A scalable distributed multimedia directory

I.W.Marshall
BT Labs
Martlesham Heath, Ipswich, IP5 3RE, UK
+44 1473 646951
Ian.W.Marshall@bt.com

Ng Chun Kiam
National University of Singapore,
10 Kent Ridge Crescent,
Singapore, 119260
eng50049@leonis.nus.edu.sg

Clare Bradford
BT Labs
Martlesham Heath, Ipswich, IP5 3RE, UK
+44 1473 649473
Clare.Bradford@bt.com

1. ABSTRACT

A prototype distributed multimedia directory based on scalable technology used in the World Wide Web has been constructed and tested. Techniques for flexibly adding multimedia features and supporting efficient user queries are discussed.

2. INTRODUCTION

The World Wide Web (WWW) is the only multimedia distribution and retrieval system which has so far been implemented on a global scale. The BT multimedia trial [1] in Colchester is still the largest non-WWW multimedia deployment that has been realised. In that trial a key finding was that some of the technology used was not very scalable. In particular the centralised techniques used for adding and indexing content proved problematic. In order to ensure that future trials use scalable technology we have undertaken a study of the WWW.

Based on our experiences in the trial it is vital to have a simple method of content addition, and index compilation, that can be easily automated, but which nevertheless provides a complete and easy to use content index. The WWW provides a simple scalable method of content addition but currently very weak index mechanisms [2].

Any scalable multimedia system will need to distribute content onto servers which are relatively close to the users in order to minimise waiting times [3]. Persistent network caches [4] enable local servers to be populated with minimal management overhead. This mechanism delivers considerable benefit to users [5] of the WWW. However, studies [6] have shown that there is a significant proportion of items in the web which are only requested once from a particular location, and caching them delivers no benefit. It can be shown [7] that many of these items are retrieved by users who are making poorly formulated, machine assisted searches and are probably largely unread. This highlights the need to provide users with an effective scalable directory (unlike existing search engines which are not keeping up with the growth of the WWW [2]) and also with tools that enable them to formulate efficient queries.

We have implemented and tested a text based index, based on a cache hierarchy equipped with catalogue servers [8]. We propose mechanisms enabling the addition of automated indexing of images and videos [9,12]. The benefits of the index were estimated from observations of traffic on an experimental cache in our Laboratory. The principal benefit results from the automation of the index build. Additional benefits arise from enabling the directory to prioritise matching items on the basis of "nearness" to the user, and thereby improving the probability that the users needs will be met from his local cache. There is also a benefit which results from the directory being a better match to users requirements for access history [10] than current browser stacks. We predict further substantial benefits when the index is associated with client tools that support users building efficient queries [11].

3. MEASUREMENTS

The distributed directory was installed on a cache hierarchy made up of single disk PCs. The directory was tested by a small community of users and a test client was used to simulate heavy loads and identify the system bottlenecks. The directory created no additional management load. The disk space taken up by the indexes corresponded to 8% of the space taken by the cached content being indexed. Users reported the directory was useful. The number of observed requests to external search engines was halved. We found that one user was using the directory in preference to his favourites file.

The resource requirement of the directory was estimated using a performance monitor tool. The results (table 1) showed that the index server was only consuming a few % of the servers resources. Responding to local queries with an average response of 50 pages had a measurable cost on disk i/o. However, the server could easily support more than 10 queries/second without any

	Running index	Each Local query	Distributed Query
Disk space used	8 Mbytes	0	100 kbytes
% processor time	1	3 over 100ms	5
% disk read time	4	80 over 100ms	100 over 200 ms
% disk write time	2	0	60 over 200ms
Avg disk queue	0.1	0	2
Memory bytes used	2.7 Mbyte	200 kbyte	600 kbyte
Page faults/sec	0	1	7

Table 1 Summary performance statistics

performance degradation which is more than adequate for the community of 50 concurrent users the server was configured to support. Adding extra disks would improve performance considerably. Responding to distributed queries was more expensive. This is not surprising as the implementation was only intended to prove the concepts.

The initial results are sufficient to indicate that a directory serving a large user community could be built using these techniques and would consume a small proportion of server resources. A large scale solution would use a cluster of web servers with a separate member of the cluster building the index and responding to queries. On the basis that typically 5% of requests are to some kind of directory, the cluster could support up to 200 – 1000 requests/s if it had a single index server supporting 10-50 requests/s. A typical active user makes 200 requests/day. If all these requests from all the users were in a busy hour this cluster could serve a community of 3600 – 18000 concurrent users.

4. DISCUSSION

The directory we have implemented has been shown to deliver a scalable automated index build, whilst only imposing a small overhead on the servers running the index. However, the index is currently only based on the text content of objects. A full multimedia directory will need to provide indexing for video, image and audio content. We believe that multimedia indexing will always be based on reducing all media to text, since the queries will always be handled as text. We are therefore working on extensions to the index server which transform audio to text, sound track to text (for video), and image metadata generation. The extensions are built as serverlets or proxylets [13] thus other extensions could also be added easily and even dynamically.

A further difficulty is that the simple queries preferred by users are extremely inefficient, since they generate a large number of matches with the index. To resolve this we are investigating structuring the index by knowledge domain, and also tools which assist users to formulate efficient queries. The current design is based on a form which offers the user more opportunity to provide information (like a library reference request). The form is backed by a rule based engine which asks additional questions until the query has a small number of matches. We feel that the query support engine will prove to be more useful than agents which attempt to predict the needs of individual users. This is based on analysis by ourselves and others showing that real user requests form a Fractal (i.e. not very predictable) sequence.

5. CONCLUSION

We have demonstrated that WWW technology forms a reasonable basis for building a scalable distributed multimedia directory. However, more work on enabling effective queries is needed.

6. REFERENCES

[1] A.Livingstone “BT interactive TV”, BT engineering, 15 page 254 (1996)

[2] K.Bharat and A.Broder “A technique for measuring the relative size and overlap of public Web search engines” Computer Networks and ISDN systems, 30, 379-388 (1998)

[3] S.Manley and M.Seltzer “Web facts and fantasy” Proc USENIX symposium on Internet technologies and systems, Monterey, Dec 1997

[4] A.Luotonen and K.Altis “World wide web proxies” Computer Networks and ISDN systems 27, 2 (1994)

[5] N.Smith “What can archives offer the World wide web” Proc 1st WWW conf Geneva 1994

[6] V.Almeida, A.Bestavros, M.Crovella and A.de Oliveira “Characterising reference locality in the WWW” Proc IEEE PDIS’96

[7] I.Marshall and C.Roadknight “Linking cache performance to user behaviour” Proc web cache workshop, Manchester 1998

[8] The harvest information discovery and access system http://harvest.transarc.com/

[9] W.Chang et al “Efficient resource selection in distributed visual information systems” Proc ACM multimedia ’97 203-213

[10] L.Tauscher and S.Greenberg “How people revisit web pages” International journal of human computer studies, 47, 1, 97-138 (1997)

[11] W.Li and J.Shim “Facilitating complex web queries through visual user interfaces” Comp nets and ISDN systems 30, 149-59 (1998)

[12] A.Ghosh and M.Fry “Application layer active networking” Proc HIPPARCH workshop London 1998