Using the Directory for Network Traffic Management

Optimising FTP Traffic in the Internet

K. Jayanthi*, Glenn Mansfield*, Y. Kimura*
Thomas Johannsen†, Y.Nemoto•, S. Noguchi⊕

* *Advanced Intelligent Communication Systems Laboratories, 6-6-3, Minami Yoshinari, Aoba-ku, Sendai 989-32, Japan.*
email: jaya@aic.co.jp
TEL: +81-22-279-3310. FAX: +81-22-279-3640

†*Dresden University of Technology, Dresden, Germany.*
•*Tohoku University, Sendai, Japan.*
⊕ *Nihon University, Koriyama, Japan.*

Abstract

Network configuration information can be very effectively used for network traffic optimization and improving network performance. FTP traffic, one of the largest consumers of network bandwidth, is taken as a case study of how effectively network configuration information can be used. The idea is essentially to provide a workable and scalable mechanism by which a user or an application can evaluate the network cost *factor corresponding to several ftp-servers. By choosing the server with the lowest cost factor,* network cost *of the retrieval is optimized. Cost could pertain to response time, tariff, or any other parameter. A network model based on which the comparative* network costs *of several servers in a network, as seen by a user or application, can actually be calculated is shown. Two implementations which realize the idea are discussed.*

1 Introduction

Internet in recent times has grown phenomenally, and there is an active interest in reviewing the framework and operations. There is a movement towards making the infrastructural and service related information available to users and applications at all places[MANS]. It is envisaged that such information can be very effectively used for network traffic optimization and improving network performance. A host of useful applications based on this directory information are coming up - from drawing the network map to finding the optimal paths between two nodes in the network.

In this paper, we concentrate on the method to optimise network traffic - more specifically, ftp (file transfer protocol) traffic, which forms the bulk of network traffic [ASABA], [MERIT]). A study [FTP] of ftp retrievals has shown that there is a substantial non-optimality involved. The trend does provide a motivation to examine the possibility of optimizing ftp retrievals. Theoretically, given the cost of all the links in a network, it is a straightforward exercise to compute the cost of a path and thence to deduce the cheapest path. However, from a practical point of view it is unreasonable to expect cost information on all the links to be available at hand; even if it was available the network cost of fetching the information would be high, further the computation costs will become unwieldy too. To alleviate this problem we make use of the hierarchical structure of the network.

To illustrate the concept, we have made a case study of a simplified map of the Japanese Internet. Two implementations, SoftPages and SmartArchie based on this concept are discussed.

2 Optimality in Network Traffic

The global network can be considered as an undirected graph, a mesh. The edges represent lines with certain properties and nodes the gateways. Usually, the optimization problem is to find the shortest path between the nodes A and B. There are several algorithms that deal with *shortest-path routing*. These algorithms fall into two categories - *distance-vector algorithms DVA* and *link state algorithms LSA*. In the DVA a node (router) knows the length of the

shortest path from each of its neighbours to every destination in the network. This information is used to compute the distance of the router itself from all destinations in the network and thereby select the shortest path. In the LSA each node (router) knows the entire network topology based on which it calculates the shortest path to the destination.

Our work is not concerned with the routing issues and algorithms. We look at the map of the network where the *shortest path* is being used by the underlying IP implementations. Thereby the reachability between the nodes is fixed - so is the communication cost for inter-node communication. Admittedly, there is an element of inaccuracy as the routing is dynamic, leading to costs that are varying. However from experience, the variation is small and the simplification is logical.

A path can be evaluated by summing the weights of all parts (edges) along this path. Routing protocols try to solve this optimization problem. In simple routing protocols, all edges have the weight 1. In that case, only the number of edges (called number of hops) counts for path evaluation and optimization. For networks, this means that the (different) bandwidth of links are not considered when estimating a path. There is, however, a significant difference between two links, when one has a capacity of say 512 kByte/s, and the other is a 9.6 kByte/s line.

For distributed network applications, path optimization is an even more difficult problem. Also with end-to-end routes optimized by lower-level protocols, a client-server-communication can be non-optimal. This can occur when the application allows a client to choose one out of several (apparently identical) servers to contact. The connection to each of these servers is in itself optimal, i.e. there is no better way to reach server X. But the link to server X might be worse than the link to server Y. Presently, applications like ftp do not provide a way to compare alternate client-server-connections with respect to network characteristics.

Users are basically concerned with a "fast" connection to remote services. It does not make a difference where a server is located or which organization it belongs to. But one can foresee that in the near future users will also have to worry about the cost of sending network traffic over specific links or using certain services.

In this study we propose a network model based on which calculation of network cost may be carried out relatively easily. A simplified version of the backbone of the Japanese Internet is taken as a case study to illustrate the application of the idea. The model makes it possible to compare the cost of retrieval from several servers in the network.

3 Cost calculation - the Network model

Strictly speaking for a mesh network consisting of n elements there will be $\frac{1}{2}n(n-1)$ *cost elements* which basically form a $[n \times n]$ symmetrical matrix. This does lead to practical computational problems in a network of reasonable size. However, fortunately enough the configuration of networks are simpler and for all practical purposes are hierarchical with a backbone connecting regionals and the regional backbone connecting locals. The hierarchy could be deep depending on the network.

This allows a very useful simplification. Instead of looking at all the n elements we can concentrate on the m regionals i.e. on the inter-region costs. Generally m is much smaller and the cost computation problem is much more tractable.

For our case study, the medium sized Japanese network, the constituent networks could be grouped into a reasonably small number of regions. We focus on the networks belonging to the 15 WIDE or JAIN regions[1]. A region in this context is the group of regional networks which are connected to the backbone by a common gateway. Thus, they all have the same entry point for wide area traffic. Regions are interconnected by backbone links as depicted in the figure below

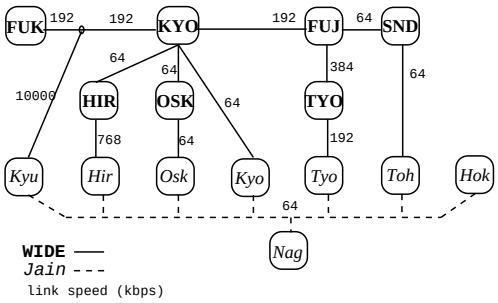

WIDE and JAIN regions.

As a result, the mesh graph with nodes and (weighted) edges can be used for further consideration of the network.

Traffic between regions is routed according to topology and policy rules. Unfortunately, there is no common database revealing all routes. This topology and policy related information or the

[1] There are several other important constituent networks of the Japan Internet. However for the sake of simplicity we have restricted this study to the WIDE and JAIN regions only

resultant IP-image of the network may be obtained from the X.500 or may be derived by running traceroutes. Also, the link properties such as bandwidth, traffic, etc. are expected to be available from the Directory or from the database of the concerned NICs.

4 The Cost Model

4.1 Simple cost calculation

There are several possible definitions of cost. Cost could reflect tariff, speed, policy, ... with varying degree of weightage on the parameters.

For our purposes - namely estimating the *network cost* of traffic - we concentrate on the time the network packet exists in the network. In the present situation lines are either leased or free so the tariff does not come into the picture (The presence of dial-up lines will require added considerations). The cost has two components:

$$cost = time_{at-nodes} + time_{in-transmission}$$

The time at nodes is of the order of micro seconds and is negligible when compared with the time in transmission along lines of max speed 1 Mbps.

Thus,

$$cost \approx time_{in-transmission}$$

As unit of cost we took the time taken for 1 Mbyte to reach from source to destination. This way, the definition of cost is reduced to

$$cost = \sum_{i=1}^{n} \frac{1000}{sp_i}$$

where sp_i is the speed of a single link i in kbps, and n the total number of links involved in the end-to-end connection.

4.2 Total cost calculation

The simplified model wherein the network comprises of 15 regions and the regions are connected by the backbone networks, greatly simplifies the total-cost calculation. It is clear that

$$\begin{aligned} total_cost_{absolute} &= cost_{inter-node} \\ &= cost_{source-region} + \\ &\quad cost_{inter-region} + \\ &\quad cost_{destination-region} \end{aligned}$$

where,

- source-region [destination-region] costs are the costs from [to] the source [destination] node to [from] the backbone router and,
- inter-region cost is the cost between the source-region router and the destination-region router on the backbone.

As far as comparative costs from one fixed starting point are concerned the source-region cost will drop - so that

$$\begin{aligned} total_cost_{compare} &= cost_{inter-region} + \\ &\quad cost_{destination-region} \end{aligned}$$

Thus to carry out the cost calculations we needed a [15×15] symmetric matrix of inter-region costs and an array of intra-region costs. Calculation of this matrix using network configuration information and the above cost model was straight forward.

Table 1 gives an extract of this matrix. Here, the connection costs from one region (FUJ) to all other regions are given.

Table 1: Connection costs as seen from jp-gate.wide.ad.jp (Fujizawa)

Target region	cost
WIDE Fukuoka (FUK)	156
WIDE Hiroshima (HIR)	208
WIDE Osaka (OSK)	208
WIDE Kyoto (KYO)	52
WIDE Tokyo (TYO)	26
WIDE Sendai (SND)	156
JAIN Kyushu (Kyu)	104
JAIN Hiroshima (Hir)	221
JAIN Osaka (Osk)	364
JAIN Kyoto (Kyo)	208
JAIN Nagoya (Nag)	234
JAIN Tokyo (Tyo)	78
JAIN Tohoku (Toh)	312
JAIN Hokkaido (Hok)	234

If we are able to identify the corresponding regions of the client and the server of an application, the relative cost of this connection may be computed. For anonymous ftp, a comparison of servers located in different regions is possible from the very point a user is situated.

For a realization of the idea essentially what is necessary is that

- for each backbone node the cost to all other backbone nodes be given
- for each non-backbone node the cost to the corresponding backbone node be given.

This information is relatively easily provided in the Directory services along with other network

details. The backbone node entries have a cost attribute that is actually a row of the $[n \times n]$ symmetrical matrix (n = number of backbone nodes) and represents the costs from the particular node to each of the other backbone nodes. The non-backbone entries have a cost attribute that represents the cost to a backbone node.

5 Implementations

So far, we have talked about our cost model; in this section, we examine two projects which have which employ cost calculation strategies to select from several servers.

5.1 SoftPages

The Soft Pages Project[1] explores a way to optimize document retrieval by using a "shortest" way approach.

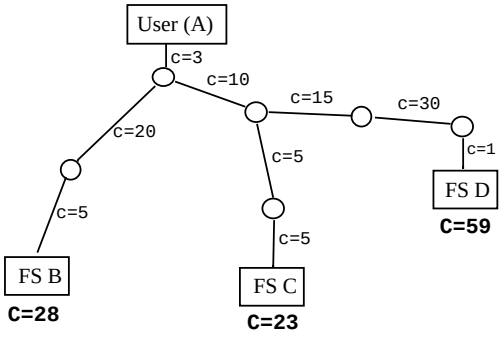

The Soft Pages Project [SPAGE] started with the wish to reduce ftp traffic on crowded overseas links and national backbones. The basic idea is to query for files always on filestores that are nearest/cheapest to the users' site. For that purpose, it is necessary to evaluate distance/cost between the users' host and possible servers. Also, it is necessary to search the contents of the servers.

For example, a user on host A looking for a certain document should under ideal situation carry out the search in the ordered list of fileservers *[C,B,D]*, the ordering being based on the potential retrieval costs.

SoftPages makes use of the network infrastructure information concerning network configuration, location of servers and their contents, the links and their charges, etc.

Once a user has entered the name of the file he/she is looking for (file name can contain wildcards and a path/subdir), Soft Pages is expected to return a host name or invoke the ftp session within a reasonable time. The steps involved are as follows. First, a set of fileservers is chosen. This set contains all servers that will be considered for a filename search. It is possible to generate such a set from the Directory by searching within a certain scope (e.g. "all file servers in Japan"). But a user can also generate his own set of fileservers and thus include some strange sites that might resolve his questions.

Next, according to the connection parameter preference the user has set (speed, cost, etc.) all fileservers of the above mentioned set are evaluated. The required information for the evaluation is gathered from the directory - and in general cached locally.

The cost evaluation will be used to compile a sorted list of fileservers. Generally, Directory enquiries will proceed along this sorted list, starting with the fileserver that has the lowest cost index. If no file name match can be found, the search goes on for the next fileserver (with the next higher cost index). This process will be continued until a file matching the search pattern has been found or information for all fileservers has been scanned unsuccessfully. A SoftPages User Agent hides all this processing from the user.

[1] at Tohoku University, AIC System Laboratories and WIDE

5.2 SmartArchie

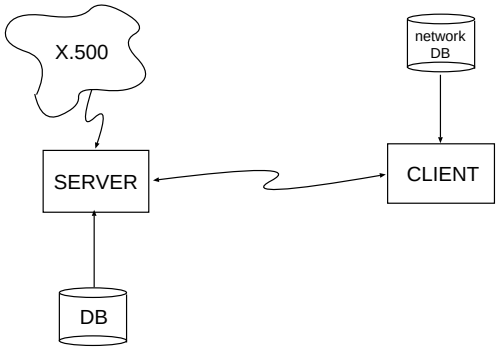

THE SMART ARCHIE

SmartArchie is a very simple improvement on the conventional Archie mechanism [Deutsch]. The conventional archie clients return a list of fileservers to the user. The list of file servers may be sorted according to some very simple ordering - for example based on a given ordering of the first-level domain components (*.jp, .us .uk, .edu .. etc.*). However, this is very inadequate - particularly when one considers the difficulty of associating any preference based on the next level (*.ac, .co .or, .ad .. etc.*) or levels below.

In its simplified form a SmartArchie client works by using the actual information on number of hops and the RTT for the various candidates. In the absence of such information - the client itself may collect the information using some simple network query tools - e.g. *ping, traceroute.*

The results are cached in a local database- to improve performance. The list of servers returned by the Archie server is sorted based on this cost.

A SmartArchie server would use information from the Directory in the same manner as the SoftPages client does to sort the potential servers depending on the cost as seen by the client/user.

Conclusion

FTP traffic, one of the largest consumers of network bandwidth, is taken as a case study of how effectively network configuration information can be used. A workable and scalable mechanism is shown using which a user or an application can evaluate the **cost** factor corresponding to several ftp-servers. By choosing the server with the lowest cost factor, network cost of the retreival is optimized. Cost could pertain to response time, tariff, or any other parameter. A network model based on which the comparative costs of several servers in a network can actually be calculated is shown. Two experimental implementations of the idea are discussed to demonstrate the effectiveness and practicality of the the proposal. The concept can be profitably employed in several other application which utilize ftp mechanisms e.g. World Wide File Server [WWFS].

6 Acknowledgements

The authors would like to thank the researchers at WIDE and Tohoku University who have been very helpful with various pieces of network related information and advice.

References

[ASABA] Asaba, T., K. Claffy, O. Nakamura, H. Kusumoto, T. Kitajima, J. Murai:
An analysis of international academic research network traffic between JP and other nations;
Technical report, WIDE, November 1991.

[MANS] G. Mansfield et. al.: Mapping Communication Networks in the Directory
Computer Networks and ISDN Systems, Vol. 26, No. 3, Nov. 1993.

[Deutsch] Deutsch, P., A. Emtage:
The archie System: An Internet electronic Directory Service;
in: Connexions, The Interoperability Report, Vol.6 No. 2, February 1992

[Satoh] Satoh,H., G. Mansfield:
SOFTPAGES: A wide area distributed file enquiry and retrieval system based on OSI directory service
Internal Report, AIC 1991.

[OSI-DS 22] Yeong, W.:
Representing Public Archives in the Directory;
Networking Working Group document osi-ds-22, November 1991.

[MERIT] Merit/NSFNet Information Services. NSFNET Traffic Distribution Highlights, *(monthly)*, 1993.

[Alex] Cate., V.:
Alex a global filesystem;
Proceedings of the Usenix File Systems Workshop, May 1992.

[UCB] Danzig, P. B., Hall, R. S., Schwartz, M.F.:
A Case for Caching Objects Inside Internetworks;
Technical Report CU-CS-642-93, Department of Computer Science, University of Colorado at Boulder, March 1993.

[SPP] Johannsen, Th., G. Mansfield, K. Takahashi, S. Noguchi:
The SoftPages Project;
Proc of the National Conf. of the Information Processing Society of Japan, October 1992.

[RFC959] Postel, J.B., J.K Reynolds:
File Transfer Protocol;
RFC 959, October 1985.

[SPAGE] The SoftPages Project;
Thomas Johannsen, G. Mansfield, K. Takahashi, S. Noguchi;
Proc of the National Conf. of the Information Processing Society of Japan;
October 1992.

[FTP] A Study on FTP Traffic: Estimation of the Non-Optimality;
Thomas Johannsen, Glenn Mansfield.
International Journal of Network Management, to be published.

[WWFS] World Wide File System, Yamaguchi S. et. al.;
Wide Technical Memo WIDE-TM-92002, July 11, 1992.

Author information

Krishnamachari Jayanthi obtained her Bachelor's degree in 1978 from Madras University, India, in the field of Electronics Engineering. Subsequently she was employed for six years in Tata Consultancy Services, Bombay. In 1988, she obtained her Master's degree in Information Engineering from Tohoku University, Japan. After working for two and half years in a graphics company in Tokyo, she is currently specialising in Network management research in AIC Systems Laboratories, Sendai, Japan. Her areas of interest include pattern recognition, graphics and network management.
e-mail: jaya@aic.co.jp

Glenn Mansfield obtained his Master's degree in 1977 from Indian Institute of Technology, Kharagpore, India in the field of Nuclear and Particle Physics followed by his Masters in Physical Engineering in 1979 from Indian Institute of Science, Bangalore, India. After working with Tata Consultancy Services, Bombay, India as a senior systems analyst for five years he obtained his Ph.D. specializing in Logic programming, from Tohoku University, Japan. He has worked as a research associate in the computer center of Tohoku University for a period of 3 years and is currently a senior visiting researcher at AIC Systems Laboratories, Sendai, Japan. His areas of interest include expert systems, logic programming, computer networks and their management. He is a member of the Internet Society, the IEEE, IEEE Communications Society, and the Information Processing Society of Japan.
e-mail: glenn@aic.co.jp

Yukio Kimura received his B.E. and M.E. degrees, in 1976 and 1978 from University of Electro-Communication, Tokyo, Japan. Subsequently he joined Nippon Telegraph and Telephone Corporation, Japan. Since 1993 he has been working as Manager of the Intelligent Network Management Systems Laboratory at AIC Systems Laboratories, Japan. He is a member of the Information Processing Society of Japan and a member of the Institute of Electronics, Information and Communication Engineers of Japan.
e-mail: kimura@aic.co.jp

Thomas Johanssen received a Dipl.-Ing. degree (M.E.) in electrical engineering from Dresden University of Technology, Germany. His interests include local and wide area computer networks. As research assistant at Dresden he was involved with OSI-LANs. From January 1992 to March 1993, he was working at Tohoku University Sendai, Japan, on network management issues using the X.500 Directory. Presently, he is continuing this research work at Dresden University.
e-mail: Thomas.Johannsen@ifn.et.tu-dresden.de

Yoshiaki Nemoto was born in Sendai City, Miyagiken, Japan, on December 2, 1945. He received the B.E., M.E., and Dr. Eng. degrees from Tohoku University, Sendai, Japan, in 1968, 1970, and 1973, respectively. From 1973 to 1984, he was a Research Associate with the Faculty of Engineering, Tohoku University. From 1984 to 1991, he was an Associate Professor wiht the Research Institute of Electrical Communication and at present he is a Professor at the Computer Center at the same university. He has been engaged in research work on microwave networks, communication system, satellite communication and computer network system. He was a co-recipient of the 1982 Microwave Prize from the IEEE Microwave Theory and Techniques Society. Dr. Nemoto is a member of the Information Processing Society of Japan.
e-mail: nemoto@cc.tohoku.ac.jp

Shoichi Noguchi received his B.E., M.E. and D.E. degree in Electrical Communication Engineering from Tohoku University, Sendai, in 1954, 1956 and 1960 respectively. He joined the faculty of the same University in 1960 and has been a Professor since 1971. Presently Director of the Research Institute of Applied Information Sciences, he has served as the vice President of the Information Society of Japan during the years 1988-1990. He is active in the fields of information science theory, computer network fundamentals, parallel processing, computer network architectures and knowledge engineering fundamentals. He has also been actively promoting the concept of a "Pan Pacific Satellite-based Academic Network". He has authored several books and is the recipient of the IPSJ 25th anniversary memorial prize winning paper award in 1985, the Niwa award in 1986, the Kahoku Scientific Culture award in 1991. He is a member of the Information Processing Society of Japan and the Institue of Electronics Information and Communication Engineers, Japan.
e-mail: nogu@riec.tohoku.ac.jp