

Optimizing Enterprise Search Performance Using EHCACHE-Backed Apache Lucene Indexing for Hybrid Caching Systems

Deepak Singh^[0009-0001-9381-8797]

Independent Researcher, USA

deepaksingh1981@gmail.com

Published

2022-10-16

Issue

[Vol. 4 No. 4 \(2022\): AJCDI](#)

Abstract

Enterprise search systems are very important in facilitating rapid retrieval and retrieval of information in large volumes. Nonetheless, with the amount of data and complexity of queries, low query latency and high throughput becomes a major issue. The classical search models with only indexing schemes usually have performance bottlenecks with regard to redundancy in query execution and low reuse of computed responses. The proposed paper is a hybrid caching and optimization search framework that will combine EHCACHE with the search indexing of Apache Lucene to improve search performance within enterprise applications. The suggested solution makes use of caching in-memories to save queries and search results that are most commonly used and thus eliminates repetitive processes and also accelerates the response times. Through integration of EHCACHE and efficient indexing and retrieval system of Lucene, the system is able to process queries faster without compromising accuracy and relevance. The framework presents smart cache invalidation methods to make sure that there is consistency in data, such as time-expired, event-based cache invalidation, and part of the cache update based on the index update. These solutions solve one of the main problems of the caching systems a tradeoff between performance progress and the appropriateness of the information. In order to gauge the solution effectiveness, comparative study is done between the cached and non-cached search systems checking benchmark and real world query work loads. Such performance metrics like query latency, throughput, cache hit ratio, and resource utilization in a system are examined. The results of the experiment indicate that the time

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal

Australian Journal of Cross-Disciplinary Innovation

Impact Factor: 16.4

ISSN-3746-757x (Online)

Acceptance Rate: below 5%

required to respond to the queries is greatly reduced and the system efficiency is enhanced when the hybrid caching approach is employed. The results indicate a good combination of EHCACHE with Apache Lucene produces a more efficient solution of scalability in enhancing enterprise search systems. This is not only a way of improving performance but also, reliability and flexibility in dynamic data environments where it is appropriate in the modern applications that demand high speeds of information retrieval.

Keywords: hybrid caching, enterprise search, EHCACHE, Apache Lucene, search optimization, query latency reduction, cache invalidation, in-memory caching, indexing, information retrieval

Introduction

Enterprise search systems have emerged as one of the foundations of contemporary digital infrastructures that allow organisations to access, analyse and use unlimited volumes of organised and unorganised data effectively. As the volume of information produced by enterprise applications, logs, documents and user interactions continues to grow exponentially, there has been a growing inclination towards need to have search solutions that are of high performance. Indexing frameworks, among other technologies, have become very important in making it possible to access the data in very little time, but with the complexity of queries and the ever-increasing data volume, ensuring reduced latency and high throughput has become an issue to concern. Even optimized traditional search systems which utilize efficient indexing methods may be troubled with the challenge of execution of repeated queries, computation intensive operations, and a longer response time when overloaded. User queries are repetitive, which is also one of the main bottlenecks of performance of enterprise search systems. In most practical situations, much of search queries is of repetitive or like-wise queries particularly in applications like e-commerce, knowledge management systems, and analytics dashboards. A query to a relational database that involves repeated processing with the help of only indexing facilities may result in unnecessary computational costs and longer latencies. This ineffectiveness is further amplified in busy environments where the system has to serve thousands of millions of queries within a given second. Consequently, there has been increased demand towards complementary methods of minimizing redundant operations as well as improving the overall system performance.

Caching has also proven to be a potent measure to solve these problems as it can store the most visited information in memory thus allowing it to be accessed much quicker than when it was stored in a disk. The in-memory caching systems are important in minimizing the computational overhead and repetitive data retrieval that would otherwise slow down response time and efficiency of the system. Through caching, enterprise search systems are able to deliver commonly requested queries out of memory, avoiding the processing that has to be done in the underlying index to obtain them. But there is no easy way to implement caching in search systems with caching introducing new complexities in the form of caching management,

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal

Australian Journal of Cross-Disciplinary Innovation

Impact Factor: 16.4

ISSN-3746-757x (Online)

Acceptance Rate: below 5%

inconsistency, and invalidation. The major issue of caching is what should be cached and how long the data should be maintained. Caching strategies should take into account the query frequency, relevance of results and freshness of the data in the case of search systems. It is essential that stale or outdated results are not stored because there are incorrect search results that can occur in most enterprise applications. So, proper cache invalidation techniques are necessary in order to make sure that data stored in a cache is in line with the underlying index. Various methods like time-based expiration, event-conditioned invalidation, and selective updates to the cache have been suggested to deal with this problem, and all have their own performance/consistency trade-offs.

The other issue to pay attention to is integration of caching and indexing frameworks. Apache Lucene is an open-source search library that is popular as a search tool that can be used to provide efficient indexing and retrieval with inverted indexes and sophisticated search algorithms. Although Lucene is best suited to complex querying and processing of numerous datasets, it is not the remedy to repeated querying. Combining the commercial Disk cache like EHCache with on-heap caching and full text search like Lucene, it can be designed to use the merits of both technologies. The main concerns in this scenario are Lucene that is used to process and index documents and EHCache that is used to store query results which are used frequently so that repetitive calculations will not be made. Hybrid caching in search systems refers to the strategic combination of indexing and caching in search system to realize optimum performance. Not only does this method reduce query latency; it also enhances throughput of the system by reducing resource consumption. Nevertheless, it is imperative to consider several factors to design an effective hybrid caching system such as a cache size, eviction and synchronization with the underlying index. As an example, over-aggressive caching strategy can result in data being stale whereas conservative approach can fail to leverage the advantages of caching. It is very important to have a balance of these factors in order to realize the expected performance.

Scalability is another important criterion of modern enterprise search systems besides optimization in terms of performance. As companies keep expanding, it has to ensure that their search systems can support the rising amount of data and user requests, without compromising performance. Distributed caching and indexing solutions can offer a solution to such scalability as a solution to such scalability is the ability to horizontally scale across many nodes. The distribution of the index and the cache also allows the systems to scale to higher workloads and offer an overall consistent performance even in varying conditions. Such a dispersed design also contributes to fault tolerance so that an able system is still operational even when some of the available components are not functioning. Performance evaluation and benchmarking are important in the process of determining the suitability of hybrid caching mechanisms. Under

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal

Australian Journal of Cross-Disciplinary Innovation

Impact Factor: 16.4

ISSN-3746-757x (Online)

Acceptance Rate: below 5%

controlled conditions, comparing the actions of a system with a cache and one without, it is possible to measure the gains made in such metrics as query latency, throughput, and resource utilization. Such analyses offer great understanding on the trade-off of various caching policies and aid in establishing best configurations in particular applications. The validity of these evaluations is further supported by real workloads in the world that are representative of the actual usage patterns and system behavior.

Although there are benefits of hybrid caching, some challenges exist on its application. These involve control of consistency of the cache, control of the dynamic updates of the data and optimization of the eviction policies in the cache. Also, caching and indexing systems integration should be smooth to prevent bringing new complexity and overhead. To deal with these challenges we need both sophisticated algorithms and effective system design as well as constant surveillance. This paper will seek to solve these issues by presenting and recommending a hybrid caching and search optimisation system to be a combination of EHCACHE and Apache Lucene indexing. The structure emphasizes query latency minimization, enhancements on throughput and data consistency ought to be provided with efficient methods of cache invalidation. Through an extensive analysis and benchmarking study, the study has shown that the hybrid caching could have a tremendous effect on improving the performance of enterprise searches. Finally, the aim is also to offer a scalable and efficient solution to support the requirements of modern enterprise applications so as to fast track and rely information access in an ever-expanding data-intensive world.

Literature Review

There is sufficient literature covering the area of enterprise search and information retrieval, although the initial research aimed at indexing methods and query optimization. The use of inverted indexing structure was significant in traditional search systems and this allowed them to retrieve the documents effectively on the basis of matching the keyword. Apache Lucene and other systems have been well known in supporting operations of large-scale text indexing and searching in a highly efficient manner. Research studies have shown that the use of inverted indexes, tokenization and ranking algorithms in Lucene enhance the accuracy and speed of search substantially. Nevertheless, even with these developments, scientists have pointed out the shortcoming in terms of repeated query processing and low-latency enforcement when the system is subjected to heavy tasks. The caching mechanisms have been proposed to overcome performance bottlenecks of performance of search systems which is an additive measure of indexing. The studies on caching in information retrieval systems are pointing out that caching is effective in lowering the time taken in query processing through the keeping of the often searched results in the memory. Initial caching methods laid emphasis on static caching where results were stored over a predefined period. Even though the approach enhanced performance, it was prone to causing problems with stale data and low accuracy. Later researchers added the

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal

Australian Journal of Cross-Disciplinary Innovation

Impact Factor: 16.4

ISSN-3746-757x (Online)

Acceptance Rate: below 5%

idea of dynamic caching strategies that change with query frequency and load on the system which enhanced both efficiency and relevance. In-memory caching products are receiving a huge focus because of their capability to offer low-latency access to data. It has been proposed in literature (several times) that combining in-memory caching and search systems can significantly decrease responding time by lowering the disk I/O activities. EHCACHE in general and specifically have been extensively employed in enterprise applications due to its ability to scale, be flexible as well as its support of various eviction policies and expiration policies. To improve the performance of a cache, researchers have examined various caching techniques, such as Least RecentlyUsed (LRU), Least FrequentlyUsed (LFU), and time-to-live (TTL)-based expiration among others. All these strategies have their own advantages and disadvantages based on the characteristics of workload and query patterns.

This is because cache invalidation has been one of the most difficult facets in caching with search systems. Research studies focus on the fact that it is extremely important that data consistency is maintained between the cache and the underlying index thus ensuring precision in the search results. Several methods of invalidation have been suggested such as time expiration, event invalidation and selective cache updates. Time-oriented techniques are easy to apply yet can result in stagnant data unless the expiration time is selected with caution. On the other hand, event based invalidation provides instant updates by invalidating cache entries whenever the underlying data changes. Selective invalidation methods seek to give a compromise between updating all the affected cache pieces and minimizing overheads that would otherwise lead to full consistency. Hybrid caching as a combination of the indexing and caching techniques has also turned out to be the bright idea in search performance enhancement. The studies in this field recommend that hybrid systems can exploit the capabilities of both methods effective search of data with indexing and fast accessibility with caching. It has been demonstrated that hybrid architectures provide considerable better query latency and throughput, especially in a high traffic environment characterized by repetitive query patterns. Nevertheless, the efficacy of the hybrid caching relies on the size of a cache, the operation of the workload, and the arrangement of the system.

One area of evaluation of search optimization methods using performance benchmarking has been a major attraction. Comparative studies have been carried out between the cached and non -cached systems by analyzing measures like response time, throughput, cache hit ratio, and resource utilization. These experiments continually show that the caching can save query latency by a significant factor and in many cases, over 50 per cent, related to the workload. Nevertheless, they also point to the necessity of trading-off cache performance and data freshness that aggressive caching strategies may result in obsolete outcomes. Another important issue that is implemented in the literature is scalability. Enterprise systems evolve

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal

Australian Journal of Cross-Disciplinary Innovation

Impact Factor: 16.4

ISSN-3746-757x (Online)

Acceptance Rate: below 5%

and grow in size, necessitating the search solutions to process more data and handle more requests. To handle this challenge, distributed search and caching systems have been suggested, which allow horizontal scaling of various nodes. It has been shown that distributed caching with distributed indexing can have a substantial performance and reliability improvement in systems. Nevertheless, these systems also bring with them other complications like paging, consistency control and network overheads.

Novel analysis has also been carried out on the combination of real-time data processing and search systems. It is necessary to have up-to-date search results in the changes of information that hasn't been uploaded recently in dynamic environments. Various techniques like incremental indexing and real time updates of the cache have been suggested in order to meet this requirement. These strategies make sure that the new data is instantly captured in search results without causing much effect on the system performance. In spite of the development of search optimization and caching, there are a number of issues. Researchers have found problems associated with the cache coherence, effective invalidation as well as optimal resource allocation. As well, adaptive mechanisms are required in which caching strategies can dynamically change according to the patterns of workload and system states. Another possible area or field of potential research was also proposed to include the integration of machine learning approaches on predictive caching and optimization. Summarizing, the literature indicates the need to integrate indexing and caching methods to enhance the performance of enterprise search. Whereas indexing systems like Apache Lucene offer effective methods of recalling and accessing data, caching systems like EHCACHE are more efficient in terms of speed and reduce unnecessary calculations. Hybrid caching solutions have a lot of potential, yet certain issues should be being considered on how to implement them so that the threats associated with consistency, scaling, and resources are being addressed. The paper is based on the previous studies and developed into proposing an optimized hybrid caching model that combines EHCACHE and Lucene aimed at optimizing query latency, throughput, and data consistency in enterprise search.

Methodology

The suggested approach includes a hybrid model consisting of the in-memory caching and search indexing techniques to maximize the performance of enterprise search. The strategy is a hybrid of Apache Lucene to provide efficient indexing and search with EHCACHE to provide fast in-memory response with low latency to ensure query latency, throughput, and data consistency. The methodology is divided into several steps and they are the architecture design of the system, data indexing, integrate a cache, cache invalidation schemes and performance testing.

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal

Australian Journal of Cross-Disciplinary Innovation

Impact Factor: 16.4

ISSN-3746-757x (Online)

Acceptance Rate: below 5%

This begins with the system architecture design that involves the creation of a layered model in which indexing, caching, and query processing units are separated. The architecture is made up of data ingestion layer, indexing layer, caching layer, and query execution layer. The incoming data gets initially treated and indexed on Lucene to generate an optimized inverted index that is easy to access. The caching layer will be combined with the indexing layer to store query results and other intermediate calculated results of the most accessed queries so that they can respond quicker upon repetition of the query. The second step is an indexing of the data with Apache Lucene. The surviving data of the enterprise (documents, logs and organized records) is tokenized, normalized and indexed in this phase with the inverted indexing feature of Lucene. Advanced methods of indexing such as stemming, removing of stop words and field-based indexing are implemented to make search accurate and efficient. The index is also updated by periodically updating it to take into account changes in underlying data so that the results of searches can be relevant and up-to-date. Minimization of overhead and real-time data updating is achieved through incremental indexing.

After indexing, the methodology involves integration of cache with the help of EHCACHE. EHCACHE is set up to act as an in-memory cache layer that is used to store the most frequently used queries as well as their results. Upon receiving a query, the system initially searches the cache to find a corresponding query. When a cache hit is found, the result is immediately returned without the index having to be executed in full against the query. When there is a cache miss, the query is then handled by Lucene and then it is stored in the cache to be used in the future. This method will also diversify and minimize unnecessary calculations and enhance the performance of the system. Cache management strategies are applied to ensure maximum performance of the cache. They encompass least recently used (LRU) and least frequently used (LFU) eviction policies which dictate how the entries in the cache are pushed out of the memory when the memory reaches its capacity. There is also time-to-live (TTL) which is used to manage the duration of data in each of the cache as time runs out, automatically eliminating old results. The size of the cache is optimally set on the use of system resources and workload related to the nature of workloads to find an appropriate balance between the use of the memory section and the performance gains.

Cache invalidation and consistency management is also a critical component of the methodology. Enterprise data is dynamic, so it is important to make sure that the results on the cache are consistent with the underlying index. The invalidation methods used in the framework include time-based and event-based invalidation, and selective cache update. Time-based expiration guarantees that cache entries are updated on a certain period of time whereas event-based invalidation keeps the cache updated every time a change is made to the indexed data. The selective invalidation only invalidates affected cache entries with a minimized

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal

Australian Journal of Cross-Disciplinary Innovation

Impact Factor: 16.4

ISSN-3746-757x (Online)

Acceptance Rate: below 5%

overhead at the cost of accuracy to the data. Additional performance optimization methods such as query optimization methods are also part of the methodology. Search efficiency and relevance enhancement are provided by query rewriting, result filtering and ranking optimization. Most of the queries that are frequently asked are cached, whereas the complicated queries are optimized so that they can take less processing time. Also, query normalization is applied to recognize similar queries and enhance the rates of cache hits.

In order to test the effectiveness of proposed framework, a performance benchmarking and evaluation stage is carried out. There are synthetic datasets and actual datasets used to test the system, simulating different query heavy loads and usage patterns. Measures and analysis of performance metrics like query latency, throughput, cache hit ratio and resource utilization are done. The hybrid caching system and a baseline non-cached Lucene system are compared and performance gains are quantified. Scalability is also handled using distributed deployment strategy whereby the indexing and caching components are scaled out and across multiple nodes. Distributed caching will ensure that the cache information is replicated across nodes enhancing availability and fault tolerance. The load balancing methods are employed to distribute query requests equally within the system without causing bottlenecks and maintaining high performance when the system is working with a high load.

Lastly, the process of methodology includes monitoring and optimization algorithms. The monitoring of system performance is carried out with the continuous measures of the metric and logs, which will allow locating the points of leakage and inefficiency. On the insights, cache configurations, indexing strategy and query optimizations are dynamically changed to ensure optimum performance. This dynamic solution makes the system capable of adapting to the dynamism in the workloads and data trends. The approach to enterprise search system optimization presented in this study, combined with hybrid indexing and caching, offers a cohesive way to optimize the search systems. The framework provides great enhancements in query performance, scaling and reliability through a combination of EHCACHE and Apache Lucene, as well as the solid use of cache management and invalidation schemes.

Case Study:

This case study evaluates the implementation of a hybrid caching framework combining **Apache Lucene indexing** with **EHCACHE-based in-memory caching** in a large enterprise knowledge management system. The system is used by a multinational organization to search across **over 20 million documents**, including reports, emails, logs, and internal knowledge articles. The platform supports thousands of concurrent users, generating high volumes of repetitive and complex search queries.

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal

Initially, the system relied solely on Lucene-based indexing for search operations. While the indexing mechanism provided accurate and efficient retrieval, the system experienced **high query latency and increased CPU utilization** during peak usage. Repeated execution of frequently used queries resulted in redundant computations, leading to performance degradation and slower response times.

To address these challenges, a hybrid caching architecture was introduced. EHCACHE was integrated as an in-memory caching layer to store frequently accessed query results and intermediate search computations. The system was designed to first check the cache before executing queries on the Lucene index. If a cache hit occurred, results were returned instantly; otherwise, the query was processed and the results were stored in the cache for future use.

The implementation also included **cache invalidation strategies** to maintain data consistency. Time-based expiration (TTL), event-driven invalidation triggered by index updates, and selective cache eviction were used to ensure that cached results remained accurate. An LRU (Least Recently Used) eviction policy was applied to manage memory usage effectively.

The evaluation was conducted over a **3-month period**, comparing system performance before and after the implementation of hybrid caching. Key metrics analyzed included query latency, throughput, cache hit ratio, system resource utilization, and user experience.

Table 1 implementation of a hybrid caching framework

Metric	Before (Lucene Only)	After (Hybrid Cache + Lucene)	Improvement (%)
Average Query Latency (ms)	480	140	↓ 71%
Throughput (queries/sec)	900	2,100	↑ 133%
CPU Utilization (%)	78	52	↓ 33%
Cache Hit Ratio (%)	N/A	76	—
Memory Utilization (%)	60	68	↑ 13%
Query Error Rate (%)	3.2	1.1	↓ 66%
Peak Load Handling Capacity (%)	70	90	↑ 28%

Australian Journal of Cross-Disciplinary Innovation

Impact Factor: 16.4

ISSN-3746-757x (Online)

Acceptance Rate: below 5%

User Satisfaction Score (out of 10)	6.8	8.9	↑ 31%
-------------------------------------	-----	-----	-------

Analysis

The results clearly demonstrate the effectiveness of the hybrid caching approach in improving search performance. The **average query latency decreased by 71%**, significantly enhancing user experience by providing faster search results. The **throughput increased by 133%**, indicating the system's ability to handle a higher volume of queries efficiently.

The introduction of caching reduced CPU utilization by 33%, as fewer queries required full index processing. The achieved **cache hit ratio of 76%** highlights that a majority of queries were served directly from memory, reducing computational overhead. Although memory utilization increased slightly due to caching, the trade-off was justified by the substantial performance gains. The implementation of cache invalidation strategies ensured that data consistency was maintained, with minimal impact on accuracy. Additionally, the system demonstrated improved reliability, as indicated by the reduction in query error rates and enhanced peak load handling capacity.

Conclusion

In this paper, a hybrid framework of caching that has been applied to optimize the performance of enterprise searches was proposed through a combination of in-memory caching with indexing frameworks. By integrating effective indexing with smart caching techniques, the study overcame some of the major difficulties in traditional search system, especially, a long query turnaround time and unnecessary processing. The suggested solution, based on Apache Lucene as the solution to efficient indexing and EHCACHE as the solution to quick access to data in memory, contributed greatly to the overall response time of the queries and system performance. The implementation showed that, caching of heavily used queries has resulted in the reduction of the computational overhead and disk I/O operations and in effect the search performance has been made fast and more efficient. Also, with the integration of the cache invalidation techniques including the time based expiration, update according to events, and selective eviction, the data consistency between the cache version and the underlying index was guaranteed. The effectiveness of the approach was confirmed in the case study as significant improvements in latency, throughput and system resource utilization were achieved at a high accuracy and reliability. In general, the study demonstrates that hybrid caching is an effective and scalable approach to the current search solutions in enterprise systems. The proposed framework will allow organizations to manage large volumes of search workload and provide a high-quality user experience by optimizing their performance without jeopardizing data consistency.

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal

Australian Journal of Cross-Disciplinary Innovation

Impact Factor: 16.4

ISSN-3746-757x (Online)

Acceptance Rate: below 5%

Future Work

Although the suggested hybrid caching framework can make a considerable performance improvement, certain areas can be researched and improved. A potential solution is the combination of adaptive caching methods with machine learning algorithms to automatically foresee pattern of queries and optimize the utilization of caches. This will help to boost cache map strike rates and also lower the latency more, by actively caching perceptually relevant data. The other area in which I intend to explore in the future is use of distributed caching architectures that are to be used in support of multi-node large-scale settings. This would provide better scalability and fault resistance as systems would be able to support even more workload and still assure the same performance. Furthermore, a study can also be undertaken to optimize the procedure of cache synchronization to maintain consistency towards the nodes which are distributed without incurring a substantial overhead. There is also the investigation of more sophisticated methods of invalidating caches. More efficient mechanisms of intelligent invalidation where directly affected data can be identified and only important cache entries updated can also enhance efficiency and minimize the number of unnecessary cache invalidations that are unnecessary. It is also possible to incorporate the freshness of search results by adding in a real-time data streaming and incremental indexing. The possibility of integrating real-time analytics and search personalization can also be investigated in the future, whereby the caching approach is customized according to the user behavior patterns and preferences. This would allow more pertinent and quicker search queries, and enhance the user experience. Finally, with the increasing importance of cloud-native architectures, the adoption of serverless and cloud-based caching solutions can be investigated to improve scalability, cost efficiency, and ease of deployment. Such innovations will also enhance the functionality of enterprise search systems, and make them more adjustable, more efficient and robust when it comes to meeting the changing data and workload demands.

References

1. McCandless, M., Hatcher, E., & Gospodnetic, O. (2010). *Lucene in action* (2nd ed.). Manning Publications.
2. Bialecki, A., Muir, R., & Ingersoll, G. (2012). Apache Solr. In *Proceedings of the ACM SIGIR Conference* (pp. 1–2).
3. Shapira, B., Rokach, L., & Freilikhman, S. (2014). Facebook single sign-on for online services. *IEEE Internet Computing*, 18(2), 38–45.
4. Ehcache. (2022). *Ehcache documentation*. Retrieved from <https://www.ehcache.org>

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal

Australian Journal of Cross-Disciplinary Innovation

Impact Factor: 16.4

ISSN-3746-757x (Online)

Acceptance Rate: below 5%

5. Dean, J., & Barroso, L. A. (2013). The tail at scale. *Communications of the ACM*, 56(2), 74–80.
6. Kleppmann, M. (2017). *Designing data-intensive applications*. O'Reilly Media.
7. Hennessy, J. L., & Patterson, D. A. (2019). *Computer architecture: A quantitative approach* (6th ed.). Morgan Kaufmann.
8. Tanenbaum, A. S., & Van Steen, M. (2017). *Distributed systems: Principles and paradigms* (2nd ed.). Pearson.
9. Aggarwal, C. C. (2015). *Data mining: The textbook*. Springer.
10. Ghemawat, S., Gobiuff, H., & Leung, S. T. (2003). The Google file system. In *Proceedings of the ACM Symposium on Operating Systems Principles* (pp. 29–43).
11. Zaharia, M., Chowdhury, M., Franklin, M. J., et al. (2010). Spark: Cluster computing with working sets. In *Proceedings of the USENIX Conference* (pp. 1–9).
12. Pugh, W. (1990). Skip lists: A probabilistic alternative to balanced trees. *Communications of the ACM*, 33(6), 668–676.
13. Podlipnig, S., & Böszörményi, L. (2003). A survey of web cache replacement strategies. *ACM Computing Surveys*, 35(4), 374–398.
14. Cao, P., & Irani, S. (1997). Cost-aware WWW proxy caching algorithms. In *Proceedings of the USENIX Symposium* (pp. 193–206).
15. Elasticsearch. (2022). *Elasticsearch reference documentation*. Retrieved from <https://www.elastic.co>

Frequency: Yearly

Indexing: Google Scholar | DOAJ | ResearchGate

Peer Reviewed Refereed Journal