voldemort [6] is a distributed key-value store, which can be scaled to a large cluster and provide high throughput. Cassandra [5], Facebook's distributed key-value store, combines the model of BigTable and Dynamo to support efficient InBox search. Although the underlying implementation may be different, the common goal of these proposals is to provide techniques that store huge datasets over a shared-nothing computer cluster or data center. These work are orthogonal to ours. We focus on providing an efficient secondary indexing technique over those data storage systems.

In [10], a distributed B^+ -tree algorithm was proposed for indexing the large-scale dataset in the cluster. The B^+ -tree is distributed among the available nodes by randomly disseminating each B^+ -tree node to a compute node (also called server node in [10]). This strategy has two weaknesses. First, although it uses a B^+ -tree based index, the index is mainly designed for simple lookup queries and is therefore not capable of handling range queries efficiently. To process a range query $[l, u]$, it must first locate the leaf node responsible for l . Then, if u is not contained by the leaf node, it needs to retrieve the next leaf node from some compute server based on the sibling pointer. Such form of retrieval continues until the whole range has been searched. Second, it incurs high maintenance cost for the server nodes and huge memory overhead in the client machines, as the client node (user's own PC) lazily replicates all the corresponding internal nodes.

The work that is most related to ours is RT-CAN [25]. RT-CAN integrates CAN [23]-based routing protocol and the R-tree based indexing scheme to support multi-dimensional queries. Different from RT-CAN, CG-index organizes computer nodes into a BATON [20] network and builds B-tree indexes to support high throughput one-dimensional queries. CG-index is a preliminary work of our project, epiC [7]. In epiC, we re-implement RT-CAN and CG-index in a unified indexing framework to support various types of queries.

3. SYSTEM OVERVIEW

Figure 1 shows the system architecture of our cluster system. A set of low-cost workstations join the cluster as compute (or processing) nodes. This is a shared nothing and stable system where each node has its own memory and hard disk. To facilitate search, nodes are connected based on the BATON protocol [20]. Namely, if two nodes are routing neighbors in BATON, we will keep a TCP/IP connection between them. Note that BATON was proposed for a dynamic Peer-to-Peer network. It is designed for handling dynamic and frequent node departure and joining. Cloud computing is different in that nodes are organized by the service provider to enhance performance. In this paper, the overlay protocols are only used for routing purposes. Amazon's Dynamo [15] adopts the same idea by applying consistent hashing for routing in clusters. BATON is used as the basis to demonstrate our ideas due to its tree topology. Details of BATON can be found in the appendix. Other overlays supporting range queries, such as P-Ring [14] and P-Grid [8], can be easily adapted as well.

In our system, data are partitioned into data shards (based on the primary key), which are randomly distributed to compute nodes. To facilitate search for a secondary key, each compute node builds a B^+ -tree for the key to index its local data (data shards assigned to the node). In this way, given a key value, we can efficiently receive

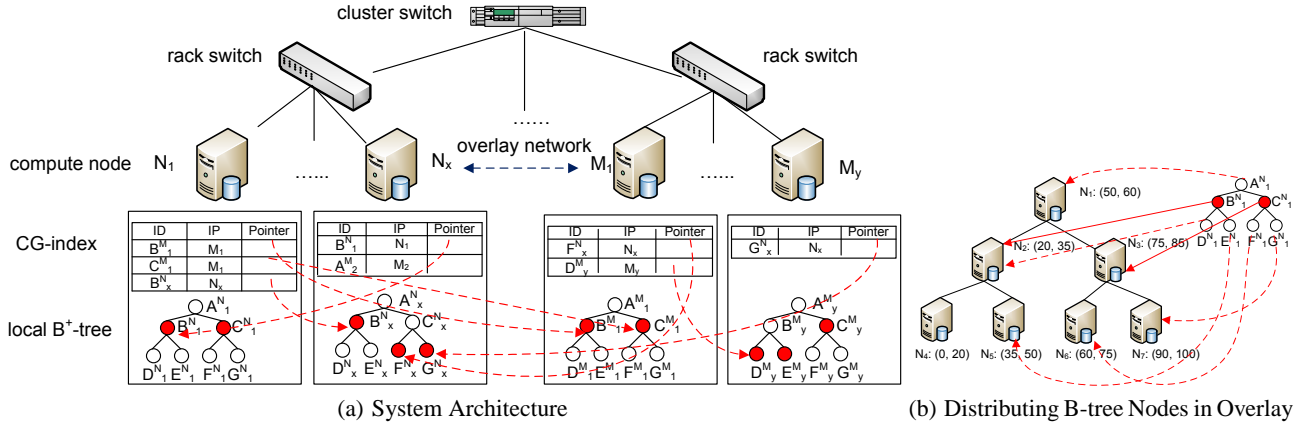


Figure 1: System Overview

its handle. The handle is an arbitrary byte string which could be used to fetch the corresponding value in the Cloud storage system. To process queries in the cluster, a traditional scheme will broadcast the queries to all the nodes, where local search is performed in parallel. This strategy, though simple, is not cost efficient and scalable. Another approach is to maintain data partitioning information in a centralized server. The query processor needs to look up the partitioning information for every query. The central server risks being the bottle-neck.

Therefore, given a key value or range, to locate the corresponding B^+ -trees, we build a global index (CG-index) over the local B^+ -trees. Specifically, some of the local B^+ -tree nodes (red nodes in Figure 1) are published and indexed in the remote compute nodes based on the overlay routing protocols. Note that to save the storage cost, we only store the following meta-data of a published B^+ -tree node: $(blk, range, keys, ip)$, where blk is the disk block number of the node, $range$ is the value range of the B^+ -tree node (we will discuss it in the next section), $keys$ are search keys in the B^+ -tree node and ip is the IP address of the corresponding compute node. In this way, we maintain a remote index for the local B^+ -trees in each compute node. These indexes compose the CG-index in our system. Figure 1(a) shows an example of the CG-index, where each compute node maintains a portion of the CG-index. Figure 1(b) gives an example of mapping B^+ -tree nodes to compute nodes in the overlay. To process a query, we first look up the CG-index for the corresponding B^+ -tree nodes based on the overlay routing protocols. And then following the pointers of the CG-index, we search the local B^+ -trees in parallel.

The CG-index is disseminated to compute nodes in the system. To improve the search efficiency, the CG-index is fully buffered in memory, where each compute node maintains a subset of CG-index in its memory. As memory size is limited, only a portion of B^+ -tree nodes can be inserted into CG-index. Hence, we need to plan our indexing strategy wisely. In this system, we build a virtual expansion tree for the B^+ -tree. We expand the B^+ -tree from the root node step by step. If the child nodes are beneficial for the query processing, we will expand the tree and publish the child nodes. Otherwise, we may collapse the tree to reduce maintenance cost and free up memory. Algorithm 1 shows the general idea of our indexing scheme. Initially, the compute node only publishes the root of its local B^+ -tree. Then, based on the query patterns and our cost model, we compute the benefit of expanding or collapsing the tree (line 4 and line 7). To reduce maintenance cost, we only publish internal B^+ -tree nodes (we will not expand the tree to the leaf level). Note that in our expanding/collapsing strategy, if a B^+ -tree node is

indexed, its ascendant and descendant nodes will not be indexed. Overlay routing protocol allows us to jump to any indexed B^+ -tree nodes directly. Therefore, we do not need to start the search from the B^+ -tree's root.

Algorithm 1 CGIndexPublish(N_i)

- 1: N_i publishes the root node of its B^+ -tree
 - 2: **while** true **do**
 - 3: N_i checks its published B^+ -tree node n_j
 - 4: **if** isBeneficial($n_j.children$) **then**
 - 5: expand the tree from n_j by indexing n_j 's children
 - 6: **else**
 - 7: **if** benefit(n_j) < maintenanceCost(n_j) **then**
 - 8: collapse the tree by removing n_j and index n_j 's parent
 if necessary
 - 9: wait for a time
-

4. THE CG-INDEX

Different from [10], where a global B^+ -tree index is established for all the compute nodes in the network, in our approach, each compute node has its local B^+ -tree, and we disseminate the local B^+ -tree nodes to various compute nodes. In this section, we discuss our index routing and maintenance protocols. The index selection scheme will be presented in next section. To clarify the representation, we use **upper-case** and **lower-case** characters to denote the compute node and B^+ -tree node, respectively.

4.1 Indexing Local B^+ -tree Nodes Remotely

Given a range, we can locate the BATON node responsible for the range (the node whose subtree range can fully cover the search range). On the other hand, the B^+ -tree node maintains the information about the data within a range. This observation provides us with a straightforward method to publish B^+ -tree nodes to remote compute nodes. We employ the *lookup* protocol in the overlay to map a B^+ -tree node to a compute node and store the meta-data of the B^+ -tree node at the compute node's memory.

To publish B^+ -tree nodes into CG-index, we need to generate a range for each B^+ -tree node. Based on their positions, the B^+ -tree nodes can be classified into two types: 1) the node is neither the left-most nor the right-most node at its level and 2) the node and its ancestors are always the left-most or right-most child.

For the first type of nodes, we can generate their ranges based on the parents' information. For example, in Figure 2, node c is node a 's second child. So its range is from the first key to the




Figure 2: B⁺-Tree Nodes and Their Index Ranges

second key of a , namely (12,35). The second type of nodes only provide an open range (no lower bound or no upper bound). We use the smallest value and the largest value in current tree as bound values. To reduce update cost, we slightly increase the range. For example, in Figure 2, we use 0 as the lower bound instead of 5, the actual smallest value. After defining the lower bound and upper bound of the tree, we can generate a range for the type 2 nodes. For example, the ranges of node r and a are (0,100) and (0,45) respectively. The lower bound and upper bound can be cached in memory, and updated when new data are inserted into the left-most or right-most leaf nodes.

To publish a B⁺-tree node, we first generate its range R . Then, based on the BATON routing protocols, we obtain the compute node N , which is responsible for the lower bound of R . Step by step, we forward the request to the ancestors of N until we reach the one whose subtree range can fully contain R . The B⁺-tree node is then indexed in that node. For additional details, please see the appendix (section A).

In the cluster system, as the processing nodes are low-cost workstations, there may be node failures at any time. Single point of failure can be handled by our replication strategy. But when a subset of nodes are offline (e.g. a rack switch is down), all replicas may be lost. To handle such problem, the compute node refreshes all its published B⁺-tree nodes occasionally.

4.2 Query Processing

Given a range query Q , we need to search the CG-index to locate the B⁺-tree nodes whose ranges overlap with Q . We can simulate the overlay’s search algorithm to process the query. Algorithm 2 shows a general range search process. Starting from the lower bound of Q , we follow the right adjacent links to search sibling nodes until reaching the upper bound of Q . However, the range search of many overlays, including BATON’s, could be further optimized. Suppose k nodes overlap with Q . The average cost of a typical range search in BATON is estimated as $\frac{1}{2} \log_2 N + k$, where N is the total number of compute nodes in the system.

Algorithm 2 Search($Q = [l, u]$)

- 1: $N_i = \text{lookup}(l)$
 - 2: perform local search on N_i
 - 3: **while** $N_i = N_i.\text{right}$ and $N_i.\text{low} < u$ **do**
 - 4: perform local search on N_i
-

The first optimization to the range search algorithm is that instead of starting the search from the lower bound, we can start at any point inside the range. Suppose the data are uniformly distributed among nodes and R is the total range, this optimization reduces the average cost of searching a node in a range Q from $\frac{1}{2} \log_2 N + k$ to $\frac{1}{2} \log_2 \frac{QN}{R} + k$.

The existing analysis ignores the effect of k , which in fact dominates search performance in a large-scale overlay network. As a simple example, in a 10,000-node network, suppose the data are uniformly partitioned among processing nodes, $k = 100$ if $\frac{Q}{R} = 0.01$. To reduce the latency of range search, the second optimiza-

tion is to increase parallelism. We broadcast the query to the processing nodes that overlaps with the search range in parallel.

Finally, the new search algorithm is summarized as:

1. Locate a random processing node in the search range (optimization 1).
2. Following the parent link, locate the root node of a BATON subtree. The subtree covers the whole search range.
3. Selectively broadcast the query to the descendants of the subtree (optimization 2).
4. In each processing node, after receiving the search request, do a local search for the CG-index.

Parallel search algorithm reduces the average cost from $\frac{1}{2} \log_2 N + k$ to $\frac{1}{2} \log_2 \frac{QN}{R} + \log_2 N$, where $\log_2 N$ is the height of the BATON tree. For detail algorithms, please refer to the appendix.

5. ADAPTIVE TUNING

In this section, we propose our adaptive indexing strategy based on the cost model of overlay routings. Our adaptive scheme selectively indexes local B⁺-tree nodes according to query patterns by expanding the local B⁺-tree from the root node dynamically.

5.1 Cost Modeling

We now consider the cost of publishing a local B⁺-tree node in the network under the adaptive approach. We do so by reviewing the procedures of query processing and index maintenance. Generally, we consider three types of costs: network routing cost, local search cost and index maintenance cost. All the costs are estimated approximately. We use α and β to denote the average cost of a random I/O operation and the cost of sending an index message, respectively. As we evaluate CG-index on Amazon’s EC2 [1], α and β are estimated based on the results of [24, 12].

In query processing, we first locate the compute nodes responsible for our search key. This incurs $\frac{1}{2} \beta \log_2 N$ cost in the structured overlays, where N is the total number of nodes in the cluster. After locating the compute nodes, we retrieve the indexed B⁺-tree nodes in the CG-index. As the index is fully buffered in memory, the local retrieval cost can be discarded.

Suppose the height of B⁺-tree \mathcal{T} is h and n is \mathcal{T} ’s node with height $h(n)$. Then, processing queries via the index of n will incur $\alpha h(n)$ cost in the local search. We save a cost of $\alpha(h - h(n))$ by searching from n instead of the root node.

On the other hand, to synchronize the local B⁺-tree index with the remote one, we need to send index update messages. The B⁺-tree leaf nodes incur much higher update cost than the internal nodes. Assume that the updates happen uniformly among the leaf nodes. To model the update cost in the local B⁺-tree, we define the parameters in Table 5.1. On average, the nodes of B⁺-tree have $\frac{3m}{2}$ keys. Synchronization is performed when an indexed B⁺-tree node splits or merges with other nodes. Thus, we need to compute the probability of splitting or merging a node n with height $h(n)$.

Table 5.1 Parameters

m	B ⁺ tree’s order
h	height of the node
p_1	probability of insertion
p_2	probability of deletion

This problem can be formalized as a random walk problem with two absorbing states. The start state is at $\frac{3m}{2}$ and the two absorbing states are m and $2m$, respectively. With probability p_1 , we move to

the state $2m$ and with probability p_2 , we move to the state m . The random walk problem can be solved by the theorems in [22], and we obtain the following result:

$$p_{split} = \frac{\left(\frac{p_2}{p_1}\right)^{\frac{3m}{2}} - \left(\frac{p_2}{p_1}\right)^m}{\left(\frac{p_2}{p_1}\right)^{2m} - \left(\frac{p_2}{p_1}\right)^m} \quad (1)$$

$$p_{merge} = \frac{\left(\frac{p_2}{p_1}\right)^{2m} - \left(\frac{p_2}{p_1}\right)^{\frac{3m}{2}}}{\left(\frac{p_2}{p_1}\right)^{2m} - \left(\frac{p_2}{p_1}\right)^m} \quad (2)$$

where p_{split} and p_{merge} are the probabilities of splitting the node and merging the node, respectively. Furthermore, based on [22] we can compute the average updates required for triggering a splitting or merging as:

$$n_u = \frac{m(p_{split} - p_{merge})}{2(p_1 - p_2)} \quad (3)$$

Thus, given the probabilities of updating the child nodes, we can compute the effect to the parent nodes. Iteratively, we can estimate the update probability of nodes at any level. p_{split} and p_{merge} of the child node equal p_1 and p_2 of the parent node, respectively. Finally, suppose there are U updates in a time unit, we can compute the number of updates for each node in the B^+ -tree. To simplify the discussion, we use $g(n_i)$ to represent the number of update messages of a B^+ -tree node n_i (we discard the complex formula of function g for simplifying the presentation). As it takes $\frac{1}{2} \log_2 N$ hops to notify the corresponding compute node, the total cost of maintaining n_i in the remote index is $\frac{1}{2} \beta g(n_i) \log_2 N$. To handle node failure, multiple replicas are kept to improve the availability of CG-index (replication strategy is discussed in Section 5). Suppose there are k replicas for an index entry. The cost of maintaining n_i and its replicas is $\frac{k}{2} \beta g(n_i) \log_2 N$.

Another kind of maintenance cost is the replication cost. As mentioned above, to handle unexpected network failures, a compute node will periodically republish its local B^+ -tree nodes. Suppose republishing happens every T time unit. The cost is estimated as $\frac{\beta \log_2 N}{2T}$. Finally, suppose there are Q queries overlapping with the B^+ -tree node n in a time unit, the total cost of indexing n is:

$$cost(n) = \alpha Q h(n) + \frac{1}{2} \beta (k g(n) + \frac{1}{T}) \log_2 N \quad (4)$$

5.2 Tuning Algorithm

The intuition of our adaptive indexing strategy is to selectively publish the local B^+ -tree nodes based on the query distribution. Initially, only the root node of the B^+ -tree is indexed. However, publishing the root node of the B^+ -tree does not provide efficient search, as its range could be big and this may result in redundant visit of the compute node and local search. To solve this problem, we remove the index of the root node and publish the nodes at the second level (root node's child nodes) when some child nodes are frequently searched over time. The query can then jump directly to the second level of the local B^+ -trees. Similarly, if we find that indexing the nodes is no longer beneficial, we remove the nodes' index and publish their parent node instead. With the same principle, we can recursively expand or shrink the range being indexed, and thereby, increasing or reducing the number of index nodes being indexed. By doing so, we build a dynamic global index based on query distribution and an adaptive expansion of the local B^+ -trees. To reduce maintenance cost, we only publish internal B^+ -tree nodes into CG-index. Consider a local B^+ -tree in Figure 3, the shaded nodes will be indexed in the CG-index based on the query patterns.

Given the cost model, the compute node can estimate the cost of a specific indexing strategy. Specifically, the compute node is




Figure 3: Example of B^+ -tree Indexing Strategy (shaded nodes are published in the CG-index)

responsible for a key range R for routing purposes in the overlay. It stores the index for remote B^+ -tree nodes, whose ranges are covered by R . As a query is routed based on the search range, the compute node must receive any query that overlaps with R . It can have a precise description about the query distribution in R . Hence, the compute node has full information to compute the cost of the current index.

Algorithm 3 and Algorithm 4 generalize our adaptive indexing strategy. In Algorithm 3, we show how to expand the indexed tree. In line 1, we collect the query statistics and generate a histogram to estimate the query patterns. We compare the cost of indexing a B^+ -tree node to the cost of indexing all its child nodes (line 5-7). If indexing the child nodes can improve the search performance, we will remove the index of the parent node and publish the indexes of the child nodes. In this way, we expand the indexed tree. The indexed B^+ -tree node should periodically report its cost status (line 9). Based on the reported status, we can decide whether to collapse the tree. In Algorithm 4, we show the process of collapsing. We group the received reports by their parent nodes (line 1-2). When we receive the reports from all the child nodes, we start to evaluate the cost of different index strategies (line 3-9). If indexing the parent node can reduce the maintenance cost, we replace the indexes of all the child nodes with the index of the parent node (line 6-8). Both Algorithm 3 and Algorithm 4 are invoked occasionally to tune the performance of CG-index.

Algorithm 3 Expand()

- 1: compute the query histogram H
 - 2: **for** $\forall B^+$ -tree node $n_i \in S_n$ **do**
 - 3: $c_1 = n_i$'s current cost
 - 4: $c_2 = n_i$'s child nodes' cost
 - 5: **if** $c_2 < c_1$ **then**
 - 6: remove n_i from S_n
 - 7: notify n_i 's owner to index the child nodes of n_i
 - 8: **else**
 - 9: statusReport(n_i)
-

Algorithm 4 Collapse(B^+ -tree node n_i)

- //receiving a status report from n_i
- 1: $n = n_i.parent$
 - 2: put n_i in n 's child list L_n
 - 3: **if** L_n is full **then**
 - 4: $c_1 = \sum_{\forall n_i \in L_n} cost(n_i)$
 - 5: $c_2 = cost$ of indexing n
 - 6: **if** $c_2 < c_1$ **then**
 - 7: remove index of nodes in L_n
 - 8: notify the owner to index the B^+ -tree node n
 - 9: clear L_n
-

To guarantee the correctness of tuning approach, the expansion and collapse operation are set to be atomic operations. E.g. in expansion operation, if node n_i tries to replace its index entry with the

entries of its children’s, either all the children’s entries are created, or the expansion operation fails and we keep the old entry.

THEOREM 1. *If the expansion and collapse are atomic operations, the adaptive indexing strategy can provide a complete result.*

PROOF. See the appendix. \square

6. MAINTENANCE

6.1 Updating CG-index

In the CG-index, updates are processed concurrently with search. To maximize the throughput and improve the scalability, we adopt the eventual consistent model, which has been adopted in distributed systems [13]. Two types of updates, lazy update and eager update, are supported. When updates of local B^+ -tree do not affect the correctness of search results, we adopt lazy update. Otherwise, eager update is applied to perform search synchronization as soon as possible.

THEOREM 2. *In CG-index, if the update does not affect the key range of a local B^+ -tree, the stale index will not affect the correctness of the query processing.*

PROOF. See the appendix. \square

A close observation reveals that only updates in the left-most or right-most nodes can violate the key range of a local B^+ -tree. Given a B^+ -tree \mathcal{T} , suppose its root node is n_r , and the corresponding range is $[l, u]$. The index strategy of \mathcal{T} is actually a partitioning strategy of $[l, u]$, as 1) each node of \mathcal{T} maintains a sub-range of $[l, u]$ and 2) for any value v in $[l, u]$, there is an indexed node of \mathcal{T} , whose key range covers v . For example, in Figure 3, the root range $[0, 100]$ is partitioned into sub-ranges of $[0, 20]$, $[20, 25]$, $[25, 30]$, $[30, 40]$, $[40, 45]$, $[45, 50]$, $[50, 80]$ and $[80, 100]$. Except left-most and right-most nodes (those nodes responsible for the lower bound and upper bound of the root range), updates in other nodes can only change the way of partitioning. Suppose in Figure 3, node i and j merge together. The sub-range $[20, 25]$ and $[25, 30]$ are replaced by $[20, 30]$. Regardless of how the root range is partitioned, the query can be correctly forwarded to the node based on the index, even if the index is stale. Therefore, if the updates do not change the lower bound or upper bound of the root range, we adopt the lazy update approach. Namely, we do not synchronize the index with the local B^+ -tree immediately. Instead, after a predefined time threshold, all updates are committed together.

Given two nodes n_i and n_j , lazy updates are processed in the following ways.

1. If n_i is merged with n_j and both of them are published into the CG-index, we replace the index entries of n_i and n_j with the index entry of the merged node.
2. If n_i is merged with n_j and only one node (suppose it is n_i) is published into CG-index, we remove all the index entries of n_j ’s child nodes and update n_i ’s index entry as the new merged one.
3. If n_i is published into the CG-index and split into two new nodes, we replace n_i ’s index entry with the index entries of the new nodes.

In the index entry, two attributes, IP address and block number, are used in query processing. Specifically, IP address is used to forward the query to a correct cluster server. And block number is applied to locate the corresponding B^+ -tree node when performing local search. Based on the above analysis, the IP address is always

correct if the updates do not change the lower bound or upper bound of the B^+ -tree. However, the block number may be invalid due to node merging and splitting. In such case, we just start searching from the root node.

On the other hand, some updates in the left-most and right-most nodes may change the lower bound and upper bound of the B^+ -tree. In that case, the old index entry may generate false positive and false negative in query processing. As an example, suppose key “0” is removed from node b in Figure 3, b ’s key range will shrink to $[5, 20]$. If applying the old index to process query $[-5, 3]$, the query will be forwarded to the cluster server, which actually cannot provide any result. That is, the index generates false positives. On the contrary, suppose a new key “-5” is inserted into node b , the key ranges of b and a are updated to $[-5, 20]$ and $[-5, 100]$, respectively. If the old index entry is applied to process query $[-10, -2]$, false negative is generated as the CG-index fails to retrieve the data from some cluster servers. False positive does not violate the consistency of the result and we adopt lazy update strategy to handle it. False negative is crucial for the consistency. Therefore, we apply eager update strategy to synchronize the index.

In eager update, we first update the indexed nodes (including their replicas) in CG-index. If all indexed nodes have been successfully updated, we update the local B^+ -tree nodes. Otherwise, we roll back the operations to keep the old indexed nodes in CG-index and trigger an update failure event.

THEOREM 3. *The eager update can provide a complete result.*

PROOF. See the appendix. \square

6.2 Replication

To guarantee the robustness of the CG-index, we create multiple replicas for a cluster server. Replications are performed in two granularities. We replicate both the CG-index and local B^+ -tree index. When a cluster server is offline, we can still access its index and retrieve the data from DFS. The replicas are built based on BATON’s replication protocol. Specifically, the index entries maintained by a BATON node (master replica) are replicated in its left adjacent node and right adjacent node (slave replicas). Therefore, each node has 3 replicas (Dynamo [15] keeps 3 replicas typically. In Starfish [17], 3 replicas can guarantee 99.9% availability, if the compute node is online for 90% of time). The master replica is used to process queries and the slave replicas are used as backups. When a BATON node fails, we apply the routing tables to locate its adjacent nodes to retrieve the replicas. We first try to access the left adjacent node and if it also fails, we go for the right adjacent node.

In either lazy update or eager update, we need to guarantee the consistency between the replicas. Suppose BATON node N_i maintains the master replica of index entry E . To update E , we send the new version of E to N_i , which will forward the update to the living replicas. The corresponding BATON nodes, when receiving the update request, will keep the new version of E and respond to N_i . After collecting all the responses, N_i commits the update and ask other replicas to use the new index entries.

In BATON, the node occasionally sends ping messages to its adjacent nodes and nodes in its routing table. That ping message can be exploited to detect node failure. If we have not received the ping response from a specific node for k times, we assume the node fails and broadcast the information to all cluster servers. When a node fails, its left adjacent node is promoted to be the primary copy. If both the node and its left adjacent node fail, the right adjacent node is promoted to be the primary copy.

Each update is assigned a timestamp. When a BATON node restarts from failure, it asks current master replicas to get the latest




Figure 4: Query Throughput




Figure 5: CG-Index VS. ScalableBTree (range search)




Figure 6: Update Throughput




Figure 7: CG-Index VS. ScalableBTree (mixed workload of exact and range search)




Figure 8: CG-Index VS. ScalableBTree (mixed workload of queries and insertions)




Figure 9: Cost of Scaling Up/Downsizing

updates. By comparing the timestamp of an index entry, it replaces the stale entry with the new one. After that, it declares to be the master replica of the corresponding index data and starts serving the search. Note that, the query processing is resilient to node’s failure as suggested by the following theorem.

THEOREM 4. *In BATON, if the adjacent links and parent-child links are up-to-date, the query can be successfully processed, even if some nodes fail or the routing tables are not correct.*

PROOF. See the appendix. \square

7. EXPERIMENT EVALUATION

To evaluate the performance of our system, we deploy it on Amazon’s EC2 [1] platform. Details of experiment settings can be found in the appendix. For comparison purpose, we implement a distributed B^+ -tree index described in [10]. We use “ScalableBTree” to denote the index. The ScalableBTree is built on HP’s Sinfonia [11], a distributed file system. As Sinfonia’s code is not publicly available¹, we use a master server (large instance of EC2 with 7.5 GB of memory and 4 EC2 compute units) to simulate its behaviors, e.g., data locating service and transaction service. In the ScalableBTree index, each processing node acts as both client and server. Servers are responsible for maintaining the index and the clients are used to generate queries. The ScalableBTree is different from the CG-Index in that it maintains a large B^+ -tree over the network, whereas in the CG-Index, each node maintains a small local B^+ -tree. For the ScalableBTree, we create a distributed B^+ -tree with 10 million keys. Therefore, the total data size of ScalableBTree is less than that of the CG-Index. This is because for a large B^+ -tree,

¹The authors could not release the codes due to HP’s copyright concerns.

the size of internal nodes may be too large to be cached at the client (the ScalableBTree index proposes to lazily buffer internal nodes in clients).

7.1 Scalability

Figure 4 shows query throughput under different search ranges. The best performance is achieved for the exact search query ($s=0$). When the search range is enlarged, throughput degrades as more nodes are involved. Scalability increases when we increase the number of processing nodes. From Figure 5 to Figure 8, we show the performance comparison with the ScalableBTree index under different workloads. Figure 5 shows the CG-Index produces much higher throughput for range queries. In the CG-Index, after locating leaf nodes, a query is processed by the local B^+ -trees in parallel, while in the ScalableBTree, we cannot apply the parallel search algorithm, because the leaf nodes are randomly distributed in the cluster.

Figure 6 shows the update throughput of the system (in **logarithmic** scale). In the CG-Index, the node generates uniform insertions for its local B^+ -tree, while in the ScalableBTree index, the node issues uniform insertions for the distributed B^+ -tree. In the CG-Index, most updates can be processed by nodes locally, because we only insert internal B^+ -tree nodes to CG-index, which has few updates when update follows uniform distribution. Only a few requests, resulting in node splitting or merging, trigger a synchronization request to the network. In contrast, each insertion request in the ScalableBTree index triggers a network round-trip. If an internal node is being split or merged, it needs to broadcast the change to every node to update the version table.

In real systems, different types of operations are processed concurrently. In Figure 7, we generate a mixed workload of exact queries and range queries with selectivity 0.04. We vary the per-

percentage of range queries from 0% to 100%. That is, when the percentage is 0%, we have all exact match queries. ScalableBTree outperforms the CG-Index in exact match search because most queries only require one round trip to retrieve the data in the ScalableBTree, while, in the CG-Index, following the routing protocols of BATON, the query needs several hops to obtain the data. In fact, we can improve the search efficiency of the CG-index by adopting the same replication strategy of the ScalableBTree. However, this will incur higher maintenance overheads for the client nodes. In other cases (mixed workload), the CG-Index performs much better than the ScalableBTree. In Figure 8, a mixed workload is generated, with varying percentage of insertions and queries (exact query : range query = 6:4). The CG-Index demonstrates that it is superior to the ScalableBTree in handling a mixed workload of queries and updates efficiently.

7.2 Cost of Scaling-up and Downsizing

In Cloud systems where the storage and compute power are elastic by design, new service nodes may join or leave the cluster in batches. In this experiment, we double (e.g., 16→32) or half (e.g., 32→16) the node number to evaluate expansion cost and collapse cost respectively to evaluate the robustness and efficiency of our proposal with respect to the system elasticity. As each node holds the same amount of data, the collapse process needs to redistribute more data than the expansion case (16→32 moves half data of 16 nodes to others while 32→16 moves all data of 16 nodes to the rest). In Figure 9, we compare the cost of the two kinds of data distribution strategies: sequential expansion/collapse and parallel expansion/collapse. The x-axis label “16↔32” implies the network is expanded from 16 to 32 nodes or collapsed from 32 to 16 nodes. In the sequential expansion/collapse strategy, nodes join or leave the network one by one while in the parallel setup, all nodes join or leave the network simultaneously. Figure 9 shows that the parallel expansion/collapse strategy is more efficient than the sequential one. In the parallel expansion strategy, new nodes and old nodes are grouped into pairs. New nodes obtain data from the old ones and then they rebuild an overlay network. The parallel collapse strategy works in the inverse direction. The result confirms the robustness and efficiency of our proposal with respect to dynamic system re-configuration due to application loads.

8. CONCLUSION

We have presented the design and implementation of a scalable and high-throughput indexing scheme for SaaS applications in the Cloud environment. We assume a local B⁺-tree is built for the dataset stored in each compute node. And to enhance the throughput of the system, we organize compute nodes as a structured overlay and build a Cloud Global index, called the CG-index, for the system. Only a portion of local B⁺-tree nodes are published and indexed in the CG-index. Based on the overlay’s routing protocol, the CG-index is disseminated to compute nodes. To save maintenance cost, we propose an adaptive indexing scheme to selectively expand local B⁺-trees for indexing. Our scheme has been implemented and evaluated in Amazon’s EC2, a real-world Cloud infrastructure. The experimental results show that our approach is efficient, adaptive and scalable.

9. REFERENCES

- [1] <http://aws.amazon.com/ec2/>.
- [2] <http://aws.amazon.com/s3/>.
- [3] <http://hadoop.apache.org>.
- [4] <http://hypertable.org>.

- [5] <http://incubator.apache.org/cassandra/>.
- [6] <http://project-voldemort.com/>.
- [7] <http://www.comp.nus.edu.sg/~epic>.
- [8] K. Aberer, P. Cudré-Mauroux, A. Datta, Z. Despotovic, M. Hauswirth, M. Puceva, and R. Schmidt. P-grid: a self-organizing structured p2p system. *SIGMOD Record*, 2003.
- [9] A. Adya, W. J. Bolosky, M. Castro, G. Cermak, R. Chaiken, J. R. Douceur, J. Howell, J. R. Lorch, M. Theimer, and R. P. Wattenhofer. Farsite: federated, available, and reliable storage for an incompletely trusted environment. In *OSDI*, pages 1–14, 2002.
- [10] M. K. Aguilera, W. Golab, and M. A. Shah. A practical scalable distributed b-tree. *VLDB*, pages 598–609, 2008.
- [11] M. K. Aguilera, A. Merchant, M. Shah, A. Veitch, and C. Karamanolis. Sinfonia: a new paradigm for building scalable distributed systems. *SIGOPS*, 2007.
- [12] bitsource.com. Rackspace cloud servers versus amazon ec2: Performance analysis. 2010.
- [13] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber. Bigtable: A distributed storage system for structured data. *OSDI*, 2006.
- [14] A. Crainiceanu, P. Linga, A. Machanavajjhala, J. Gehrke, and J. Shanmugasundaram. P-ring: an efficient and robust p2p range index structure. In *SIGMOD*, 2007.
- [15] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels. Dynamo: Amazon’s highly available key-value store. *SIGOPS*, 2007.
- [16] D. DeWitt and J. Gray. Parallel database systems: the future of high performance database systems. *Commun. ACM*, 1992.
- [17] E. Gabber, J. Fellin, M. Flaster, F. Gu, B. Hillyer, W. T. Ng, B. Özden, and E. A. M. Shriver. Starfish: highly-available block storage. In *USENIX*, pages 151–163, 2003.
- [18] S. Ghemawat, H. Gobioff, and S.-T. Leung. The Google file system. In *SOSP*, 2003.
- [19] H. V. Jagadish, B. C. Ooi, K.-L. Tan, Q. H. Vu, and R. Zhang. Speeding up search in peer-to-peer networks with a multi-way tree structure. In *SIGMOD*, 2006.
- [20] H. V. Jagadish, B. C. Ooi, and Q. H. Vu. Baton: A balanced tree structure for peer-to-peer networks. In *VLDB*, 2005.
- [21] J. Kubiawicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, C. Wells, and B. Zhao. Oceanstore: an architecture for global-scale persistent storage. *SIGARCH*, pages 190–201, 2000.
- [22] E. Parzen. *Stochastic processes*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1999.
- [23] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Schenker. A scalable content-addressable network. In *SIGCOMM*, 2001.
- [24] G. Wang and T. S. E. Ng. The impact of virtualization on network performance of amazon ec2 data center. In *INFOCOM*, 2010.
- [25] J. Wang, S. Wu, H. Gao, J. Li, and B. C. Ooi. Indexing multi-dimensional data in a cloud system. In *SIGMOD Conference*, pages 591–602, 2010.
- [26] S. Weil, S. A. Brandt, E. L. Miller, D. D. E. Long, and C. Maltzahn. Ceph: A scalable, high-performance distributed file system. 2006.




Figure 10: A BATON Tree Overlay

APPENDIX

A. APPENDIX

A.1 BATON Overlay

In this paper, BATON is applied to organize the compute nodes. Detailed description of BATON protocols can be found in [20]. In BATON, each node is responsible for a key range, and each node maintains routing pointers to predecessor, successor, parent, child and sibling nodes. The BATON form of indexing is similar in spirit to that of the B-tree. If we travel the BATON tree in an inorder manner, we end up searching the key range sequentially.

In Figure 10, we show a BATON overlay, where dotted lines connect the routing neighbors and we mark the key range of each BATON node. Specifically, the nodes in H 's left or right routing table are H 's sibling nodes with a distant of 2^x to H ($0 \leq x \leq H.level - 1$). To lookup a specific key, the node will first check its own range. If the key is bounded by the range, it will do a local search. Otherwise, it will search its left or right routing table to locate a node most close to the key and forward the request to the node. If no such routing node exists, the lookup request will be forwarded to the parent, child or predecessor/successor node. In BATON, the search cost and maintenance cost are bounded by $O(\log_2 N)$ hops, where N is the number of nodes. A more efficient variant of BATON (BATON* [19]) reduces the search cost to $O(\log_b N)$ with a larger fan-out b at the expense of incurring much more maintenance overheads. Therefore, we use BATON in this paper. To support the range index, we only need to slightly extend the BATON overlay by recording the subtree range of each internal node.

The compute node acts as a BATON node in the overlay. Using the following interfaces (Table A.1) provided by BATON, we can organize the cluster system as a BATON overlay and search the index based on the routing protocols.

Table A.1 BATON Interface

join(IP)	Join the BATON network
leave()	Leave the BATON network
lookup(key)	Lookup for the node responsible for the key
store(key, value)	Publish a value using the key
remove(key)	Remove the values with a specific key

A.2 Proofs

Proof of Theorem 1.

PROOF. The adaptive indexing algorithm starts from publishing the root node of each local B^+ -tree. The initial state of the CG-index provides a correct snapshot of the local indexes, as the root nodes represent an overview of the local B^+ -trees. The initial state can be changed to different indexing states via expansion and collapse operations. If both expansion and collapse operation are

atomic operations, the indexing states must satisfy the following property for a local B^+ -tree \mathcal{T} (suppose its root range is $[l, u]$).

- Given a key k in $[l, u]$, we can find one and only one indexed node of \mathcal{T} in CG-index.

This is because we always replace an indexed node with all its child nodes and vice versa. Therefore, the adaptive indexing strategy can return a complete result in the query processing. \square

Proof of Theorem 2.

PROOF. In the CG-index, we generate a key range for each B^+ -tree node and publish the node based on the key range. The query algorithm also routes the query based on the key range of an indexed node. If the key range is not affected by the update, the index is still valid, because all queries involving the node is still sent to the correct node. \square

Proof of Theorem 3

PROOF. If no node fails, the eager update strategy can guarantee the CG-index is consistent with local B^+ -tree. If the node responsible for the indexed node fails, we cannot update all replicas of the CG-index node successfully. Therefore, the eager update will keep the old CG-index nodes and the local B^+ -tree node. The index is still consistent. If the node responsible for local B^+ -tree fails after all CG-index replicas are updated, the CG-index may not be consistent with local index. However, it only triggers false positive. Therefore, the query processing is still correct. \square

Proof of Theorem 4.

PROOF. In BATON, when a node joins the system, it obtains its adjacent links and parent-child links directly from its contacting node, while its routing neighbors are obtained via a stabilization process. The routing neighbors are used to facilitate the routing. However, without them, BATON can still route queries based on adjacent/parent-child links. Even if we route the query based on an inconsistent routing link, we can correct the routing process via adjacent/parent-child links. If adjacent links or parent-child links are incorrect due to node joining, the routing process will fail. However, in Cloud system, node will not frequently join or leave the system. \square

A.3 Algorithms

Algorithm 5 shows the publication process in BATON. We first obtain the compute node responsible for the lower bound of the B^+ -tree node based on the BATON routing protocols (line 1). Then, step by step, we forward the request to the upper level nodes until we reach the one whose subtree range can fully contain the B^+ -tree node's range. In line 4, the compute node stores the meta-data of remote B^+ -tree node in the disk and buffers it in memory. The stored information can be used to process queries. In lines 6 and 10, we tag the nodes with two values, indicating whether the query should be forwarded to the parent node (e.g. the parent node stores an index node whose range overlaps with the child node).

Algorithm 6 shows our parallel search algorithm. We first find a node that can fully contain the search range (lines 1-3). Function $lookup([l, u])$ returns a compute node that overlaps with the search range $[l, u]$. As discussed in section 4, instead of returning the node responsible for the lower bound of the range, we return a node that overlaps with the search range. This optimization reduces the overhead of routing. Then, we broadcast the query message to the nodes within the subtree (line 4-7). The broadcast messages are sent to the nodes in parallel. After a node receives the search request, it starts searching its local CG-index. Besides the node itself, we need to search the possible results in the ancestor nodes (lines 8-10). Finally, the index search result (a set of indexed nodes, S_b) is

Algorithm 5 Publish(n)

// n is a B^+ -tree node for indexing

- 1: $N_i = \text{lookup}(n.\text{low})$
- 2: **while** TRUE **do**
- 3: **if** N_i .subtree contains $(n.\text{low}, n.\text{up})$ **then**
- 4: store n at N_i
- 5: **if** n 's range overlaps with N_i 's right subtree **then**
- 6: tagSet(N_i .rightchild)
- 7: break
- 8: **else**
- 9: **if** N_i .parent! = null **then**
- 10: Update the tag value of N_i
- 11: $N_i = N_i$.parent
- 12: **else**
- 13: break

returned to the query sender. Algorithm 7 shows the broadcast process. The broadcasting is processed in a recursive way. To reduce the network overheads, only the nodes within the search range will receive the query (e.g. we do not invoke the broadcast algorithm for the subtrees outside the search range).

Algorithm 6 ParallelSearch($Q = [l, u]$)

- 1: $N_i = \text{lookup}([l, u])$
- 2: **while** N_i 's subtree range cannot contain Q **do**
- 3: $N_i = N_i$.parent
- 4: **if** N_i .leftchild.subtree overlaps with Q **then**
- 5: broadcast(N_i .leftchild, $Q = [l, u]$)
- 6: **if** N_i .rightchild.subtree overlaps with Q **then**
- 7: broadcast(N_i .rightchild, $Q = [l, u]$)
- 8: **while** N_i 's tag values in the search range **do**
- 9: local search on N_i and put the indexed nodes overlapping with Q into set S_b
- 10: $N_i = N_i$.parent
- 11: forward S_b to query requestor

Algorithm 7 Broadcast(compute node N_i , $Q = [l, u]$)

- 1: local search on N_i
- 2: **if** N_i is not leaf node **then**
- 3: **if** N_i .leftchild.subtree overlaps with Q **then**
- 4: broadcast(N_i .leftchild, $Q = [l, u]$)
- 5: **if** N_i .rightchild.subtree overlaps with Q **then**
- 6: broadcast(N_i .rightchild, $Q = [l, u]$)

Suppose each BATON node can share M bytes memory and each B^+ -tree node's index entry requires E bytes, we can only support up to $\frac{M}{E}$ indexed nodes. If the corresponding compute nodes have enough memory for storing the child nodes' index, the index is built successfully. Otherwise, index replacement is triggered.

Algorithm 8 generalizes the index replacement process. The index replacement happens when a new B^+ -tree node is inserted into CG-index. If there are more memory for the incoming nodes, we just accept it and notify the requester (line 2-4). Specifically, let S_n be the nodes in local CG-index. The B^+ -tree nodes in S_n can be classified into two types. Suppose $n_j \in S_n$, and n_p is n_j 's parent node. If $n_p \in S_n$, and n_j does not have sibling nodes in S_n , replacing n_j with n_p cannot reduce the storage load of N_i . Thus, n_j is not a candidate for index replacement. By removing such type of nodes from S_n , we get a candidate set S'_n (line 7-11). The new B^+ -tree node n is inserted into S'_n , and we rank the nodes in S'_n based on the query histogram (line 12). Let $\min(S'_n)$ be the node with the least rank. If $\min(S'_n) = n$, N_i rejects the index request for the new node n (line 13-14). Otherwise, N_i replaces $\min(S'_n)$ with n in its index and triggers a tree collapse (line 16-17).

To guarantee atomic indexing, if the node receives a reject notification, we will roll back the index (keep the current indexed node and remove the newly built index).

Algorithm 8 IndexReplace(B^+ -tree node n , compute node N_i)

// n is a new node to be indexed at N_i

- 1: $S_n = N_i$'s index set
- 2: **if** S_n is not full **then**
- 3: accept n in S_n
- 4: notify the sender with success
- 5: **else**
- 6: $S'_n = \{n\}$
- 7: **for** $\forall n_j \in S_n$ **do**
- 8: **if** getsibling(n_j, S_n) = null **then**
- 9: $n_p = n_j$'s parent
- 10: **if** N_i .range cannot cover n_p .range **then**
- 11: $S'_n = S'_n \cup \{n_j\}$
- 12: rank nodes in S'_n by query histogram
- 13: **if** $\min(S'_n) == n$ **then**
- 14: reject n and notify the sender
- 15: **else**
- 16: remove $\min(S'_n)$ and triggers a tree collapse
- 17: notify the sender with success

A.4 Details of Tuning Algorithm

Let S_n represent the remote B^+ -tree nodes indexed at the compute node N_i . Then, in a time period of T_1 , N_i records a query set S_q for S_n . Based on S_q , we can estimate the cost of the current indexing strategy and perform some optimization. For this purpose, we build a query histogram at each compute node. Basically, suppose N_i 's subtree range is $[l_i, u_i]$, we partition the range into k equal-length cells. Thus, cell j covers the range $[\frac{j(u_i-l_i)}{k}, \frac{(j+1)(u_i-l_i)}{k})$. Given a query $q \in S_q$, suppose there are x cells involved in q , we increase the counter of these cells by $\frac{1}{x}$. Finally, we get a counter array $H = \{c_0, c_1, \dots, c_{k-1}\}$ for the query distribution.

Given an indexed B^+ -tree node n_i , we can compute its query cost by searching H . Let R_i denote the histogram cells overlapping with n_i . The current cost of n_i is estimated as:

$$\text{cost}(n_i) = \alpha \frac{h(n_i)}{T_1} \sum_{x \in R_i} c_x + \frac{\beta}{2} (kg(n_i) + \frac{1}{T}) \log_2 N \quad (5)$$

As mentioned before, we have two alternative indexing strategies: indexing the child nodes of n_i and indexing the parent node of n_i . Let n_{ij} represent n_i 's child node. Suppose n_i has m_0 child nodes, the strategy of indexing the child nodes incurs a cost of:

$$\text{cost}(n_i.c) = \frac{(h(n_i) - 1)\alpha}{T_1} \sum_{x \in R_i} c_x + \frac{\beta}{2} \left(\sum_{j=0}^{m_0-1} kg(n_{ij}) + \frac{m_0}{T} \right) \log_2 N \quad (6)$$

Suppose n_i has m_1 sibling nodes, the strategy of indexing the parent node incurs a cost of:

$$\text{cost}(n_i.p) = \alpha \frac{(h(n_i) + 1)}{T_1} \sum_{i=0}^{m_1} \sum_{x \in R_i} c_x + \frac{\beta}{2} (kg(n_i.p) + \frac{1}{T}) \log_2 N \quad (7)$$

Equation 5 and 6 can be computed by the node's local information while Equation 7 needs information from the sibling nodes. Figure 3 illustrates a possible indexing strategy in the system, where the shaded rectangles represent the B^+ -tree nodes being indexed. If node i wants to estimate the cost of indexing its parent f , it needs to obtain the query distribution from its sibling node j . Given that node i does not know the details of its siblings, it is difficult to collect the necessary information. An alternative is to collect the

status of the child nodes in the parent node, e.g., node f periodically checks the status of node i and j . As node f is not indexed, the “pull” scheme is not applicable. Instead, we use the “push” method. The indexed B^+ -tree nodes will periodically report their query distribution information to the compute node that handles their parent’s range. After collecting all the information, the compute node decides on the indexing strategy. If it can save cost by indexing the parent B^+ -tree node, the compute node will issue a process to delete all indexes about the child B^+ -tree nodes and notify the corresponding compute node to publish the parent tree node.

The major cost of the adaptive approach is the cost of reporting the status of the child nodes. To reduce overhead, we propose an efficient optimization. As observed from Figure 3, node a does not have efficient information to change the indexing strategy, unless all its child nodes b, c, d and e are indexed. Based on Theorem 5, only nodes i, j, k and l need to report the status to their parents. We therefore greatly reduce the communication cost.

THEOREM 5. *The indexed B^+ tree node needs to report its status to the parent node, i.f.f. none of its siblings has an indexed descendant node.*

PROOF. In the tuning algorithm, the B^+ tree node n is indexed in the network if it does not have an indexed descendant node. If all sibling nodes do not have an indexed descendant node, all the siblings are indexed. Hence, the parent node can receive reports from all the child nodes to decide whether to change the indexing strategy. \square

A.5 System Optimizations

To further improve the performance of CG-index, we propose three optimizations. Routing buffer is used to reduce the routing overhead of the overlay. Selective expansion is used to reduce the maintenance overheads of CG-index. And single local search is proposed to reduce the index search cost.

A.5.1 Routing Buffer

Locating a specific key in the overlay incurs a cost of $O(\log_2 N)$, where N is the total number of nodes in the overlay. To reduce routing cost, we apply a buffering approach. In a successful *lookup(key)* operation in the overlay, the compute node responsible for the *key* will notify the requester about its key range and IP address. The requester, once receiving the information, will store the key range and IP address in its routing buffer. The routing buffer is limited to s entries and is maintained with the LRU strategy. In the future routing process, the node will check both its routing table and routing buffer. The node nearest to the search key is selected as the next hop. As the network is stable in Cloud systems, the routing buffer can efficiently reduce the routing overheads. Even if the routing buffer is not consistent with the network, the query can be routed to the destination based on the routing table. To detect the stale routing buffer, the sender attaches the expected destination in the message. The receiver will check its status against the expected one. And it will notify the sender to update its routing buffer if it is not the expected receiver.

A.5.2 Selective Expansion

The adaptive scheme can effectively tune the index based on query distribution. It expands the B^+ -tree step by step. Heavily queried nodes have a high probability of being published. However, if the query distribution is skewed, we do not need to publish every child nodes. In B^+ -tree, the order of nodes are always set to a large value (based on the disk block size). In our experiments,




Figure 11: Example of Selective Expansion Strategy

each node can support up to 100 child nodes. If most queries focus on a small number of child nodes, we can save the indexing cost by only publishing the corresponding nodes.

In Algorithm 3, we compare the cost of indexing the current node with the cost of indexing all the child nodes. As a matter of fact, the order of B^+ -tree may be quite large. Indexing all the child nodes is not necessary and incurs too much overhead. In Figure 3, suppose the queries focus on the range $[20, 50]$, we do not need to index the nodes b, d and e . Instead of indexing all the child nodes, we only select the beneficial ones for indexing.

Figure 11 shows the selective expansion tree of Figure 3. In the selective expansion strategy, the parent node is kept in the index if not all of its child nodes are indexed. For example, in Figure 11, node a is responsible for the ranges of its three none-indexed child nodes, b, d and e .

Given an indexed B^+ -tree node n_i with m child nodes (denoted as $\{n_{ij} | 0 \leq j \leq m - 1\}$), we define an m -element vector $V = \{v_0, \dots, v_{m-1}\}$. v_j is 1, if the node n_{ij} is selected to be indexed. Otherwise, v_j is 0. We can compute the indexing cost for a specific V .

The optimal solution is to find a vector V that can minimize the above cost. Using brute force to search the solution is not practical as there are 2^m possibilities. If we further consider memory size, the optimal indexing problem is reduced to a 0-1 knapsack problem. Instead of searching for the optimal solution, we use a simple but efficient heuristic method.

In fact, the cost of indexing a child node can be considered to comprise two parts. First, the indexing benefit of query processing is computed as:

$$benefit(n_{ij}) = \alpha \frac{1}{T_1} \sum_{x \in r_{ij}} c_x \quad (8)$$

Then, the cost of maintenance is estimated as:

$$cost_m(n_{ij}) = \frac{\beta}{2} (kg(n_{ij}) + \frac{1}{T}) \log_2 N \quad (9)$$

A greedy heuristic method is to index the child node if its benefit is greater than its maintenance cost until memory is full. This can provide us with a good enough indexing plan. Algorithm 9 shows the selective expansion scheme. The parent node decides whether to index each specific child node individually. If a child node is indexed, the parent node needs to be split (line 6). Let $[l_i, u_i]$ and $[l_{ij}, u_{ij}]$ represent the ranges of the parent node n_i and its child node n_{ij} , respectively. After indexing node n_{ij} , we split the range of n_i into $[l_i, l_{ij}]$ and $[u_{ij}, u_i]$. We remove the current index of n_i and insert two new index entries based on the new ranges. The insertion of the new index entries must be atomic. If it fails due to memory limitation, we roll back the indexing operation and keep the old index. In an extreme case, if all child nodes are beneficial to indexing, the selective expansion scheme evolves into the full expansion scheme.

In the selective expansion scheme, we keep a record of how the B^+ -tree is indexed in the owner node. We generate a $2m$ -length

Algorithm 9 `SelectivelyExpand()`

```
1: compute the query histogram H
2: for  $\forall$  B+-tree node  $n_i \in S_n$  do
3:   for  $\forall n_i$ 's child  $n_{ij}$  do
4:     if  $n_{ij}$ 's benefit is greater than its cost then
5:       index  $n_{ij}$ 
6:       split the index range of  $n_i$ 
7:   if  $n_i$  does not have an indexed descendant then
8:     if  $n_i$ 's benefit is less than its cost then
9:       remove  $n_i$ 's index
10:    statusReport( $n_i$ )
```

bitmap for each B⁺-tree node, where m is the order of the tree. If the subtree rooted at i th child has been expanded for indexing, we mark the i th bit of the bitmap to 1. Based on the bitmap, the owner node can help collapse the tree if necessary. Algorithm 10 shows the collapse operation for the selective collapse strategy. On receiving an index removal notification, the owner node checks the corresponding bitmap and combines the index entries if necessary. First, it searches for the index entries that can be combined with the removed child index (lines 4 to 7). Let $I(i, j)$ denote the index entry for the range from i th child to j th child. The removed child index are combined with the left or right adjacent index entries (lines 8-14).

Algorithm 10 `SelectivelyCollapse(B+-tree node n_i)`

```
//recieve the status report from node  $n_i$ 
1: map= $n_i$ 's parent bitmap
2: if map[i]==1 then
3:   map[i]=0, x=i, y=i
4:   while  $x-1 \geq 0$  and map[x-1]==0 do
5:     x=x-1
6:   while  $y+1 < m$  and map[y+1]==0 do
7:     y=y+1
8:   if  $x \neq i$  and  $y \neq i$  then
9:     combine  $I(x, i-1), I(i, i)$  and  $I(i+1, y)$  into  $I(x, y)$ 
10:  else
11:   if  $x \neq i$  then
12:     combine  $I(x, i-1)$  and  $I(i, i)$  into  $I(x, i)$ 
13:   else
14:     combine  $I(i+1, y)$  and  $I(i, i)$  into  $I(i, y)$ 
```

Index replacement can be handled in the same way as the full expansion case. As a matter of fact, the selective expansion strategy reduces the cost of index replacement. As child nodes are indexed individually in the selective expansion strategy, once it is decided that a B⁺-tree node is to be removed from the index, we do not need to find and remove all its siblings. The selective expansion strategy makes our adaptive indexing scheme more flexible.

A.5.3 Single Local Search

The adaptive indexing scheme allows us to process queries with indexed B⁺-tree nodes. After locating the indexed B⁺-tree nodes, we forward the query to the corresponding local B⁺-trees to complete data retrieval. Given a query $Q = [l, u]$, let S_b be the set of indexed B⁺-tree nodes returned by Algorithm 6. We group the nodes in S_b by their owners. Suppose $S_b(N_i)$ denote the B⁺-tree nodes from the compute node N_i . We need to access N_i 's local B⁺-tree based on the B⁺-tree nodes in $S_b(N_i)$. A close analysis reveals that only one node in $S_b(N_i)$ is required to be accessed.

In a B⁺-tree, to retrieve the data within a continuous range, we first locate the leaf node responsible for the lower bound of the search range. Then, we scan the corresponding leaf nodes by following the leaf nodes' links. All the involved internal nodes reside

in the path from the root to the first searched leaf node. The other internal nodes, though overlapping with the search range, are not searched. This observation motivates an optimization.

LEMMA 1. *For a range query $Q = [l, u]$, the indexed B⁺-tree nodes from the same compute node N_i (e.g. $S_b(N_i)$) involved in the query can be sorted into a continuous range based on their responsible ranges.*

PROOF. Our adaptive indexing scheme guarantees that there is no overlap between the B⁺-tree nodes' responsible ranges, and that for any search point in the domain, there is an indexed B⁺-tree node whose responsible range contains it. Thus, nodes in $S_b(N_i)$ can be sorted into a continuous range based on their ranges. \square

LEMMA 2. *For a range query $Q = [l, u]$ and the B⁺-tree node set involved $S_b(N_i)$, we sort the nodes in $S_b(N_i)$ by their ranges. Only the first B⁺-tree node in $S_b(N_i)$ triggers a local search.*

PROOF. Directly derived from Lemma 1 and the B⁺-tree's search algorithm. \square

Given a specific query $Q = [l, u]$ and an indexed B⁺-tree node n_i , the compute node can decide whether to issue a local B⁺-tree search based on Theorem 6.

THEOREM 6. *The B⁺-tree node n_i with range $[l_i, u_i]$ incurs a local search for query $Q = [l, u]$, only if $l_i \leq l \wedge l \leq u_i$ or n_i is the left-most node of its level and $l \leq l_i \wedge l_i < u_i$.*

A.6 Experiment Settings

The compute unit (small instance) in EC2 is a virtual server with a 1.7 GHz Xeon processor, 1.7 GB memory and 160 GB storage. Compute units are connected via a 250 Mbps network. Our system is implemented in Java 1.6.0. Table A.6 lists the experiment settings. In our system, each node hosts 500k tuples. The tuple format is (*key*, *string*). The *key* is an integer key with the value in the range of $[0, 10^9]$ and the *string* is a randomly generated string with 265 bytes. The data are sorted by the keys and grouped into 64M chunks. Therefore, each compute node hosts two chunks. We generate exact queries and range queries for the keys in zipfian distribution. When the skew factor is 0, the queries are uniformly distributed. The major metrics in the experiment are query throughput and update throughput. Based on the reports of [24, 12], we set $\frac{\alpha}{\beta} = 0.5$ (the random disk read is slower than TCP/IP message sending).

Table A.6 Experiment Settings

Name	Default Value
node number	256
memory size	1M
use routing buffer	false
skew factor (sf)	0.8
default selectivity (s)	0.04
adaptive period	10 sec

In our implementation, the page size of the local B⁺-tree is set to 2K and the maximal fan-out is about 100. Before the experiments begin, we load 500K keys into each local B⁺-tree. The total number of tuples therefore varies from 8 million to 128 million. We use a simulator to act as clients. The processing nodes receive queries from the simulator continuously. After processing one query, a node will ask the simulator to obtain a new query. Thus, users' queries are processed in parallel. In each experiment, 1000*N* queries are injected into the system, where *N* is the number of nodes in the system. Each experiment is repeated for 10 times and we take the average result.