Network Analysis for Distributed Information Retrieval Architectures

Fidel Cacheda[1], Victor Carneiro[1], Vassilis Plachouras[2], and Iadh Ounis[2]

[1] Dept. ICT, University of A Coruña, Campus de Elviña s/n, 15071 A Coruña, Spain
{fidel, viccar}@udc.es
[2] Department of Computing Science, University of Glasgow, Glasgow, G12 8QQ, UK
{vassilis, ounis}@dcs.gla.ac.uk

Abstract. In this study, we present the analysis of the interconnection network of a distributed Information Retrieval (IR) system, by simulating a switched network versus a shared access network. The results show that the use of a switched network improves the performance, especially in a replicated system because the switched network prevents the saturation of the network, particularly when using a large number of query servers.

1 Introduction

This study is a continuation of our previous work, introduced in [1] and extended in [2], on the choice of optimal architectures for building a distributed large-scale IR system. The SPIRIT collection (94,552,870 documents and 1 terabyte (TB) of text) [3] was used in these previous studies to simulate a distributed IR system using a local inverted file strategy, with the aim of measuring the performance for different configurations (distributed, replicated and clustered systems). In the local inverted file strategy, each query server is responsible for a disjoint subset of documents and has an independent local index. Tomasic and Garcia-Molina [6] proved that this strategy uses the system resources effectively and provides a good query throughput while being more resilient to failures.

The main objective of this work is to improve the interconnection network of a distributed system, by defining a switched network in order to analyse the improvements in performance as compared to the shared access network from our previous study.

2 Simulation Model

To explore the performance of different architectures for a distributed large-scale IR system, we implemented a discrete event-oriented simulator using the JavaSim simulation environment [4]. The simulated distributed IR system is an extension of the Terrier IR system described in [5]. Moreover, we use the analytical model described in [1] and [2] for the simulation of the querying process in the distributed IR system. The SPIRIT collection [3] is simulated (94,552,870 documents and on average 456

D.E. Losada and J.M. Fernández-Luna (Eds.): ECIR 2005, LNCS 3408, pp. 527 – 529, 2005.

words per document). In order to test the performance, we generate 50 queries, following the skewed query model [1] [2]. The performance is measured using 5 different simulations and calculating the corresponding average throughput.

In our previous studies in [1] and [2], the simulated distributed IR system contained a single shared access LAN, which had certain limitations that reduced the capacities of the simulated IR systems. To improve these limitations, we have defined a new network model equivalent to a switched network FastEthernet 100BASE-T at 100Mbps. This new model represents the interconnection using switches, assuming that each switch has a capacity for 64 hosts. Furthermore, the network overhead is analysed exhaustively, considering the network protocol headers, IP fragmentation, and even the propagation delay.

We analyse the similarity between the proposed model and a real network, by measuring the time to send messages between two PCs (AMD Athlon, 2 GHz and 512 MB RAM), through a switched LAN. According to the Mann-Whitney and Kolmogorov-Smirnov two sample tests, the correspondence between the real and estimated transmission times is statistically significant, with p-values higher than 0.90.

3 Network Experiments

We examine the differences in performance between a switched and a shared access network (both operating at 100 Mbps), for both a distributed system and a replicated system. In these experiments, the collection of documents is distributed using the local inverted file strategy over N query servers (N = 1, 2, …, 1024).

Table 1. Throughput (queries/second) for different replicated IR systems with the optimal number of bropokers using a shared (*sh*) and switched (*sw*) LAN. The obtained improvement is also indicated (*%*)

Query	*R=1*			*R=2*			*R=3*			*R=4*		
servers	*sh*	*Sw*	*%*	*sh*	*sw*	*%*	*sh*	*sw*	*%*	*sh*	*sw*	*%*
1	0.02	0.05	170.6%	0.03	0.09	179.1%	0.05	0.13	177.7%	0.06	0.19	212.6%
2	0.03	0.09	161.7%	0.06	0.17	183.1%	0.09	0.25	183.2%	0.11	0.32	179.4%
4	0.06	0.16	147.7%	0.11	0.30	171.1%	0.15	0.42	180.3%	0.19	0.53	176.1%
8	0.11	0.25	129.1%	0.20	0.46	132.0%	0.27	0.71	166.3%	0.36	0.90	152.9%
16	0.18	0.37	103.8%	0.34	0.72	108.8%	0.47	1.00	114.2%	0.64	1.39	118.9%
32	0.29	0.49	69.7%	0.53	0.95	77.4%	0.77	1.38	79.1%	0.99	1.85	87.0%
64	0.41	0.58	38.9%	0.78	1.15	46.1%	1.18	1.67	41.6%	1.48	2.17	46.2%
128	0.53	0.64	21.4%	0.98	1.26	29.0%	1.42	1.83	29.0%	1.93	2.38	23.4%
256	0.60	0.67	11.8%	1.16	1.33	14.9%	1.70	1.93	13.8%	2.18	2.51	14.9%
512	0.64	0.69	6.6%	1.24	1.35	9.4%	1.78	1.95	9.5%	2.05	2.53	23.0%
768	0.66	0.69	3.9%	1.24	1.34	8.4%	1.42	1.93	35.4%	1.45	2.46	69.8%
1024	0.66	0.69	3.4%	0.99	1.32	33.9%	1.09	1.89	73.9%	1.13	2.40	112.0%

In Table 1 we detail the throughput obtained for a distributed and replicated IR system (with *R* replicas) with the optimal number of brokers, using both a shared and a switched network (obtained empirically as *2R + 1* and *3R*, respectively). In a distributed system (*R=1*), the switched network improves the performance in all cases. However, this improvement is reduced when more query servers are added to the system, because the brokers, merging all the partial results, become the bottleneck [1] [2] and the network has a less significant impact on the final performance.

In a system with more replicas (*R*=2, 3 and 4), the percentage of improvement decreases as the number of query servers increases (following the pattern described for the distributed system), except for the maximum number of query servers where this percentage augments. This is motivated by the fact that, in a replicated system using a shared LAN, the bottleneck is the network, which becomes saturated as the number of query servers increases [1] [2]. The switched network prevents this saturation and therefore increases considerably the performance.

4 Conclusions

This paper is the continuation of our previous work on different architectures for a distributed IR system [1] [2]. In this work, we show that the use of a switched network for a distributed system improves the throughput performance in all cases. However, this improvement in performance is more important in replicated systems, as in these cases the bottleneck is the interconnection network. In our future work, we intend to apply these results to the clustered system analysis and we will study different techniques to reduce the brokers' bottleneck.

References

1. Cacheda, F., Plachouras & V., Ounis, I.: Performance Analysis of Distributed Architectures to Index One Terabyte of Text. In *Proceedings of 26th European Conference on Information Retrieval Research (ECIR'03)*, (2003), pp. 394-408.
2. Cacheda, F., Plachouras, V. & Ounis, I.: A Case Study of Distributed Information Retrieval Architectures to Index One Terabyte of Text. *Information Processing and Management Journal*, Articles in Press, (2004).
3. Jones, C. B., Purves, R., Ruas, A., Sanderson, M., Sester, M., van Kreveld, M. & Weibel, R.: Spatial information retrieval and geographical ontologies an overview of the SPIRIT project. In *Proceedings of the 25th ACM-SIGIR Conference on Research and Development in Information Retrieval*, pp. 387-388 (2002). New York: ACM Press.
4. Little, M. C.: *JavaSim User's Guide. Public Release 0.3, Version 1.0*. University of Newcastle upon Tyne. http://javasim.ncl.ac.uk/manual/javasim.pdf (2001).
5. Plachouras, V., Ounis, I., Amati, G. & van Rijsbergen C.J.: University of Glasgow at the Web track of TREC 2002. In Voorhees E.M. & Buckland, L.P. (Ed.), *Proceedings of The Eleventh Text REtrieval Conference*, 645-651. (2002) NIST Special Publication 500-251.
6. Tomasic, A. & Garcia-Molina, H.: Performance of inverted indices in shared-nothing distributed text document information retrieval systems. In *Proceedings of the 2nd International Conference on Parallel and Distributed Information Systems*, pp: 8-17 (1993).