

Efficient distributed selective search

Yubin Kim¹ · Jamie Callan¹ · J. Shane Culpepper² ·
Alistair Moffat³

Received: 27 May 2016 / Accepted: 1 November 2016 / Published online: 25 November 2016
© Springer Science+Business Media New York 2016

Abstract Simulation and analysis have shown that *selective search* can reduce the cost of large-scale distributed information retrieval. By partitioning the collection into small *topical shards*, and then using a resource ranking algorithm to choose a subset of shards to search for each query, fewer postings are evaluated. In this paper we extend the study of selective search into new areas using a fine-grained simulation, examining the difference in efficiency when term-based and sample-based resource selection algorithms are used; measuring the effect of two policies for assigning index shards to machines; and exploring the benefits of index-spreading and mirroring as the number of deployed machines is varied. Results obtained for two large datasets and four large query logs confirm that selective search is significantly more efficient than conventional distributed search architectures and can handle higher query rates. Furthermore, we demonstrate that selective search can be tuned to avoid bottlenecks, and thus maximize usage of the underlying computer hardware.

Keywords Selective search · Distributed search · Load balancing · Efficiency

This is an extended version of work first presented as a short paper at the 2016 ACM-SIGIR International Conference on Research and Development in Information Retrieval (Kim et al. 2016b).

✉ Yubin Kim
yubink@cmu.edu

Jamie Callan
callan@cs.cmu.edu

J. Shane Culpepper
shane.culpepper@rmit.edu.au

Alistair Moffat
ammoffat@unimelb.edu.au

¹ Carnegie Mellon University, Pittsburgh, PA, USA

² RMIT University, Melbourne, Australia

³ The University of Melbourne, Melbourne, Australia

1 Introduction

A *selective search* architecture divides a document corpus or corpus tier into P topic-based partitions (*shards*), and assigns them to C processing *cores*, typically with $P \gg C \geq 1$. When a query arrives, a resource ranking algorithm (also known as “resource selection” or “shard ranking”) selects a small number of shards to be interrogated for that query. The search results from those shards are combined to produce the final results for the query. Selective search has been shown to provide similar effectiveness to exhaustive search when measuring early precision, and can provide additional efficiency-effectiveness trade-offs when working in low resource environments (Kulkarni and Callan 2010a, b; Kulkarni 2013; Aly et al. 2013; Kulkarni and Callan 2015).

Previous investigations argued that selective search is substantially more efficient than a typical distributed search architecture based on the number of postings processed when evaluating a small, single query stream. While this metric is suitable for comparing the work done between different architectures, it does not consider how work is divided across processors, or behavior when multiple query streams are evaluated in parallel. The traditional distributed search architecture using a random assignment of documents to shards tends to spread the workload evenly, and is relatively immune to bottlenecks. In contrast, a selective search architecture, which deliberately concentrates similar documents into a few index shards, might be vulnerable to uneven workloads, and hence leave processing resources idle. Selective search might also be more sensitive to tuning parameters.

Here we use an event-based simulator to investigate the efficiency of selective search, the details of which are discussed in Sect. 4. A simulator makes it possible to investigate a wider range of machine configurations than would be practical in a live system; in our case, we provide realistic measurements of query waiting times, query processing costs, query latency, system throughput, and hardware utilization under a query processing environment representative of a practical real-world implementation. Our investigation extends prior work by defining a more realistic experimental methodology for studying efficiency that employs similarly sized index shards, long query streams extracted from web search logs, and varying query arrival rates. In particular, we present a detailed study of the computational costs, load distribution, and throughput of selective search so as to address four research questions:

RQ 1 Is selective search more efficient than exhaustive search in a parallel query processing environment? (Sect. 5.1)

RQ 2 How does the choice of resource selection algorithm affect throughput and load distribution in selective search, and how can any imbalances originating from resource selection be overcome?(Sect. 5.2)

RQ 3 How do different methods of allocating shards to machines affect throughput and load distribution across machines? (Sect. 5.3)

RQ 4 Does selective search scale efficiently when adding more machines and/or shard replicas? (Sect. 5.4)

This paper includes and extends work initially presented in preliminary form in a short conference paper (Kim et al. 2016b), and also provides greater detail in regard to the experimental conditions, so as to make the experiments reproducible by others. Although the primary focus is on search efficiency, this work also describes simple improvements

that deliver improved search accuracy compared to prior results (Kulkarni 2013; Kulkarni and Callan 2015), meaning that the system we study is more similar to what would be used in practice. Experiments that report search efficiency include information about how queries are affected at the 50, 75 and 99% percentiles, and how much time a typical query spends in different parts of the system. Experiments with different resource selection algorithms show the average utilization of machines in environments discussed by prior research (Kulkarni 2013; Kulkarni and Callan 2015), and investigate the load imbalances that can arise. Experiments that investigate the policies for assigning shards to machines include information about the average loads on different machines for varying query arrival rates, the overall throughput for different shard assignment policies, and how log-based shard assignment compares to ten different instances of the more common random assignment policy. Finally, this paper includes an investigation of how selective search behaves as a system is scaled to a larger number of machines, which the preliminary work did not address.

Prior to our investigation, selective search efficiency was reported presuming that queries are evaluated sequentially. The investigation reported here provides a thorough study of the efficiency and load balancing characteristics of selective search in a parallel processing environment, to both broaden and deepen our understanding of this retrieval architecture.

2 Related work

Selective search integrates ideas and techniques from cluster-based retrieval, distributed retrieval, and federated search into a new architecture with distinct characteristics, some of which have not yet been explored in depth. We begin by describing these roots, and then summarize the state-of-the-art in selective search research.

2.1 Cluster-based retrieval

Cluster-based retrieval systems organize a corpus into hierarchical or flat clusters during indexing. When a query is received, the most appropriate clusters are selected using various methods. For example, cluster selection can be done by comparing the cosine similarity of the query to cluster centroids. Systems using small collections and clusters may return all of the documents in selected clusters; systems using larger clusters that contain many more documents than a user would browse rank the documents in the selected clusters and return just the highest-scoring documents. Early cluster-based retrieval systems aimed to improve the accuracy of search, but were unable to achieve consistent improvements in effectiveness (Croft 1980; Voorhees 1985; Griffiths et al. 1986; Willett 1988).

Typically cluster-based retrieval systems have many small clusters. For example, Can et al. (2004) used clusters with average sizes of 8–128 documents per cluster across five datasets, and Altinogvde et al. (2008) used average cluster sizes of 128–313 documents per cluster across three datasets. When clusters are small, many must be selected to maintain acceptable accuracy; Can et al. searched 10% of the clusters, a heuristic threshold that was also used in earlier investigations.

Clusters are stored in a single index, because so many must be searched for each query. Cluster membership information is stored in inverted list data structures, and postings are

grouped by cluster membership so that large portions of an inverted list may be skipped during query processing (Can et al. 2004; Altingovde et al. 2008). This architecture must bear the I/O costs of reading complete inverted lists, and the computational costs of processing them (albeit, efficiently). Can et al. (2004) note that storing each cluster in its own index would reduce the computational costs of processing long inverted lists, but incur prohibitive I/O costs (primarily disk seeks) due to the large number of clusters selected if the data collection cannot fit in memory.

2.2 Distributed retrieval

When the index becomes too large to search quickly with a single machine, the index is partitioned, and each partition is assigned to a distinct machine.

In *term-based partitioning*, each partial index is responsible for a non-overlapping subset of the terms in the vocabulary (Moffat et al. 2006; Lucchese et al. 2007; Zhang and Suel 2007; Cambazoglu et al. 2013). When the collection is searched, only indexes that contain the query terms are searched. Because queries are typically short, only a few indexes are required for each query, allowing multiple queries to be evaluated in parallel. This style of index has largely fallen out of favor because it is prone to load imbalances (Moffat et al. 2007). Cambazoglu et al. (2013) provide more details about term-based partitioning approaches.

In *document-based partitioning*, each partial index is responsible for a non-overlapping subset of the documents in the collection. There are two major approaches to creating the document subsets: tiering and sharding. *Tiering* creates partitions that have different priorities. Search begins at the top tier, and progresses to lower tiers only if necessary (Risvik et al. 2003; Baeza-Yates et al. 2009b; Cambazoglu et al. 2010; Brefeld et al. 2011). Tiers can be defined based on document characteristics such as geographic location, popularity, or assessed quality. The alternative approach, *sharding*, creates partitions that have the same priority and are searched in parallel (Tomasic and Garcia-Molina 1993; Cahoon et al. 2000; Puppin et al. 2006; Badue et al. 2007). Documents are usually assigned to shards randomly or in a round-robin approach. However, the assignment can also be based on document characteristics such as source (for example, URL).

Tiering and sharding are complementary methods that can be combined. For example, the corpus might be divided into tiers based on document popularity or authority, and then each tier divided into shards, with the shards distributed across a cluster of computers (Orlando et al. 2001; Barroso et al. 2003; Baeza-Yates et al. 2007, 2009a; Brefeld et al. 2011; Francès et al. 2014). A tiered document-partitioned index is the most common architecture for web search and other large-scale search tasks.

A variety of work has explored the efficiency of sharded search systems, covering topics including: reducing the communications and merging costs when large numbers of shards are searched (Cacheda et al. 2007); load balancing in mirrored systems (Macdonald et al. 2012; Freire et al. 2013); query shedding under high load to improve overall throughput (Broccolo et al. 2013); and query pruning to improve efficiency (Tonello et al. 2013). Other work focuses on addressing the load imbalances that arise when non-random shards are used, including the development of techniques for strategic assignment of index postings to shards, and strategic replication of frequently-accessed elements (Moffat et al. 2006, 2007). A common theme is that when a tier is searched, *all* of the tier's shards are searched.

2.3 Federated search

Sharded indexes and their search engines are a special case of *federated search systems*, which integrate heterogeneous search services (for example, *vertical* search engines) or search engines controlled by different organizations into a single service (Shokouhi and Si 2011). Usually the goal of federation is to send queries to as few of the underlying search services as possible, so a *resource selection* algorithm is used. Three types of resource selection have been proposed for federated search: term-, sample-, and classification-based algorithms.

In *term-based* methods, each search service is represented as a bag of words, with document ranking algorithms adapted to the task of ranking resources or services; GIOSS (Gravano et al. 1999), CORI (Callan 2000), and the query likelihood model (Si and Callan 2004a) are all examples of this approach. The simplest term-based resource selection algorithms are not very different from the cluster selection algorithms used for cluster-based retrieval. Algorithms developed specifically for resource ranking usually model the distribution of vocabulary across search services (Yuwono and Lee 1997; Hawking and Thistlewaite 1999; Callan et al. 1999; Aly et al. 2013). Term-based algorithms typically only support bag-of-words queries, but some also support corpus-level or cluster-level preferences, or Boolean constraints (Gravano et al. 1999; Xu and Croft 1999; Callan 2000; Liu and Croft 2004). A recent survey by Markov and Crestani (2014) provides a detailed analysis of several popular term-based methods.

Each search service can also be represented by a small *sample* of its documents. The samples from all search services are combined to create a *centralized sample index*, or *CSI*. When a query is received, the CSI is searched, and each top-ranked document found in the CSI is treated as a vote for the resource from which it was sampled. Many different methods for weighting votes from different resources have been described (Si and Callan 2003, 2004b, 2005; Shokouhi 2007; Paltoglou et al. 2008; Thomas and Shokouhi 2009; Kulkarni et al. 2012).

Classification-based algorithms represent each search service using a model combining various features and trained through supervised learning. Examples of features include the presence of specific words in the query, the scores of term-based and sample-based algorithms, and the similarity of the query to a resource-specific query log (Arguello et al. 2009; Kang et al. 2012).

Sample-based algorithms have been regarded as being a little more effective than term-based algorithms (Shokouhi and Si 2011); however, recently Aly et al. (2013) argued that the term-based *Taily* algorithm is more effective than the best sample-based algorithms. Term-based and sample-based algorithms are effective when the search engines are mostly homogeneous. Both types of algorithm are unsupervised, meaning that training data is not required. Supervised classification-based algorithms can be more effective than unsupervised methods; however, their main advantage is their ability to select among heterogeneous resources (for example, “vertical” search engines), and exploit a wide range of evidence.

Our focus in this paper is on the use of resource selection algorithms to select a few (of many) search engines for a particular query. However, some resource selection algorithms are very general and have been applied to a variety of other tasks, such as blog search, desktop search, and personal metasearch (Seo and Croft 2008; Elsas et al. 2008; Kim and Croft 2010; Thomas and Hawking 2009).

2.4 Selective search

Selective search is an architecture for large-scale search that combines ideas from cluster-based retrieval, document-partitioned distributed search, and federated search architectures (Kulkarni and Callan 2010a, b; Kulkarni 2013; Kulkarni and Callan 2015). The corpus is divided into *topic-based shards* that are stored in separate indexes and distributed across the processing resources. When a query arrives, a resource selection algorithm identifies a subset of shards that are most likely to contain many of the relevant documents. The selected shards are then searched and their answer lists merged to form an overall answer.

Puppin et al. (2006) and Cacheda et al. (2007) were among the first to study the combination of topic-based partitioning and resource selection to improve distributed search. They showed that partitioning a corpus into just a few large topics (11 and 17 topics, respectively) stored in separate indexes produced more efficient search than a traditional replicated distributed system. Kulkarni and Callan (2010a, b) increased the number of clusters (for example, 50–1000), placed greater emphasis on resource selection due to the larger number of index shards, and named the architecture *selective search*. They suggested that the desired number of topics is not an inherent property of a corpus, but instead a parameter to be set based upon efficiency considerations.

Classic cluster-based retrieval systems produce many small clusters. Selective search systems produce and search a smaller number of large topic-based clusters. For example, Kulkarni (2013) and Aly et al. (2013) used clusters that contained approximately 500,000 documents, and queries typically searched 3–5 clusters. This difference makes it practical to store each cluster in its own index, which has two important implications. First, selective search systems have much lower I/O and computational costs than cluster-based retrieval systems because they read from disk only a small fraction of each term's total postings. Second, index shards can be assigned to different machines, as is typical for distributed retrieval systems.

Previous studies showed that selective search systems and typical distributed retrieval architectures have similar accuracy, but that selective search systems are much more efficient because they search only a few index shards (Kulkarni and Callan 2010a, b; Aly et al. 2013; Kulkarni 2013; Kulkarni and Callan 2015). The topic-based indexes used for selective search are also more compatible with query processing optimizations such as WAND than are the randomly-partitioned indexes often used for distributed retrieval (Kim et al. 2016a; Kulkarni and Callan 2015). However, those studies determined the computational cost by counting the number of postings processed (Aly et al. 2013; Kulkarni and Callan 2015), or by measuring execution time on proof-of-concept implementations that were deployed on just one or two multi-core machines and processed queries sequentially (that is, just a single query active at any time) (Kulkarni 2013). Cacheda et al. (2007) compared the efficiency of a cluster-based retrieval system to a traditional replicated distributed system, but the cluster-based system had many fewer shards than a typical selective search system, and the experiments did not consider the cost of resource selection.

The fact that selective search systems search only a few shards for each query creates architectural choices that received little attention in prior research. When any single query will access only a few shards, it is practical for each processor to service query traffic for multiple shards. Kulkarni and Callan (2015) studied an environment in which dozens or hundreds of index shards were assigned to 16 processors. However it is an open question how many shards to assign to each processor, how to assign shards to processors, and how to balance computational loads when different index shards attract different amounts of query

traffic. In addition, it is unclear whether selective search architectures are prone to bottlenecks or load imbalances, especially when deployed across a larger number of machines.

2.5 Resource selection for selective search

Selective search uses resource selection algorithms similar to algorithms used for federated search. Our investigation makes use of the *Taily* and *Rank-S* resource selection mechanisms that were developed for selective search, thus we provide additional detail about these algorithms.

In *Rank-S*, the query is used to rank documents in the centralized sample index (CSI), which is a small, random sample of the total document collection. Document scores are decayed exponentially and then treated as votes for the shards the documents were sampled from Kulkarni et al. (2012). The exponentially-decayed vote of a document for its parent resource is computed as:

$$Vote(d) = Score_{CSI}(d) \times base^{-Rank_{CSI}(d)}$$

where $Score_{CSI}(d)$ and $Rank_{CSI}(d)$ are the document score and rank obtained by searching the CSI; and $base$ is a configurable parameter. The final score of a resource is $\sum_{i=1}^k Vote(d_i)$, the sum of the votes of the top- k documents retrieved from that resource. Resources with a total score above 0.0001 are then selected, as originally described by Kulkarni et al. (2012).

Taily assumes that the distribution of document scores for a single query term is approximated by a Gamma distribution. The allocation algorithm uses two parameters, n_c and v , where n_c is roughly the depth of the final ranked list desired, and v is the number of documents in the top n_c that a resource must be estimated as contributing in order to be selected. Term scores are calculated from simple corpus statistics and fitted to a Gamma distribution for each shard-term pair. *Taily*'s resource selection database records these fitted Gamma distributions for each term t in resource i , describing the term's score distribution in that shard. Gamma distributions are represented by two parameters, the shape parameter k_i^t and the scale parameter θ_i^t . At query time, the cumulative distribution function of the Gamma distribution is used to estimate the number of documents from each resource that will have a score above a threshold derived from n_c . Each resource that provides v or more documents is selected (Aly et al. 2013).

When a query is received, *Taily* looks up two floating point numbers (the parameters of the fitted Gamma distribution) for each index shard per query term, whereas *Rank-S* must retrieve an inverted list for each query term. *Taily*'s computational costs are linear in the number of index shards $|S|$, and nearly identical for each query of length $|q|$. The computational costs for *Rank-S* vary depending on the document frequency (df) of each query term in the centralized sample index. The *Rank-S* approach is more efficient only when $df_t < |S|$ for query term t . That is, for most applications *Taily* is faster to compute than *Rank-S* (Aly et al. 2013).

3 Simulation model

Our goal is to study the computational characteristics of selective search in more realistic and varied distributed processing environments than have been reported previously. However, acquiring, configuring, and producing repeatable experiments with a variety of

real hardware configurations is expensive and time-consuming (Webber and Moffat 2005). Simulation is recognized as a viable way of modeling computational costs for complex systems, and simulation has played an important role in a wide range of IR-related investigations (Burkowski 1990; Cahoon et al. 2000; Ribeiro-Neto and Barbosa 1998; Cacheda et al. 2007; Tonello et al. 2013).

3.1 Simulation parameters

We developed a selective search simulator¹ based on DESMO-J² a discrete event modeling framework. The simulator has the benefit of allowing us to model a range of hardware configurations, and provides precise estimates of a broad suite of performance indicators. The implementation models a selective search system that incorporates a cluster of multi-core machines, and mimics parallel query execution across those machines. Figure 1 describes the computational model embedded in the simulator and Table 1 lists the quantities that are manipulated. The hardware is assumed to consist of M machines, with the i th of those, machine m_i , providing c_i CPU cores ($c_i = 8$ throughout the paper). Machines may be configured to act as a *broker*, a *searcher*, or to handle both roles.

A *broker* machine holds a copy of the resource selection database and performs two tasks: resource selection, and result merging. For resource selection, the machine has access to a shared *central query queue*, from which it extracts incoming queries, determines which shards need to be searched, and then assigns shard search tasks to other machines. Each broker machine also has a job queue for pending *result merge* processes. This queue contains results returned by the searcher machines, now waiting to be merged to produce a final result list for some query.

A machine m_i is a *searcher* if it is allocated $p_i > 0$ shards. A searcher also has a job queue that holds *shard search* requests pertinent to the shards hosted on that machine. Each of the available cores on the machine can access any of the shards assigned to the machine, and hence can respond to any request in that machine's search queue. When a search job is finished, the result is returned to the result merge queue of the originating broker. The assignment of shards and copies of the resource selection database to machines is assumed to be fixed at indexing time, and machines cannot access shards that are not hosted locally. A key factor for success is thus the manner in which the P shards are partitioned across the M machines.

3.2 Selective search process

Algorithm 1 describes the actions that take place in each of the machines. First, if the machine is a broker, the local result merge queue is checked for queries for which all shard searches are complete, and merged output lists are generated if any queries can now be finalized. Otherwise, if the machine is a searcher, the local shard queue is checked to see if there are any searches pending; if so, the next one is actioned, and the results directed to the merge queue for the machine that acted as broker for that query, which might be the same machine. A process on machine m_i can search any shard assigned to that machine.

If neither of these two activities are required, and if the machine is a broker, the next query (if one exists) is taken from the central queue and resource selection carried out. The result is a list of shards to be searched in order to resolve the query; that list is mapped to a

¹ <http://boston.lti.cs.cmu.edu/appendices/jir17-yubink/>.

² <http://desmoj.sourceforge.net/>.

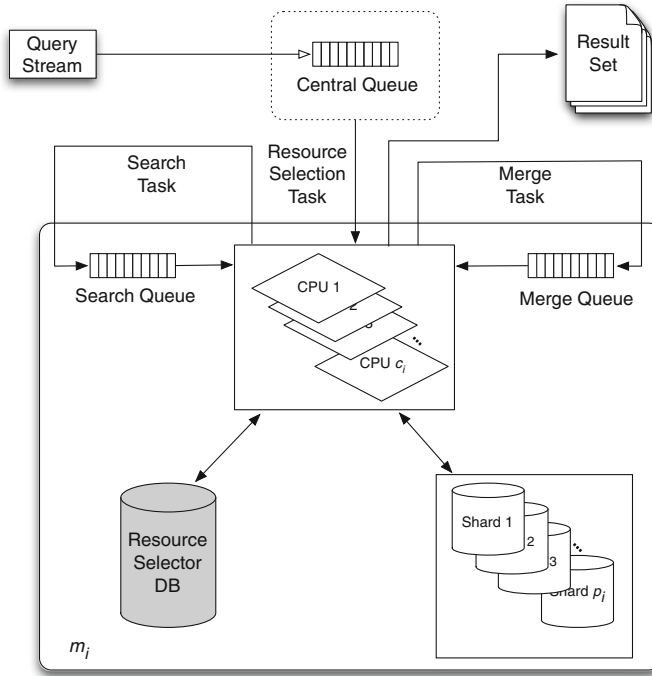


Fig. 1 Architecture of the selective distributed search system. The i th of the M machines m_i has c_i cores, each of which can be used for resource selection and result merging, or for shard search across any of the p_i shards allocated to this machine. Only a defined subset of the machines are able to perform resource selection

Table 1 Simulation parameters

M	Number of machines; m_i is the i th of these
c_i	Number of cores on m_i ; the default is $c_i = 8$
C	Total number of cores, $\sum_{i=1}^M c_i$
p_i	Number of shards assigned to m_i
P	Total number of shards. When each shard is assigned to just one machine, $P = \sum_{i=1}^M p_i$
B	Number of broker machines
S	Number of searcher machines
T	Query arrival rate described by an exponential distribution with mean $1 / T$
t_s	Seek plus latency access time, msec/postings list, $t_s = 4$ throughout
t_p	Processing cost, msec/posting, $t_p = 9 \times 10^{-4}$ throughout
t_m	Merging cost, msec/item, $t_m = 5 \times 10^{-5}$ throughout

set of machine identifiers, and the query q is added to the shard search queues for those machines. Query completion is prioritized over query initiation, minimizing the processing time for each query, and ensuring that no query has an infinite time to completion.

Algorithm 1 – Processing loop for each core on machine m_i .

```

while forever do
  if  $isBroker(m_i)$  and  $|mergequeue_i| > 0$  and
    all shard responses have been received for some query  $q$  then
    remove those responses from  $mergequeue_i$ 
    finalize the output for query  $q$  and construct a
    document ranking
  else if  $isSearcher(m_i)$  and  $|searchqueue_i| > 0$  then
    remove a query request  $(q, p, b)$  from  $searchqueue_i$ 
    perform a shard search for query  $q$  against shard  $p$ 
    append the results of the search to  $mergequeue_b$ 
  else if  $isBroker(m_i)$  and  $|centralqueue| > 0$  then
    remove a query  $q$  from  $centralqueue$ 
    perform resource selection for  $q$ 
    for each partition  $p$  to be searched for  $q$  do
      determine the machine  $m_h$  that is host for  $p$ 
      append  $(q, p, i)$  as a search request to  $searchqueue_h$ 
endwhile

```

The simulation assumes that queries arrive at the central queue at random intervals determined by an exponential distribution governed by a mean query arrival rate T . The number of machines permitted to host broker processes may be less than M , the total number of machines, but is always at least one. Query processing costs at the shards are computed based on the number of postings read from disk, plus an overhead cost to account for initial latency for a disk seek. The number of postings processed and system response times are known to have a strong correlation (Macdonald et al. 2012). While more accurate methods exist (Wu and Fang 2014), the main advantage comes from correctly estimating the postings pruned from the *total* postings of a query.

A postings list of ℓ postings is thus presumed to require $t_s + \ell \cdot t_p$ milliseconds, where $t_s = 4$ milliseconds³ and $t_p = 9 \times 10^{-4}$ milliseconds per posting. The processing rate—around a million postings per second—is based on measurement of the cost of handling posting lists in the open source Indri search engine (<http://lemurproject.org/>) on a machine with a 2.44 GHz CPU, and encompasses I/O costs as well as similarity calculation costs.

These parameters can be varied to explore different hardware architectures. For example, t_s can be set to zero and t_p can be reduced to investigate behavior when the index is stored in RAM or on a solid state device. Selective adjustment of t_p can also be used to mimic caching of postings lists. For example, when the experiments in this paper were done with $t_s = 0$ and $t_p = 9 \times 10^{-5}$ milliseconds per posting, to mimic an architecture with indexes stored in RAM and a hardware/software architecture that processed inverted lists an order of magnitude faster, queries were processed more quickly, but the relationships between different types of systems were unchanged.

3.3 Resource selection and result merging

Resource selection costs differ according to the approach used. For sample-based algorithms such as Rank-S, the cost is dominated by the need to process postings from the central sample index (CSI). For these approaches, the same computational model is used as for shard search, and a cost of $t_s + \ell \cdot t_p$ milliseconds is assumed for a list of ℓ postings. On

³ <http://www.anandtech.com/show/3636/western-digitals-new-velociraptor-vr200m-10k-rpm-at-450gb-and-600gb>, accessed 29/10/14.

the other hand, term-based algorithms such as Taily process statistics from each shard that contains a query term. The cost for term-based approaches is thus equivalent to processing a posting list of length equal to the number of shards that contain the query term, which is always less than or equal to P , the number of shards.

Result merging requires network transfer if the results are returned to a broker that is not located within the same machine. This requires transferring of up to k $\langle doc - id, score \rangle$ results from each shard searched, where k is either fixed on a system-wide basis, or is determined as part of the resource selection step. Network messages are also generated when brokers request searches for shards that are not stored within the same machine. To ensure that the simulation was accurate, the cost of network communication over a Gigabit switched network was modeled as described by Cacheda et al. (2007). The cost of merging was measured on the same 2.44 GHz machine, and an allocation of $t_m = 5 \times 10^{-5}$ milliseconds per document was found to be appropriate.

3.4 System output

The simulation models a system that returns the top-ranked 1,000 documents, thereby supporting applications such as learning to rank algorithms, text-mining applications, and TREC evaluations. Cacheda et al. (2007) showed that when shards are formed *randomly*, only a small number of documents need to be returned from each shard for the true top- k documents to be returned with a high probability. Therefore, in the exhaustive search system used as a baseline, each shard only returns the number of documents that results in a 10^{-5} probability of missing a result in the top 1,000. For example, this equates to returning 102 documents per shard for a 16 shard configuration and 19 documents per shard for 512 shards. The assumption that documents are distributed randomly across shards does not apply to the topical indexes used by selective search; clustering concentrates similar documents in a small number of shards. Thus, Cacheda's technique cannot be used with selective search and each search shard selected returns the full $k = 1,000$ documents. Since the exhaustive search baseline accesses all shards, whereas selective search typically accesses 3–5 shards, the number of documents that must be merged by the two architectures remains roughly comparable. In our experiments, the total number of documents merged by the two architectures varied between 1600 and 9700, depending upon the number of machines (exhaustive search) and query (selective search). Generally, exhaustive search merges fewer documents than selective search in configurations with fewer randomly-assigned shards, and more documents than selective search in configurations with more random shards.

Overall, the simulator takes as input a list of queries, the resource selection cost for each query, the shards to be searched for the query, and the search cost for each shard. The cost is described by the sum of the lengths of the posting lists retrieved for all query terms for the shards specified by the resource selection method. The simulator then converts these posting list costs into “simulator milliseconds”. The overall elapsed time required to process each query is taken as the difference between the arrival time of that query in the central queue, and the moment at which all processing of that query is completed. This cost includes time spent waiting in queues, network delays, and computation time. The median end-to-end elapsed query processing time is used as an aggregate measure of query latency. The median was preferred over the mean because it is less affected by the very small number of slow queries, but the conclusions of the paper remain the same with both metrics. In addition to the median latency, we also show the latency distributions in some

key cases. The primary variable in the simulator is the query arrival rate, which determines the load in the system, and hence the extent to which query response times are affected by queuing delays.

The load on each simulated machine is also tracked. The fraction of available processors utilized is measured at simulation event boundaries and summarized in a time-weighted average. This statistic is used to evaluate the evenness of load across the modeled hardware setup.

3.5 Other factors

The simulator models the main components of distributed and selective search architectures. However, in order to manage the complexity of the investigation (and hence this paper), it does not include every possible optimization for large-scale search, for example: dynamic pruning during query evaluation; postings list caching; and other network and connectivity arrangements. Indeed, it is unclear how to simulate some of these optimizations, for example dynamic pruning techniques, without fully implementing important parts of the search engine. What we can be assured of is that results presented in this paper are a reliable upper bound on the total processing cost for each method, because all postings in the selected shards contribute to the final computation. Actual processing costs in an optimized system would be somewhat lower than what we report.

As one step towards understanding these interactions, we note that recent work by the authors has demonstrated that the WAND dynamic pruning technique (Broder et al. 2003) retains its advantages when searching within topical shards; indeed, WAND was found to be somewhat *more* efficient on topic-based shards than randomly-assigned shards, meaning that selective search has a somewhat larger advantage over exhaustive search in a more optimized system (Kim et al. 2016a). Investigations into the effects of other optimizations typical of large-scale search environments would be a welcome addition to the literature, but are outside the scope of this paper.

4 Experimental methodology

The simulator was applied to two large experimental collections. This section describes those data resources, including a revised partitioning scheme that constructs approximately equal-sized sized partitions; and gives details of the effectiveness baselines we compare against.

4.1 Document collections

Two web datasets, ClueWeb09-A English (abbreviated to ClueWeb09 in all of the tables) and Gov2 were used in the experimentation. These are large collections, covering more than 500 million and 25 million web pages respectively. The selective search configurations applied to them are derived from the shard definitions for ClueWeb09 and Gov2 generated as part of the investigation carried out by Kulkarni (2013).

One issue with the ClueWeb09 shard map produced by Kulkarni is that the shard sizes are not uniform. Figure 2 plots the distribution of shard sizes in decreasing order and reveals a large skew. The largest shard contained 7.8 M documents, 12 times the average size of 620 k documents per shard. This imbalance is problematic for two reasons. Firstly,

there is a moderate correlation between the size of the shard and the frequency at which it is selected for querying (Pearson correlation of 0.5). That is, the larger the shard, the more likely it is to be selected for querying, which is a well-known bias of many resource selection algorithms. Secondly, when they do get selected, large shards typically take longer to process queries, because of their size. In combination, these two effects produce an imbalance in computation load which is not present in the Gov2 shards. To address the imbalance of shard sizes in ClueWeb09, the largest shards in the original shard partitioning of Kulkarni were divided; split via a random assignment if the resulting sub-shards would be closer to the average shard size (by number of documents) than the original one. Note that while smaller-than-average shards exist, they do not need to be explicitly combined; instead, many small shards can be assigned to the same machine to combine their load. A total of 51 shards were split into two or more smaller shards, resulting in an increase of 77 shards, for a total of 884.

A range of other statistics for the two datasets are listed in Table 2.

4.2 Resource selection parameter settings

As noted above, two resource selection techniques were used in our experiments: the sample-based Rank-S method (Kulkarni et al. 2012); and the term-based Taily approach (Aly et al. 2013). Both require that values be set for two system parameters. For Rank-S the two key parameters are the size of the document sample for the centralized index (the CSI), and the quantity *base* used for the exponential discounting. Taily requires values for *n*, conceptually the depth of the result list used for estimating relevant document counts; and for *v*, the score cut-off. The parameters used in our experiments are as recommended by Aly et al. (2013) and Kulkarni (2013), and are summarized in Table 3.

4.3 Confirming retrieval effectiveness

Tables 4 and 5 list the overall effectiveness of selective search using the default parameters, spanning multiple years of the TREC Terabyte and Web Tracks. Note that these effectiveness results are independent of the simulation process used to measure efficiency, in that there is no influence on them of the number of machines, and of how shards are assigned to machines—once the assignment of documents to shards has been completed, and the resource selection parameters determined, selective search generates the same final ranked list regardless of machines or shard allocations. The effectiveness reported for exhaustive search was achieved using structured queries and the sequential dependency

Fig. 2 Sizes of ClueWeb09 shards used in experiments of Kulkarni (2013)




Table 2 Datasets and shards used in the experiments, with “M” standing for million, “B” standing for billion, and document lengths measured in words

Dataset	# of docs.	# of words	Vocab. size	Avg. doc. length	# of shards
Gov2	25M	24B	39M	949	50
ClueWeb09	504M	381B	1,226M	757	884

Table 3 Resource selection parameter settings, based on TREC Terabyte Track and TREC Web Track queries for which relevance judgments are available, as developed by Kulkarni et al. (2012) and Aly et al. (2013). Those investigations used a variety of CSI sizes for Gov2, ranging from 0.5 to 4%. We use 1%, for consistency with ClueWeb09. The final column lists the average number of shards per query selected by the listed parameter settings

Dataset	Algorithm	Parameters	Avg. shards
Gov2	Taily	$n = 400, v = 50$	3.0
	Rank-S	$CSI = 1\%, base = 3$	4.7
ClueWeb09	Taily	$n = 400, v = 50$	3.3
	Rank-S	$CSI = 1\%, base = 5$	4.2

model (SDM) (Metzler and Croft 2005), removing spam documents at a Waterloo Fusion spam score threshold of 50%, similar to Aly et al. (2013). The selective search runs used a similar set-up whenever possible. For Rank-S, the CSI was created using a random sample of the document collection and was searched using SDM queries. For both Taily and Rank-S, the selected shards were searched using SDM queries and retrieval scores were calculated using global term frequency data; this produces scores that are comparable across shards, thus simplifying the result list merging process. The final result list was filtered for spam. Note that the spam filtering was only performed for the effectiveness results in order to generate results comparable to prior work and was not done for the efficiency experiments. These arrangements give rise to effectiveness better than was attained by Kulkarni (2013), and comparable to the levels reported by Aly et al. (2013).

Results for the four Web Track query sets are listed separately rather than combined across query sets, in order to distinguish the effects of different tasks on selective search. In the results reported by Kulkarni (2013) and Aly et al. (2013), effectiveness was averaged

Table 4 Effectiveness of selective search with various resource selection algorithms on the Gov2 dataset, using three available pairings of queries and relevance judgments

Algorithm	Terabyte Track 2004		Terabyte Track 2005		Terabyte Track 2006	
	P@10	NDCG@30	P@10	NDCG@30	P@10	NDCG@30
Baseline	0.56	0.43	0.62	0.49	0.57	0.48
Rank-S	0.57	0.42	0.59	0.45	0.54	0.44
Taily	0.51	0.37	0.48	0.35	0.50	0.42
Oracle	0.64	0.47	0.64	0.51	0.61	0.51

The “Baseline” method is an exhaustive search of all shards. The “Oracle” shard ranking assumes that the most useful four shards are selected for each query

Table 5 As for Table 4, but for the ClueWeb09 dataset, using four available pairings of queries and relevance judgments

Algorithm	Web Track 2009		Web Track 2010		Web Track 2011		Web Track 2012	
	P@10	NDCG@30	P@10	NDCG@30	P@10	NDCG@30	P@10	NDCG@30
Baseline	0.30	0.21	0.27	0.19	0.36	0.28	0.27	0.15
Rank-S	0.28	0.19	0.32	0.20	0.30	0.21	0.25	0.14
Taily	0.28	0.18	0.31	0.20	0.24	0.16	0.23	0.13
Oracle	0.37	0.25	0.45	0.32	0.42	0.35	0.33	0.20

across the 2009 and 2010 Web Track query sets; unpacking them as we have done here reveals that selective search is not as uniformly competitive as previously reported, particularly for the 2009 queries. On the other hand, on the 2010 queries a partitioned index and good resource selection are more effective than exhaustive search, supporting an observation originally made by Powell et al. (2000). When the 2010 results are averaged with the 2009 results, selective and exhaustive search have comparable performance. While the accuracy results are reported with our modified shard maps, these relationships also occur when using the original shard maps. Selective search is also less effective than exhaustive search across the two newer query sets from 2011 and 2012.

The effectiveness results reported in Tables 4 and 5 suggest that the parameters used for Taily and Rank-S might not be optimal. To address that concern, we performed a parameter sweep to check for better settings, but were unable to identify parameters that both yielded accuracy comparable to exhaustive search and also searched a moderate number of shards. We also experimented with a clairvoyant “oracle” shard ranking, in which the four shards containing the greatest number of relevant documents were presumed to always be selected, a number of shards similar to the average searched by Rank-S and Taily. Those hypothetical results are shown in the final row of each of Tables 4 and 5. The substantial gap between the oracle, and Rank-S and Taily, reinforces that the accuracy of selective search depends heavily on the resource selection process, and makes it clear that further algorithmic improvements may be possible. However, our emphasis here is on efficiency, and we leave that prospect for future work.

4.4 Query streams

A key purpose of the simulation is measurement of the performance of selective search under realistic query loads, including determination of the point at which each configuration saturates. For these experiments a much larger query log is needed than the Web Track and Terabyte Track query sets used in the effectiveness validation shown in Tables 4 and 5. The main experiments that are reported in Sect. 5 make use of the AOL query log and the TREC Million Query Track query set. These query logs are from a different time period than the two collections, but we used the queries (only) to measure efficiency and not effectiveness, so the discrepancy in their temporal coverage is not a concern.

In order to simulate a live query stream, the AOL log was sorted by timestamp, and deduplicated to only contain unique queries. Deduplication has the effect of simulating a large answer cache, and more closely reflects what would happen in a production server. For ClueWeb09, the first 1,000 queries were used as training data, to set configuration parameters such as the number of brokers (Sect. 5.2), and assignments of shards to

machines (Sect. 5.3). The next 10,000 queries were used for testing. Together, the queries cover a time period of 2006-03-01 00:01:03 to 2006-03-01 01:31:45. For Gov2, the timestamp-sorted AOL query stream was further filtered selecting only queries that had at least one .gov-domain result click recorded. The first 1000 queries were again used as training data, and the next 10,000 queries used for testing, covering the time period 2006-03-01 00:01:08 to 2006-03-01 20:09:09. Two further test query sets were also extracted from the AOL log: 10,000 queries starting from 2006-03-08, one week after the main query stream (Test_W); and another 10,000 queries commencing 2006-04-01, one month after the main query stream (Test_M). These query sequences were used in the experiments described in Sect. 5.3.

Two final query streams were created from the TREC Million Query Track (MQT). The MQT queries have no timestamps and were used in the order they appear in the files provided by NIST. For Gov2, the first 1000 queries of the 2007 query set were used for training, and 10,000 queries from the 2008 query set were used for testing. For ClueWeb09, both the 1000 training and 10,000 testing queries were extracted from the TREC 2009 MQT sequence.

In total, 84,000 queries were used at various stages of the evaluation, split across the categories summarized in Table 6.

5 Experimental results

We now use the query execution simulator to investigate the efficiency and load characteristics of the selective search architecture under a variety of conditions. The first set of experiments evaluates an environment similar to the one considered by Kulkarni (2013), but allows parallel execution of queries and tasks (Sect. 5.1). The second suite of experiments explores resource selection costs, resource selection algorithms, and how to overcome computational load imbalances arising from resource selection (Sect. 5.2). The third round of experiments investigate the load and throughput effects of two different policies for assigning shards to machines (Sect. 5.3). The last experiment compares indexing and mirroring strategies when using additional computational resources (Sect. 5.4). We reiterate that retrieval effectiveness remains constant, as reported in Tables 4 and 5, and that the simulation is used to determine execution times only, based on long streams of actual queries for which relevance judgments are not available.

Table 6 Sizes of query logs

Dataset	Log	Train	Test	Test_W	Test_M
Gov2	AOL_G	1K	10K	10K	10K
Gov2	MQT_G	1K	10K		
ClueWeb09	AOL_W	1K	10K	10K	10K
ClueWeb09	MQT_W	1K	10K		

The “Train” column indicates queries used to set parameters; the “Test” column queries used for reporting results. The “ Test_W ” and “ Test_M ” queries were sampled starting one week and one month after the Train queries, respectively

5.1 Selective search efficiency

Research Question 1 asks *Is selective search more efficient than exhaustive search in a parallel query processing environment?* We start by comparing the two architectures in small-scale environments similar to those examined by Kulkarni (2013) and Kulkarni and Callan (2010a). The test collections were divided into topical shards and the shards randomly distributed across all machines, with each machine receiving the same number of shards. The same collections were also used to build an exhaustive search baseline by constructing C (the total number of cores) evenly sized shards via a random assignment of documents, and then allocating one shard per core to each machine; $c_i = 8$ in our experiments, thus eight shards were assigned to each machine. This allows exhaustive search to make use of all cores for every query and hence maximize throughput. In both configurations, all machines accepted search tasks, but only one machine ($B = 1$) accepted broker tasks, so as to emulate previous arrangements as closely as possible. Table 7 summarizes the average number of shards selected by the resource selection algorithms with parameters as discussed in Sect. 4.

The high variance in the number of ClueWeb09 shards selected by Taily for the AOL query set is because the distribution of Taily shard scores levels off quickly and becomes flat, which makes it difficult for Taily to distinguish between shards. The AOL query set used with Gov2 does not experience the same issues; these queries are cleaner due to the filtering process that was applied. Even with these variances affecting behavior, selective search examines only a small fraction of the shards, and produces significant efficiency gains regardless of the resource selection method deployed.

Figure 3 shows the simulated throughput of two selective search variants, compared to exhaustive search. The vertical axis shows the median time to process a single query, plotted as a function of the query arrival rate on the horizontal axis. Alternative summary metrics such as the mean or the 95% percentile processing times were also explored, and produced similar patterns of behavior. The two frames correspond to the two collections, with each curve representing one combination of query set and processing regime. The right-hand end of each plotted curve is truncated at a point at which the system configuration is approaching saturation, defined as occurring when the median processing time for queries exceeds twice the median processing time for queries in a corresponding unloaded system (at an arrival rate of 0.01 queries per second, at the left-hand end of each curve). Query arrival rates were chosen dynamically for each curve, so that the region of greatest gradient can be clearly distinguished. Configurations in which the curve is lower have

Table 7 Average number of shards selected by Taily and Rank-S for the two Test logs, measured using the same system configurations as are depicted in Fig. 3

Queries	Res. sel.	Gov2		ClueWeb09	
		avg.	sddev.	avg.	sddev.
MQT	Taily	2.5	1.4	3.6	2.9
	Rank-S	4.2	1.8	4.4	1.7
AOL	Taily	2.9	1.6	11.9	31.3
	Rank-S	4.6	1.9	4.2	1.7

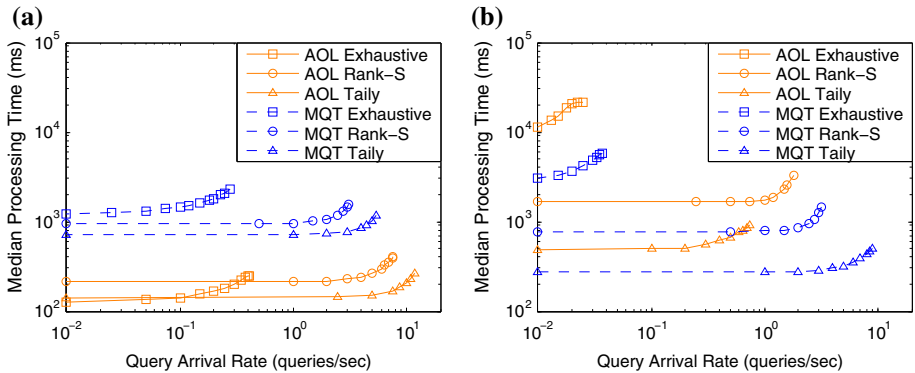


Fig. 3 Exhaustive search and selective search using ClueWeb09 and Gov2 datasets and the AOL_W Test, MQT_W Test, AOL_G Test, and MQT_G Test query sets, which are shown as AOL and MQT in the labels for brevity. The shards were randomly assigned to each machine. In these experiments, $M = 2$, $B = 1$, and $S = 2$. Note that both scales are logarithmic

Table 8 Dispersion of per-query processing times for selected configurations, at the last point on each of the corresponding curves in Fig. 3

Collection	Queries	Method	$T(q/s)$	Percentile			Mean	Max
				50%	75%	99%		
Gov2	MQT	Exhaustive	0.28	2,268	4,409	11,676	2,936	20,868
		Rank-S	3.20	1,632	2,651	6,206	1,883	10,229
		Taily	5.50	1,189	1,901	4,579	1,366	8,145
	AOL	Exhaustive	0.42	250	1,126	5,016	768	9,135
		Rank-S	7.50	397	865	2,606	590	5,352
		Taily	12.00	265	612	1,868	407	4,329
ClueWeb09	MQT	Exhaustive	0.04	5,819	26,510	110,232	17,784	206,595
		Rank-S	3.20	1,468	3,537	13,082	2,516	32,888
		Taily	9.00	510	1,012	3,340	719	6,709
	AOL	Exhaustive	0.03	21,736	50,525	172,978	34,019	427,116
		Rank-S	1.80	3,263	6,220	18,351	4,208	76,103
		Taily	0.75	902	3,869	21,350	3,032	46,770

Columns display the 50th, 75th, 99th percentiles, the mean, and the maximum value of individual query processing times, where time is measured in simulation milliseconds. In these experiments, $M = 2$. Recall that the simulation was tuned to model the costs associated with Indri, a research search engine; a commercial system would be likely to have lower processing times

lower latency under light load; configurations with elbows that are further to the right require fewer resources per query, and attain higher throughput rates when the system is approaching saturation.


To describe the variation of query processing times, Table 8 lists the 50th, 75th and 99th percentile, mean, and max values of query processing time for several configurations of selective search. The values reported are all taken at the last points in the plotted curves of Fig. 3—the point where each system is deemed to have approached saturation. Table 9

Table 9 Breakdown of elapsed query time, in percentages, at the last point on each of the corresponding curves in Fig. 3

Collection	Queries	Method	T (q/s)	Central queue (%)	Network (%)	Internal queues (%)	Resource selection (%)	Merge (%)
Gov2	MQT	Exhaustive	0.28	21.88	0.06	14.91	–	0.09
		Rank-S	3.20	32.16	<0.01	3.35	13.38	0.01
		Taily	5.50	30.38	<0.01	3.44	1.58	0.01
	AOL	Exhaustive	0.42	21.41	0.07	16.92	–	0.01
		Rank-S	7.50	25.67	0.01	2.58	14.67	0.04
		Taily	12.00	20.55	0.01	2.68	2.50	0.03
ClueWeb09	MQT	Exhaustive	0.04	31.57	<0.01	25.24	–	<0.01
		Rank-S	3.20	25.55	<0.01	1.72	53.44	0.01
		Taily	9.00	26.36	0.01	5.10	1.66	0.02
	AOL	Exhaustive	0.03	30.86	<0.01	31.08	–	<0.01
		Rank-S	1.80	22.40	<0.01	1.93	59.67	<0.01
		Taily	0.75	36.77	<0.01	7.54	0.46	0.02

Central queue indicates the percentage of time spent in the central query queue; *network* indicates the delays caused by network congestion; *internal queues* indicates time spent in system queues such as a machine’s shard search queue or merge queue; *resource selection* indicates time spent performing resource selection; and *merge* is time spent performing the final merge operation. The balance was spent in shard search

Fig. 4 Average utilization of machines using Rank-S, the ClueWeb09 dataset, the AOL_w Test queries, $M = 2$, $B = 1$ and $S = 2$. Each point is the mean over 10 sequences of 1000 queries; *error bars* represent 95% confidence intervals



breaks down the average costs at the same set of system saturation points. As expected of a system under load, queries spend a significant fraction of their time in the central query queue, and in queues internal to each system; network and merging costs are minimal. Note the differences in the time spent in resource selection between Rank-S and Taily, particularly on ClueWeb09. This is what causes the higher latency for Rank-S systems, evident in Fig. 3, and is as expected—sample-based resource selection methods are more expensive than term-based methods. This topic is explored further in Sect. 5.2.

Selective search outperforms exhaustive search by a factor of more than ten on the ClueWeb09 dataset (Fig. 3b), because only a small fraction of the 884 shards are searched for each query. Query latency is also lower in selective search, despite the two-step process of resource selection followed by shard search. This is due to fewer resources being used

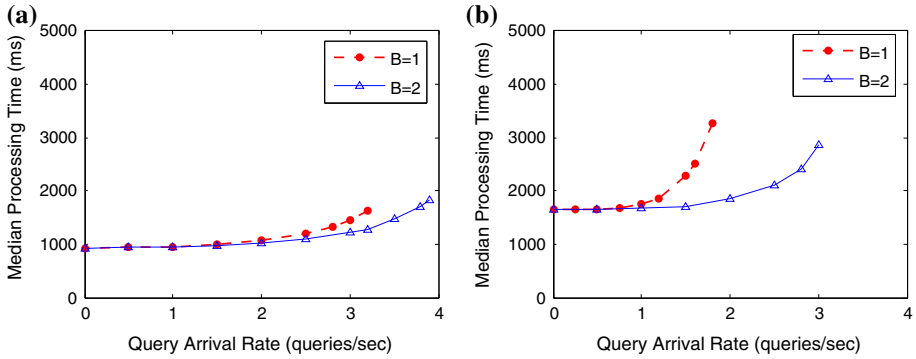


Fig. 5 Varying the resources assigned to broker processes using Rank-S, with $M = 2$, and random shard assignment


by selective search, and the fact that at low query loads, latency is largely determined by the slowest partition that is polled. Topical shards are smaller than random shards, thus they can be searched more quickly. A larger fraction of the fifty possible shards are searched in the Gov2 collection (Fig. 3a), so the performance improvement is not as large. Even so, selective search of Gov2 handles four to five times the rate of queries as exhaustive search before saturating.

Taily has better throughput and latency than Rank-S for the majority of settings tested. The only exception is for the AOL_W Test queries on ClueWeb09, where the larger number of shards searched by Taily eclipses the lower resource selection costs, resulting in better at-load throughput for Rank-S. More broadly, selective search delivers markedly better performance characteristics than exhaustive search in all of the configurations investigated, extending the findings of Kulkarni (2013) and Kulkarni and Callan (2010a) that selective search improves the efficiency of parallel query processing.

5.2 Resource allocation

The difference between Rank-S and Taily is a direct consequence of the two approaches to resource selection (Aly et al. 2013). In response, we turn to Research Question 2: *How does the choice of resource selection algorithm affect throughput and load distribution in selective search, and how can any imbalances originating from resource selection be overcome?*

Fig. 6 Popularity of the most-frequently accessed 10% of ClueWeb09 shards for the two training query sets, as determined by Rank-S. The vertical axis indicates the total number of times a shard is selected by that query set



In the experiments illustrated in Fig. 3, only one machine acted as broker, and the shards were evenly distributed across the two machines ($B = 1, S = 2$). This is similar to the configuration described by Kulkarni (2013). However, this configuration is not optimal for selective search, and produces an uneven machine load, especially when using Rank-S to select shards. Figure 4 plots the machine utilization of this configuration, using Rank-S, $M = 2$ machines, and searching ClueWeb09 with the AOL_W Test query set. All broker tasks are handled by machine m_1 , which is also a searcher. Consequently, m_1 is heavily loaded compared to m_2 , with more than 70% of the load caused by resource selection when the query arrival rate is 2.5 queries per second. Similar imbalances are observed when other simulator parameters such as query streams, collections, and hardware are varied.

The situation changes if more broker processes are permitted. Figure 5 compares the previous $B = 1$ outcomes to $B = 2$, that is, a configuration in which both machines perform resource selection. The change results in a moderate gain in throughput on the Gov2 collection with the MQT_G queries, and a marked improvement in throughput on ClueWeb09 with the AOL_W queries; with the difference in behavior due to the size of the corresponding CSIs. Rank-S uses a 1% sample of the corpus to make decisions, and for Gov2, a 1% sample is about half the size of an average shard. But for ClueWeb09, the 1% CSI is about eight times the size of an average shard, and hence a greater fraction of the search time is spent on shard selection, as is also shown in Table 9.

Taily requires far less computation for resource selection than does Rank-S, and the best setting was $B = 1$ for both datasets and all query logs. At $M = 2$ and $B = 1$, resource selection for Taily accounted for less than 2% of m_1 's processing capability, reinforcing the natural advantage of term-based algorithms. As is shown in Fig. 3b, the resulting performance gains can be large. On the other hand, sample-based algorithms have other advantages, including the ability to run structured queries.

5.3 Shard assignment

With two machines ($M = 2$), random assignment of shards to machines distributes query traffic evenly across machines, since there are many more shards than machines ($P \gg M$). That natural balance is eroded as M increases, because selective search deliberately creates




Fig. 7 Utilization of machines for selective search on Gov2, using Taily, the MQT_G Test queries, $M = 4$, $B = 1$, $S = 4$, and two methods of assigning shards to machines. Each point is the mean over 10 sequences of 1000 queries; error bars represent 95% confidence intervals. The broker resides on m_1 in both configurations

Table 10 Average $load_i$ and range ($\max load_i - \min load_i$) on the test query set, where $load_i$ is the time-averaged CPU load of machine m_i , for two shard allocation policies and a range of query arrival rates, using $M = 4$, $B = 1$, and $S = 4$, and with Tailly resource selection throughout

Dataset	Queries	Allocation method	Average $load_i$ and range of $load_i$ for each query arrival rate T							
			6 qry/s		8 qry/s		10 qry/s		12 qry/s	
			avg.	rnge.	avg.	rnge.	avg.	rnge.	avg.	rnge.
Gov2	MQT _G	Random	0.32	0.30	0.42	0.40	0.53	0.49	0.61	0.58
		Log-based	0.32	0.10	0.42	0.14	0.53	0.17	0.63	0.21
			20 qry/s		30 qry/s		40 qry/s		50 qry/s	
Gov2	AOL _G	Random	0.40	0.19	0.59	0.28	0.73	0.34	0.78	0.36
		Log-based	0.40	0.14	0.59	0.20	0.73	0.25	0.78	0.26
			10 qry/s		15 qry/s		20 qry/s		25 qry/s	
CW09	MQT _W	Random	0.36	0.08	0.53	0.11	0.70	0.15	0.81	0.17
		Log-based	0.36	0.05	0.53	0.07	0.70	0.09	0.82	0.10
			1.5 qry/s		2.0 qry/s		2.5 qry/s		3.0 qry/s	
CW09	AOL _W	Random	0.39	0.07	0.52	0.10	0.65	0.12	0.76	0.14
		Log-based	0.39	0.01	0.52	0.02	0.65	0.02	0.77	0.03

CW09 stands for ClueWeb09

Table 11 Average $load_i$ and range ($\max load_i - \min load_i$) as the training data ages, using the Log-based shard allocation policy, and all other simulation settings as for Table 10

Dataset	Query log	Average $load_i$ and range of $load_i$ for each query arrival rate T							
		avg.	rnge.	avg.	rnge.	avg.	rnge.	avg.	rnge.
		20 qry/s		30 qry/s		40 qry/s		50 qry/s	
Gov2	AOL _G Test	0.40	0.14	0.59	0.20	0.73	0.25	0.78	0.26
	AOL _G Test _W	0.38	0.13	0.56	0.19	0.71	0.24	0.76	0.25
	AOL _G Test _M	0.32	0.14	0.47	0.20	0.62	0.26	0.72	0.30
		1.5 qry/s		2.0 qry/s		2.5 qry/s		3.0 qry/s	
CW09	AOL _W Test	0.39	0.01	0.52	0.02	0.65	0.02	0.77	0.03
	AOL _W Test _W	0.39	0.01	0.51	0.02	0.63	0.02	0.75	0.02
	AOL _W Test _M	0.40	0.02	0.53	0.02	0.66	0.02	0.77	0.03

The Test queries begin immediately after the training queries; the Test_W queries begin one week after the training queries; and the Test_M queries begin one month after the training queries. The MQT queries do not have timestamps. CW09 stands for ClueWeb09

shards that have skewed term distributions and topical coverage. This section examines Research Question 3: *How do different methods of shard allocation affect throughput and load distribution across multiple machines?*

Figure 6 shows the relative popularity of the 88 most frequently-accessed ClueWeb09 shards. For example, when the Rank-S resource selection algorithm is used to select topic-based shards for the Gov2 queries from AOL_G Train, the five most popular shards account for 29% of all shards selected—1302 out of 4491 shard selections. The MQT query set displays a similar, but more moderate, skew. This unevenness of access has the potential to

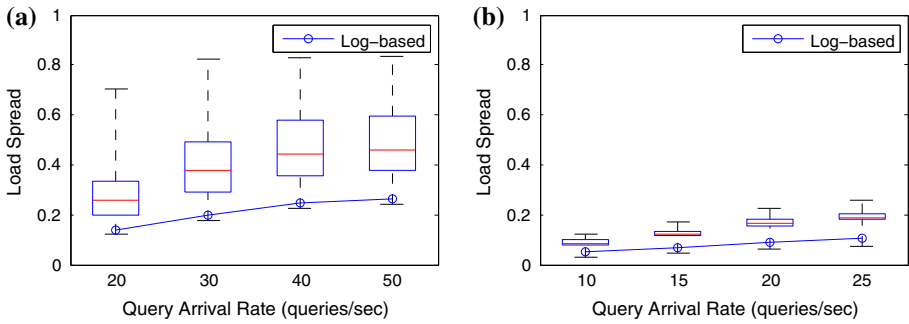


Fig. 8 Distribution of load spread (max $load_i - \min load_i$) for ten Random shard allocations, with $M = 4$ and $B = 1$. The mid-line of the box is the median, the outer edges the 25th and 75th percentile values, and the whiskers extend to the most outlying values. The load spread for Log-based shard allocation is also plotted, as a reference point. Note the different horizontal scales in the panes.

overload the machines that are responsible for popular shards, and thereby create bottlenecks that starve other machines of work.

The next set of experiments compare the performance of the usual *Random* shard assignment and an alternative *Log-based* mechanism, which uses training queries to estimate and balance the average load across machines. The training queries are fed through the resource selection process, so that the sum of the posting list lengths for each shard can be computed and used as an estimate of the shard’s workload. Shards are then ordered from most to least loaded, and assigned to machines one by one, in each case choosing the machine that currently has the lowest estimated load. Any remaining shards that were not accessed by the training queries are assigned similarly, based on their size. In these experiments, four sets of 1000 training queries were employed, AOL_G Train, MQT_G Train, AOL_W Train and MQT_W Train, as described in Sect. 4.4. Using Taily’s shard selections, Gov2 had no unaccessed shards for either query set, and for the ClueWeb09 collection, 7 and 4% of shards were unaccessed by AOL_W Train and MQT_W Train respectively.

Figure 7 shows the effect of shard assignment policy on workload across $M = 4$ hosts for the Gov2 dataset, as query arrival rates are varied, with the vertical axis showing machine utilization in the range 0.0 and 1.0. The wide spread of loads in the left-hand pane shows that the Random policy produces an uneven utilization of machines, and saturates

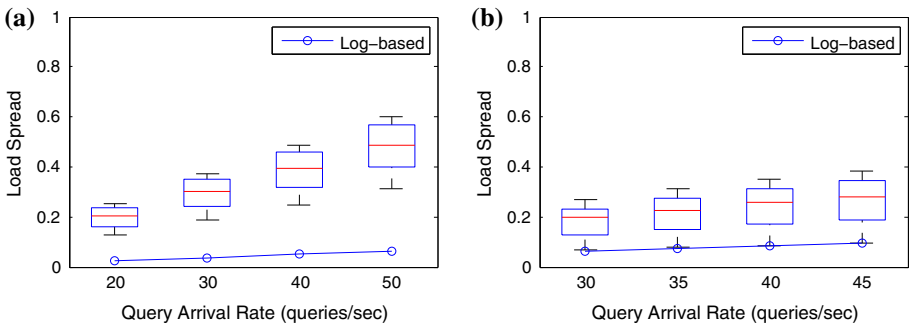


Fig. 9 As for Fig. 8, but with $M = 8$ and $B = 8$

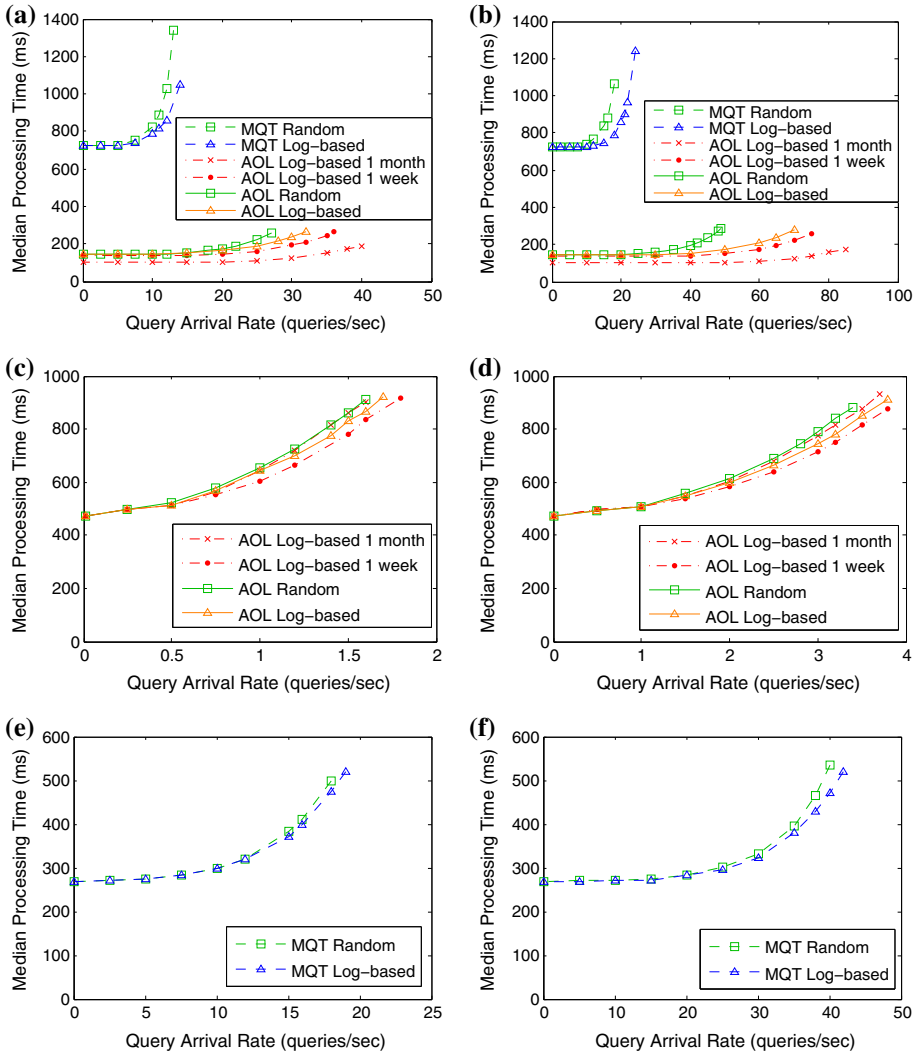


Fig. 10 The effect of shard assignment policies on throughput of selective search, using Taily with $M = 4$, $B = 1$ (all left column panes) and $M = 8$, $B = 8$ (all right column panes). Note the differing horizontal scales in each pair

relatively quickly due to a bottleneck on m_4 . In comparison, the Log-based policy markedly reduces the load variance, and allows higher query throughput rates to be attained. Numeric results for this and other configurations are given in Table 10, with Log-based assignment consistently producing more uniform resource utilization than Random assignment. Results for Rank-S resource selection are similar, and are omitted.

The risk of using Log-based allocations is that the learned attributes may become dated as a result of changes in the query stream. Table 11 investigates this potential shortcoming by showing machine usage for three time-separated query sets each containing 10,000 queries: one from immediately after the training queries; a second set from one week after the training queries; and a third set from one month after the training queries. Average

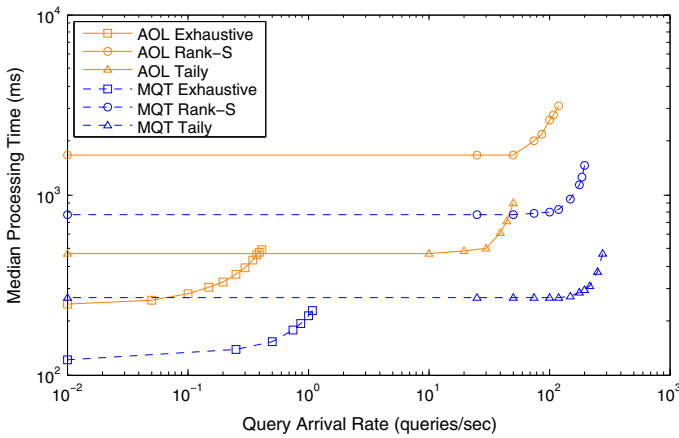


Fig. 11 Selective search and exhaustive search for ClueWeb09 and the AOL_W Test and MQT_W Test queries (shown as AOL and MQT in the legend), with $M = 64$ and $B = 64$, and Log-based shard assignments

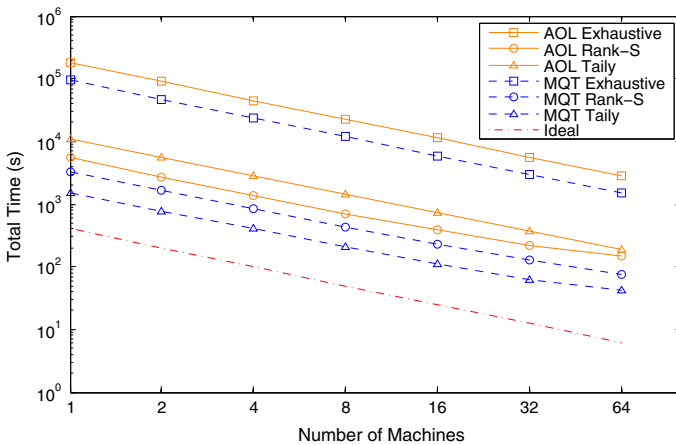


Fig. 12 Total time required to process 10,000 queries as a function of M , the number of machines, assuming an infinite query arrival rate and $B = M$. The two query streams used are MQT_G Test and AOL_G Test, denoted as MQT and AOL respectively. The *dash-dot line* represents the gradient expected for ideal linear speedup, where throughput is doubled with doubled machine capacity

utilization is similar in each case, and variance increases marginally as the query stream evolves, but the changes are generally small. Results for the AOL log for Gov2 one month after assignment are an exception and have a markedly lower utilization due to a burst of shorter queries that occurred at this time (2.18 words on average versus 2.41 and 2.44 for Test_W and Test queries respectively), meaning that at any given arrival rate less total work is required; but the variance still remains similar. Shard assignments should be periodically revised to maximize throughput, but it is not necessary to do it frequently, and the cost of refreshing shard assignments can be amortized.

To demonstrate that the random assignment results shown in Fig. 7 and Table 10 are not due to outliers, ten different allocations were generated per combination of collection

and query stream, and variance of the loading measured. Figure 8 presents the results for the AOL_G Test query set of Gov2 and the MQT_W Test query set of ClueWeb09. The load spread of the Log-based allocation is also plotted. While there was at least one Random allocation that produced a load spread comparable to that of the Log-based method, the median values were all well above it. Figure 9 shows that when the number of machines is increased to $M = 8$, Log-based allocation produces the best load distribution even when compared to ten different Random allocations. That is, the Log-based method reliably picks shard allocations with low load spread, and hence leads to consistent utilization across machines.

Balanced assignment of shards and higher utilization of machines can lead to higher query throughput, as shown in Fig. 10. The Random policies shown in this figure are the median performing assignment from Figs. 8 and 9. When $M = 4$, in the left column, the Log-based policies are measurably better than the Random policy for the Gov2 dataset. The difference widens when $M = 8$ machines are used. The gap between the allocation strategies for ClueWeb09 are smaller, due to the smaller differences in load variance between the Random and Log-based assignment, as already noted in connection with Table 10, and Figs. 8 and 9. Correcting larger load imbalances leads to more substantial throughput gains.

Finally, in most of the settings explored the Log-based assignment continues to yield better utilization than Random allocation even after a one month interval, further supporting the stability of the assignment. It is clear that the method of allocating shards to machines has an impact in load distribution of selective search, and hence throughput.

5.4 Scalability

The main focus of this paper is on low-resource environments, typically involving two to eight multi-core machines. But one advantage of a simulation model is that other configurations can also be explored, and we have examined larger-scale environments too. Distributed IR systems typically achieve scalability in two ways: by shard replication, and/or by increasing the number of machines available in the cluster so that each is responsible




Fig. 13 Throughput achieved for ClueWeb09, with $M = 64$, $B = 64$, and Log-based shard assignment. The mirrored configuration is a doubled $M = 32$, $B = 32$ system, including disk space. The two query streams are MQT_G Test and AOL_G Test.

for a smaller volume of data. We explore the scalability of selective search using these two approaches, and consider Research Question 4: *Does selective search scale efficiently when adding more machines and/or shard replicas?* Note that the experiments in this section are for ClueWeb09 alone, because Gov2 only has 50 shards.

Figure 11 compares selective search and exhaustive search when $M = 64$ machines are used, with load balancing as described in Sects. 5.2 and 5.3. With more machines in use, the latency of exhaustive search is greatly decreased, because the shards are all smaller. Selective search's latency remains the same because the number of topical shards searched is not a function of M . In the configuration shown there are 512 random shards and 884 topical shards, and the query terms' posting lists are now shorter on average in the random shards than in the topical shards selected by Taily or Rank-S. Note, however, that selective search still provides substantially higher throughput than exhaustive search; that is, while selective search takes longer to search each accessed shard, it continues to process a much smaller total volume of index data.

Figure 12 further demonstrates the throughput advantage held by selective search. In this experiment, the total time required to process 10,000 queries is measured as a function of the number of machines available. In all configurations, selective search remains a better option for throughput, requiring less elapsed time to process a fixed number of queries. This experiment also shows that selective search scales nearly linearly with the number of machines available. Both exhaustive and selective search are approximately parallel to the dash-dot line, plotted to show the gradient that would be observed in an ideal case where throughput is doubled with doubled machine capacity.

In Figs. 11 and 12 an *index spreading* strategy is assumed, in which the shards are distributed across the larger number of machines, and then the number of machines allocated to resource selection and shard search are adjusted to make the best use of the new hardware. When there are more machines than shards, index spreading is no longer feasible and a different method must be used. One alternative is to increase throughput by using additional space and adopting a *mirrored* architecture. For example, if the number of machines is doubled, two copies of each shard, twice as many resource allocation cores, and twice as many shard search cores can be used. Either approach can roughly double throughput.

Figure 13 compares the effect of index spreading and mirroring when $M = 64$ machines are employed, again plotting median query latency as a function of query arrival rate. In most configurations, mirroring has a small throughput advantage, derived from the fact that as the number of machines increases, it becomes increasingly less likely that an evenly balanced shard assignment will occur. For example, at $T = 280$, the standard deviation of the machine load in the ClueWeb09 MQT Test queries is 5.2% for 64 machines, and 3.9% for a 32 machine mirrored configuration using the index spreading strategy with Log-based allocation. While mirroring requires double the disk space of index spreading and the throughput differences are small, mirroring provides additional benefits, such as fault tolerance in case of failure. Furthermore, some amount of replication is necessary when there are more available machines than total shards.

The best overall solution in environments with high query traffic, or many machines may be a mix of index spreading and strategic shard mirroring (replication), an option also noted by Moffat et al. (2006). This is may be an interesting topic for future work.

6 Conclusion

Selective search is known to be substantially more efficient than the standard distributed architecture if computational cost is measured by counting postings processed in a one-query-at-a-time environment (Kulkarni and Callan 2010a, b; Kulkarni 2013). We have confirmed and extended those findings via a simulator that models a realistic parallel processing environment and a wide range of hardware configurations; using two large datasets and long query streams extracted from the logs of web search engines.

Previous investigations also demonstrated that selective search is as effective as exhaustive search (Kulkarni and Callan 2010a, b; Aly et al. 2013; Kulkarni 2013). Although our work here has focused primarily on efficiency, we refined the experimental methodology used to measure effectiveness, and achieved stronger baseline results for the ClueWeb09 dataset. One consequence of these changes was the discovery that in some cases selective search is less effective than exhaustive search, regardless of which resource selection algorithm is used. Effectiveness was not our focus in this work, but our findings are recorded to benefit other researchers that wish to pursue this topic.

Our investigation makes it clear that selective search is more efficient than conventional distributed query-processing architectures. For many hardware configurations the two-step selective search architecture—resource selection followed by shard access—delivers greater total throughput (number of queries processed per time interval) *as well as* lower latency (faster response time) than conventional exhaustive architecture. This somewhat surprising outcome is a consequence of the small size of the shards used by selective search, about 500 K documents each in our experiments. When many machines are available, exhaustive search latency decreases relative to selective search and it may become the faster of the two approaches, but selective search always has substantially higher throughput.

Other studies have shown that sample-based and term-based resource selection algorithms have different advantages and costs (Aly et al. 2013). Our experiments investigated the effects of the resource selection algorithm on load distribution, latency, and throughput. We found that Rank-S was more accurate than Taily, but also usually resulted in higher latency and lower throughput. Moreover, if Rank-S is chosen because of its greater effectiveness, the computational costs cause workload skews that must be corrected by replicating resource selection (and the CSI) on multiple machines.

Selective search uses topical shards that are likely to differ in access rate. Typical random assignments of shards produce imbalances in machine load, even when as few as $M = 4$ machines are in use. A Log-based assignment policy using training queries provided higher throughput and more consistent query processing times than the previous random assignment approach. The Log-based assignment is also resilient to temporal changes, and even after a delay of a month, throughput degradation was slight.

Previous studies investigated selective search in computing environments with a relatively small number of machines. With the aid of the simulator we also examined the behavior of selective search using clusters of up to $M = 64$ machines, each with eight processing cores. We found that selective search remains a viable architecture with high throughput in these larger-scale computing environments. When additional processing resources are available, mirroring (replicating all index shards) provides slightly better throughput than index spreading. Replication also has other advantages such as fault-tolerance and the ability to use more machines than there are shards. Replication of just a few high-load shards (rather than the entire index) might also be attractive, potentially

saving on storage costs without eroding throughput. We reserve this interesting problem for future work.

After investigating the efficiency of selective search architectures, we conclude that it is highly attractive in low resource environments typical of academic institutions and small businesses, and that the load imbalances of a naive configuration can be readily addressed. At larger scale, the latency advantages are lost unless smaller shards are formed, but even using the original shards, selective search continues to provide substantially higher throughput than exhaustive search.

Acknowledgments We thank the three referees for their detailed and helpful input. This work was supported by the National Science Foundation (IIS-1302206); and by the Australian Research Council (DP140101587 and DP140103256). Shane Culpepper is the recipient of an Australian Research Council DECRA Research Fellowship (DE140100275). Yubin Kim is the recipient of the Natural Sciences and Engineering Research Council of Canada PGS-D3 (438411). Any opinions, findings, conclusions or recommendations expressed in this paper are those of the authors, and do not necessarily reflect those of the sponsors.

References

- Altingovde, I. S., Demir, E., Can, F., & Ulusoy, O. (2008). Incremental cluster-based retrieval using compressed cluster-skipping inverted files. *ACM Transactions on Information Systems*, 26(3), 15:1–15:36.
- Aly, R., Hiemstra, D., & Demeester, T. (2013). Taily: Shard selection using the tail of score distributions. In *Proceedings of the 36th international ACM SIGIR conference on research and development in information retrieval* (pp. 673–682)
- Arguello, J., Callan, J., & Diaz, F. (2009). Classification-based resource selection. In *Proceedings of the 18th international ACM conference on information and knowledge management* (pp. 1277–1286)
- Badue, C. S., Baeza-Yates, R., Ribeiro-Neto, B., Ziviani, A., & Ziviani, N. (2007). Analyzing imbalance among homogeneous index servers in a web search system. *Information Processing and Management*, 43(3), 592–608.
- Baeza-Yates, R., Castillo, C., Junqueira, F., Plachouras, V., & Silvestri, F. (2007). Challenges on distributed web retrieval. In *Proceedings of the 23rd IEEE international conference on data engineering* (pp. 6–20)
- Baeza-Yates, R., Gionis, A., Junqueira, F., Plachouras, V., & Tello, L. (2009a). On the feasibility of multi-site web search engines. In *Proceedings of the 18th international ACM conference on information and knowledge management* (pp. 425–434)
- Baeza-Yates, R., Murdock, V., & Hauff, C. (2009b). Efficiency trade-offs in two-tier web search systems. In *Proceedings of the 32nd annual international ACM SIGIR conference on research and development in information retrieval* (pp. 163–170)
- Barroso, L. A., Dean, J., & Hölzle, U. (2003). Web search for a planet: The Google cluster architecture. *IEEE Micro*, 23(2), 22–28.
- Brefeld, U., Cambazoglu, B. B., & Junqueira, F. P. (2011). Document assignment in multi-site search engines. In *Proceedings of the 4th ACM international conference on web search and data mining* (pp. 575–584)
- Broccolo, D., Macdonald, C., Orlando, S., Ounis, I., Perego, R., Silvestri, F., & Tonello, N. (2013). Query processing in highly-loaded search engines. In *Proceedings of the 20th international symposium on string processing and information retrieval* (pp. 49–55)
- Broder, A.Z., Carmel, D., Herscovici, M., Soffer, A., & Zien, J. (2003). Efficient query evaluation using a two-level retrieval process. In *Proceedings of the 12th international ACM conference on information and knowledge management* (pp. 426–434)
- Burkowski, F. J. (1990). Retrieval performance of a distributed database utilising a parallel process document server. In *Proceedings of the 2nd international symposium on databases in parallel and distributed systems* (pp. 71–79)
- Cacheda, F., Carneiro, V., Plachouras, V., & Ounis, I. (2007). Performance analysis of distributed information retrieval architectures using an improved network simulation model. *Information Processing and Management*, 43, 204–224.

- Cahoon, B., McKinley, K. S., & Lu, Z. (2000). Evaluating the performance of distributed architectures for information retrieval using a variety of workloads. *ACM Transactions on Information Systems*, 18(1), 1–43.
- Callan, J. (2000). Distributed information retrieval. In *Advances in information retrieval* (pp. 127–150)
- Callan, J., Connell, M., & Du, A. (1999). Automatic discovery of language models for text databases. In *Proceedings of the 1999 ACM SIGMOD international conference on management of data* (pp. 479–490)
- Cambazoglu, B. B., Kayaaslan, E., Jonassen, S., & Aykanat, C. (2013). A term-based inverted index partitioning model for efficient distributed query processing. *ACM Transactions on the Web*, 7(3), 15:1–15:23.
- Cambazoglu, B. B., Varol, E., Kayaaslan, E., Aykanat, C., Baeza-Yates, R. (2010). Query forwarding in geographically distributed search engines. In *Proceedings of the 33rd international ACM SIGIR conference on research and development in information retrieval* (pp. 90–97)
- Can, F., Altıngövdü, I. S., & Demir, E. (2004). Efficiency and effectiveness of query processing in cluster-based retrieval. *Information Systems*, 29(8), 697–717.
- Croft, W. B. (1980). A model of cluster searching based on classification. *Information Systems*, 5(3), 189–195.
- Elsas, J. L., Arguello, J., Callan, J., & Carbonell, J. G. (2008). Retrieval and feedback models for blog feed search. In *Proceedings of the 31st annual international ACM SIGIR conference on research and development in information retrieval* (pp. 347–354)
- Francès, G., Bai, X., Cambazoglu, B. B., & Baeza-Yates, R. (2014). Improving the efficiency of multi-site web search engines. In *Proceedings of the 7th ACM international conference on web search and data mining* (pp. 3–12)
- Freire, A., Macdonald, C., Tonello, N., Ounis, I., & Cacheda, F. (2013). Hybrid query scheduling for a replicated search engine. In *Proceedings of the 35th European conference on information retrieval* (pp. 435–446)
- Gravano, L., García-Molina, H., & Tomic, A. (1999). GLOSS: Text-source discovery over the internet. *ACM Transactions on Database Systems*, 24, 229–264.
- Griffiths, A., Luckhurst, H., & Willett, P. (1986). Using inter-document similarity information in document retrieval systems. *Journal of the American Society for Information Science*, 37, 3–11.
- Hawking, D., & Thistlewaite, P. (1999). Methods for information server selection. *ACM Transactions on Information Systems*, 17(1), 40–76.
- Kang, C., Wang, X., Chang, Y., & Tseng, B. (2012). Learning to rank with multi-aspect relevance for vertical search. In *Proceedings of the 5th ACM international conference on web search and data mining* (pp. 453–462)
- Kim, J., & Croft, W. B. (2010). Ranking using multiple document types in desktop search. In *Proceedings of the 33rd international ACM SIGIR conference on research and development in information retrieval* (pp. 50–57)
- Kim, Y., Callan, J., Culpepper, J. S., & Moffat, A. (2016a). Does selective search benefit from WAND optimization? In *Proceedings of the 38th European conference on information retrieval* (pp. 145–158)
- Kim, Y., Callan, J., Culpepper, J. S., & Moffat, A. (2016b). Load-balancing in distributed selective search. In *Proceedings of the 39th international ACM SIGIR conference on research and development in information retrieval* (pp. 905–908)
- Kulkarni, A. (2013). *Efficient and effective large-scale search*. PhD thesis, Carnegie Mellon University
- Kulkarni, A., & Callan, J. (2010a). Document allocation policies for selective searching of distributed indexes. In *Proceedings of the 19th ACM international conference on information and knowledge management* (pp. 449–458)
- Kulkarni, A., & Callan, J. (2010b). Topic-based index partitions for efficient and effective selective search. In *SIGIR workshop on large-scale distributed information retrieval*
- Kulkarni, A., & Callan, J. (2015). Selective search: Efficient and effective search of large textual collections. *ACM Transactions on Information Systems*, 33(4), 17:1–17:33.
- Kulkarni, A., Tigelaar, A., Hiemstra, D., & Callan, J. (2012). Shard ranking and cutoff estimation for topically partitioned collections. In *Proceedings of the 21st ACM international conference on information and knowledge management* (pp. 555–564)
- Liu, X., & Croft, W. B. (2004). Cluster-based retrieval using language models. In *Proceedings of the 27th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 186–193)
- Lucchese, C., Orlando, S., Perego, R., & Silvestri, F. (2007). Mining query logs to optimize index partitioning in parallel web search engines. In *Proceedings of the 2nd international conference on scalable information systems* (pp. 43:1–43:9)

- Macdonald, C., Tonello, N., & Ounis, I. (2012). Learning to predict response times for online query scheduling. In *Proceedings of the 35th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 621–630)
- Markov, I., & Crestani, F. (2014). Theoretical, qualitative, and quantitative analyses of small-document approaches to resource selection. *ACM Transactions on Information Systems*, 32(2), 9:1–9:37.
- Metzler, D., & Croft, W. B. (2005). A Markov random field model for term dependencies. In *Proceedings of the 28th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 472–479)
- Moffat, A., Webber, W., Zobel, J. (2006). Load balancing for term-distributed parallel retrieval. In *Proceedings of the 29th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 348–355)
- Moffat, A., Webber, W., Zobel, J., & Baeza-Yates, R. (2007). A pipelined architecture for distributed text query evaluation. *Information Retrieval*, 10(3), 205–231.
- Orlando, S., Perego, R., & Silvestri, F. (2001). Design of a parallel and distributed web search engine. In *Proceedings of the international conference on parallel computing* (pp. 197–204)
- Paltoglou, G., Salamapasis, M., & Satratzemi, M. (2008). Integral based source selection for uncooperative distributed information retrieval environments. In *Proceedings of the 2008 ACM workshop on large-scale distributed systems for information retrieval* (pp. 67–74)
- Powell, A. L., French, J. C., Callan, J., Connell, M., & Viles, C. L. (2000). The impact of database selection on distributed searching. In *Proceedings of the 23rd annual international ACM SIGIR conference on research and development in information retrieval* (pp. 232–239)
- Puppini, D., Silvestri, F., & Laforenza, D. (2006). Query-driven document partitioning and collection selection. In *Proceedings of the 1st international conference on scalable information systems* (p. 34)
- Ribeiro-Neto, B. A., & Barbosa, R. A. (1998). Query performance for tightly coupled distributed digital libraries. In *Proceedings of the 3rd ACM conference on digital libraries* (pp. 182–190)
- Risvik, K. M., Aasheim, Y., & Lidal, M. (2003). Multi-tier architecture for Web search engines. In *Proceedings of the 1st Latin American web congress* (pp. 132–143)
- Seo, J., & Croft, W. B. (2008). Blog site search using resource selection. In *Proceedings of the 17th international ACM conference on information and knowledge management* (pp. 1053–1062)
- Shokouhi, M. (2007). Central-rank-based collection selection in uncooperative distributed information retrieval. In *Proceedings of the 29th European conference on information retrieval* (pp. 160–172)
- Shokouhi, M., & Si, L. (2011). Federated search. *Foundations and Trends in Information Retrieval*, 5(1), 1–102.
- Si, L., & Callan, J. (2003). Relevant document distribution estimation method for resource selection. In *Proceedings of the 26th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 298–305)
- Si, L., & Callan, J. (2004a). The effect of database size distribution on resource selection algorithms. In *Distributed multimedia information retrieval* (pp. 31–42). LNCS volume 2924
- Si, L., & Callan, J. (2004b). Unified utility maximization framework for resource selection. In *Proceedings of the 13th international ACM conference on information and knowledge management* (pp. 32–41)
- Si, L., & Callan, J. (2005). Modeling search engine effectiveness for federated search. In *Proceedings of the 28th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 83–90)
- Thomas, P., & Hawking, D. (2009). Server selection methods in personal metasearch: A comparative empirical study. *Information Retrieval*, 12(5), 581–604.
- Thomas, P., & Shokouhi, M. (2009). SUSHI: Scoring scaled samples for server selection. In *Proceedings of the 32nd annual international ACM SIGIR conference on research and development in information retrieval* (pp. 419–426)
- Tomasic, A., & Garcia-Molina, H. (1993). Caching and database scaling in distributed shared-nothing information retrieval systems. In *Proceedings of the 1993 ACM SIGMOD international conference on management of data* (pp. 129–138)
- Tonello, N., Macdonald, C., & Ounis, I. (2013). Efficient and effective retrieval using selective pruning. In *Proceedings of the 6th ACM international conference on web search and data mining* (pp. 63–72)
- Voorhees, E. M. (1985). The effectiveness and efficiency of agglomerative hierarchic clustering in document retrieval. Technical report, Cornell University
- Webber, W., & Moffat, A. (Dec. 2005). In search of reliable retrieval experiments. In *Proceedings of the 10th Australasian document computing symposium* (pp. 26–33)
- Willett, P. (1988). Recent trends in hierarchic document clustering: A critical review. *Information Processing and Management*, 24(5), 577–597.

- Wu, H., & Fang, H. (2014). Analytical performance modeling for top-k query processing. In *Proceedings of the 23rd ACM international conference on information and knowledge management* (pp. 1619–1628)
- Xu, J., & Croft, W.B. (1999). Cluster-based language models for distributed retrieval. In *Proceedings of the 22nd annual international ACM SIGIR conference on research and development in information retrieval* (pp. 254–261)
- Yuwono, B., & Lee, D. L. (1997). Server ranking for distributed text retrieval systems on internet. In *Proceedings of the 5th international conference on database systems for advanced applications* (pp. 41–49)
- Zhang, J., & Suel, T. (March 2007). Optimized inverted list assignment in distributed search engine architectures. In *Parallel and distributed processing symposium* (pp. 1–10)