# Modeling Performance-Driven Workload Characterization of Web Search Systems


Claudine Badue[1] claudine@dcc.ufmg.br
Ricardo Baeza-Yates[2] ricardo@baeza.cl
Berthier Ribeiro-Neto[3] berthier@dcc.ufmg.br
Artur Ziviani[4] ziviani@lncc.br
Nivio Ziviani[5] nivio@dcc.ufmg.br

[1,3,5]Computer Science Dept., Federal University of Minas Gerais, Belo Horizonte, Brazil
[4]National Laboratory for Scientific Computing (LNCC), Petrópolis, Brazil
[2]Yahoo! Research, Barcelona, Spain & Santiago, Chile
[3]Google Engineering Belo Horizonte, Brazil



**Categories and Subject Descriptors:** H.3.3 [Information Search and Retrieval]:Search process

**General Terms:** measurement, performance.

**Keywords:** parallel query processing, search engines, workload characterization, performance analysis.


## 1. INTRODUCTION

Previous work on workload characterization for Web search systems mainly focuses on the characterization of user search behavior [4, 7]. In this paper, however, we model workloads for a Web search engine from a system performance point of view, analyzing both the distribution of the interarrival times of queries and the per-query execution time. This is crucial for performance evaluation and capacity planning purposes. Our results are derived from experiments in an information retrieval testbed fed with a log of a real-world search engine. Previous analytical models for search systems considered that interarrival times and service times are exponentially distributed [6]. In contrast, we verify in practice that there is a high variability in both interarrival times of queries reaching a search engine and service times of queries processed in parallel by a cluster of index servers. We also show that this highly variable behavior of interarrival times and service times is very well captured by hyperexponential distributions.

## 2. SEARCH SYSTEM ARCHITECTURE

Search engines typically adopt a computational cluster as a platform for parallel query processing [5]. Usually, this cluster is composed of a single broker and a set of index servers. The whole collection of documents is partitioned among the index servers, such that each server stores its own local subcollection [2]. We assign each document to an index server randomly, a policy that works fine in balancing storage space utilization among servers [3]. The broker accepts queries from a user and forwards them to the index servers, triggering the parallel query processing. Each index server searches its own local subcollection and produces a partial ranked answer. These partial ranked answers are collected by the broker and combined through an in-memory merging operation. The final list of ranked documents is then sent back to the user. We adopt an inverted index as the indexing structure for each subcollection, and the standard vector space model to rank the selected documents.

For the experiments, we use a cluster of 8 identical index servers. The collection used in our tests is composed of 10 million Web pages collected by the TodoBR[1] search engine from the Brazilian Web in 2003. The query set used in our tests is composed of 100 thousand queries extracted from a partial log of queries submitted to the TodoBR search engine in September 2003.

## 3. CHARACTERIZING INTERARRIVAL AND SERVICE TIMES

First, we evaluate the distribution of interarrival times of queries that reach the cluster of index servers. In our test collection, the coefficient of variation of interarrival times of queries is equal to 2.64, suggesting that interarrival times of queries might be well fitted by an hyperexponential distribution, which is characterized by a coefficient of variation larger than one. The hyperexponential distribution is an special case of the phase-type (PH) distribution [9] that results from the combination of a certain number of exponential distributions.

Figure 1 shows the empirical distribution of interarrival times of queries, which is fitted using an exponential distribution and using an hyperexponential distribution composed of 3 exponential phases. We observe that the hyperexponential distribution captures better the highly variable behavior of the experimental data and does a good job in matching the distribution of interarrival times. Conversely, we observe that the exponential fitting is poor and is not capable of capturing the high variability present in interarrival times. To evaluate the accuracy of the fitting, we use the Kolmogorov-Smirnov goodness-of-fit test to indicate

[1]TodoBR is a trademark of Akwan Information Technologies, which was acquired by Google in July 2005.


Copyright is held by the author/owner(s).


*CIKM'06,* November 5–11, 2006, Arlington, Virginia, USA.
ACM 1-59593-433-2/06/0011.

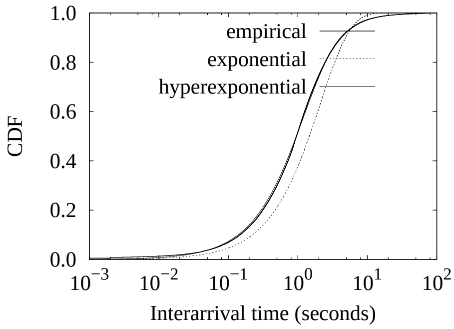


**Figure 1: Distribution of interarrival times.**

that the maximum ordinate difference between the distribution of interarrival times and the fitted exponential distribution is 15.40%, while this difference for the fitted hyperexponential distribution with 3 phases is only 1.37%.

We also found evidence that interarrival times presents self-similarity and long-range dependence (LRD) [8]. LRD characterizes processes that keep statistically significant correlations across large time scales, while a process is said to be self-similar if its behavior is roughly the same across different spatial or time scales. The Hurst parameter $H$ in the range $0.5 < H < 1$ is a measure of the level of self-similarity of a time series that exhibits LRD. A value of 0.5 indicates the absence of self-similarity. The closer $H$ is to 1, the greater the degree of LRD. In our experimental data, interarrival times of queries presented $H = 0.88$.

Second, we characterize the distribution of service times of queries processed in parallel by the cluster of index servers. In our architecture, characterized by a local partitioning of the document collection, the per-query service time is determined by the maximum execution time among index servers participating in the parallel processing of this particular query. Our experimental results show that the coefficient of variation of the service times of queries in our cluster with 8 index servers is equal to 1.39. This high variability in service times of queries is mainly due to the presence of a certain level of imbalance in per-query execution times among homogeneous index servers participating in the parallel processing. The primary source of the per-query imbalance is the heterogeneous use of disk caching among the homogeneous index servers, as pointed out in [1]. Thus, the coefficient of variation indicates the hyperexponential distribution as an appropriate kind of distribution for fitting the service times of queries.

Figure 2 shows the distribution of service times of queries fitted using an hyperexponential distribution composed of 3 exponential phases. Similarly to the case of interarrival times, we observe that the fitting provided by the hyperexponential distribution with 3 phases is very accurate. The Kolmogorov-Smirnov goodness-of-fit test verifies that the maximum ordinate difference between the distribution of service times and the fitted hyperexponential distribution with 3 phases is only 2.29%, while this difference for the fitted exponential distribution is equal to 20.43%.

Likewise the distribution of interarrival times, the distribution of per-query service times also presents self-similarity and LRD characteristics. In the case of the considered distribution of per-query service times, the Hurst parameter is $H = 0.74$.

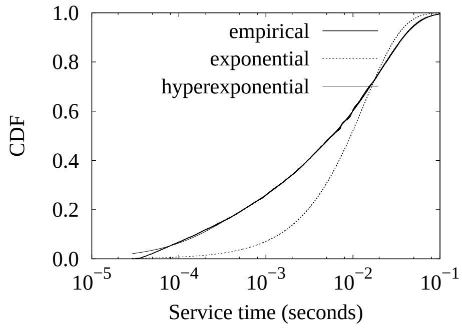


**Figure 2: Distribution of service times.**

Our findings thus suggest that typical workloads from a real-world Web search system presents high variability, self-similarity, and LRD characteristics. This finding is important because exponential-based models tend to underestimate the performance of the system, leading to unrealistic results, as the prediction without taking into account the high variability, self-similarity, and LRD in interarrival times and service times might underestimate queueing delay and, as a consequence, response time of the system as a whole.

As future work, we intend to investigate analytical models for performance prediction and capacity planning of Web search systems, while taking into account the high variability, self-similarity, and LRD characteristics found in interarrival and service processes.

# 4. REFERENCES


[1] C. Badue, R. Baeza-Yates, B. Ribeiro-Neto, A. Ziviani, and N. Ziviani. Analyzing imbalance among homogeneous index servers in a web search system. *Information Processing and Management*, to appear.

[2] C. Badue, R. Baeza-Yates, B. Ribeiro-Neto, and N. Ziviani. Distributed query processing using partitioned inverted files. In *Proceedings of the 8th Symposium on String Processing and Information Retrieval (SPIRE'01)*, pages 10–20, Laguna de San Rafael, Chile, 2001.

[3] C. Badue, R. Barbosa, P. Golgher, B. Ribeiro-Neto, and N. Ziviani. Basic issues on the processing of web queries. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'05)*, pages 577–578, Salvador, Bahia, Brazil, 2005.

[4] R. Baeza-Yates, C. Hurtado, M. Mendoza, and G. Dupret. Modeling user search behavior. In *Proceedings of the 3rd Latin American Web Congress (LA-WEB'05)*, pages 242–251, Buenos Aires, Argentina, 2005.

[5] L. A. Barroso, J. Dean, and U. Holzle. Web search for a planet: The google cluster architecture. *IEEE Micro*, 23(2):22–28, 2003.

[6] A. Chowdhury and G. Pass. Operational requirements for scalable search systems. In *Proceedings of the 12th International Conference on Information and Knowledge Management (CIKM'03)*, pages 435–442, New Orleans, LA, USA, 2003.

[7] B. Jansen and A. Spink. An analysis of web searching by european alltheweb.com users. *Information Processing and Management*, 41(2):361–381, 2005.

[8] T. Karagiannis, M. Molle, and M. Faloutsos. Long-range dependence: Ten years of Internet traffic modeling. *IEEE Internet Computing*, 8(5):57–64, 2004.

[9] A. Riska, V. Diev, and E. Smirni. An EM-based technique for approximating long-tailed data sets with PH distributions. *Performance Evaluation Journal*, 55(1-2):147–164, 2004.