Öneri. C.3. S.14.
Haziran 2000. ss.63-67.

A SWITCHING TECHNIQUE AND AN APPLICATION EXAMPLE OF THE ENTERPRISE NETWORK

Turgay KALAYCI

İstanbul Üniversitesi , Bilgisayar Bilimleri Uyg.Arş.Mer. , Yard. Doç. Dr.

Özet: *Bilgisayar ağlarında anahtarlama kullanımının artması ile birlikte "Sanal Ağlar" gündeme gelmiştir. Anahtarlar basit işlemler kullanarak yüksek performans seviyelerine ulaşırlar. Yönlendiricilerden çok köprülere benzerler. Bir anahtar ikinci katman (MAC) hedef/kaynak adreslerini kullanarak anahtarlama işlemini gerçekleştirir. Bu da yönlendirme işleminden çok daha kolaydır. Hub ve yönlendirici içeren bir ağda paketler Hub içerisinde tekrarlanır, hublar arasında ise yönlendirilir. Sanal ağlar içeren anahtarlamalı bir ağda ise paketler sanal ağ içerisinde anahtarlanır sanal ağlar arasında ise yönlendirilir. İyi bir sanal ağ çözümü tüm bir binayı, kampusu veya şehri kapsayabilir. Böylece yönlendirme ihtiyacı düşerken veri çok daha hızlı iletilebilir.*

Bu bildiride, uygulama olarak verilen problemde bazı kuruluşların bölümlerinde (veya katlarda) farklı topoloji kullanılarak yapılan çeşitli kablolama ile Ethernet ve Token Ring ağlarının bütünleştirilmesi sorun olmuştur. Bu ağların omurga teknolojisi olarak 100 mbps FDDI kullanılmıştır.

Problemimizde bu topolojileri birleştirmek için bilinen köprü/yönlendirici tekniği kullanılmıştır. İş hacmi ve bandgenişliği isteği arttıkça ağ problemleri ortaya çıkmıştır. Veri bir yerden bir yere giderken birçok yönlendiriciden geçmektedir. Yönlendiriciden her paket geçişinde belli bir latency söz konusudur. Veri yolu içindeki birden fazla yönlendiricinin sebep olduğu gecikme (latency) bir ağı çarpıcı bir şekilde yavaşlatabilir. Örnek problemimizde bu tıkanıklığı ortadan kaldırmak için söz konusu kurumsal ağda anahtarlama tekniği kullanılarak çözüme ulaşılmıştır.

I. INTRODUCTION

I.1. Mainframe Networks

When large central machines dominated computing, a series of proprietary network architectures evolved. These provided reliable, stable connections between a manufacturer's terminals and hosts. All applications ran on the mainframe, and the great majority of the data that passed across the network was textual.

The most essential elements in a mainframe network were front-end processors and cluster controllers (Figure 1). Groups of terminals were attached to controllers, were connected to the front-end processor through point to point cables (for local connections) or leased telephone lines (for remote connections). The idea of "shared bandwidth" was used for remote connections, with controllers taking turns with the bandwidth in the telephone line. Given the type of data these networks handled, they provided an excellent combination of fairness of access, throughput and cost. The dominant mainframe vendor was IBM. Their networks from this period ran so well that some of them are still relatively unchanged. Most, however, migrated into token ring networks linked with bridges and/or routers, or were completely replaced. Token ring is complex, but it provides many of the same performance characteristics as a mainframe-based controller/terminal network, with the flexibility and bandwidth of a LAN access method.[1].

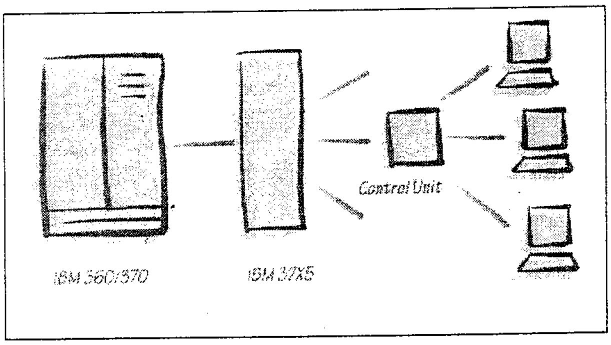

Figure.1. A typical Mainframe Layout

I.2. Minicomputer Networks

As minicomputers became technically feasible and cost-effective, many organizations shifted engineering and business applications to them. Terminal access started in a very basic way, using asynchronous terminals connected directly to a port on the mini. Statistical multiplexers evolved to provide wide area line sharing and error protection. Data PBXs were central to many networks, allowing terminal users to select computers and contend for expensive computer ports (Figure.2).

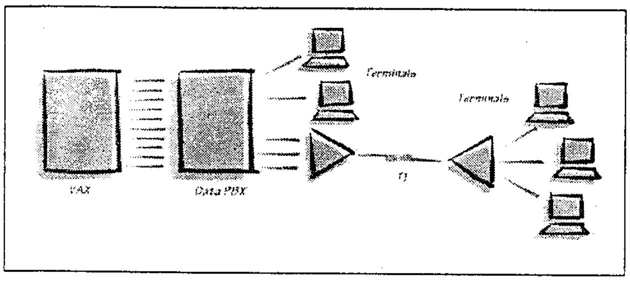

Figure.2. The Mini Computer Layout

A key change in networking had taken place: it was now typical for multiple vendors to provide the various components of a network. This rapidly accelerated the rate of change in networking technology as an increasing number of small manufacturers entered the market.

The largest of the minicomputer manufacturers, Digital Equipment Corporation, was an early leader in a variant of asynchronous minicomputer networks. Existing terminals were connected to terminal servers, rather than directly to the minicomputers or via data PBXs. The terminal servers were connected to minicomputers across a 10 Mbps Ethernet local area network. In this architecture the terminal servers replaced data PBXs, and remote bridges replaced statistical multiplexers.

A similar evolutionary change was made by IBM. The point-to-point cables which had previously attached local cluster controllers to mainframes were replaced with tokenring local area networks.

I.3. Shared-Bandwidth LANs

Terminal-based LANs from IBM, DEC, and a number of smaller manufacturers established LANs as a viable medium for computer networking. But a more fundamental shift occurred with the explosion in personal computers. Desktop computers proliferated throughout organizations, and users needed to share physical sources, such as printers, and transfer files. To meet this need, LAN-based network operating systems emerged. These supported the standalone applications that users already had, and provided a basis for a rapid, steady evolution of LAN-based applications that made use of shared databases (Figure 3).[2].

Fundamental to these networks was the concept of shared bandwidth. Many PCs and other devices were attached to, and took turns with, a single Ethernet segment or a single token ring. Since early PCs had very limited processing capacity, and could only sustain limited bursts of data on and off a network, this process worked quite well. There were two fundamental problems with these networks:

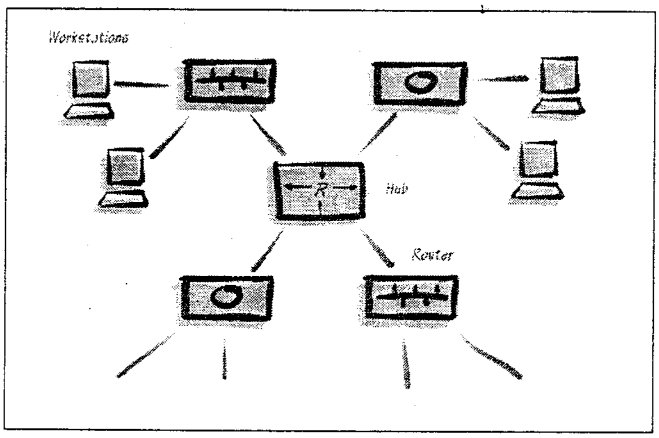

Figure.3. Shared Bandwidth LANs

- The thick cables were awkward and expensive to install, and (in the case of coaxial Ethernet) didn't follow the convenient star topology wiring that had become standard for voice and terminal connections. In a star topology, twisted-pair cables are run from each desk to a wiring closet; usually one or more wiring closets are placed on each floor of a building.

- Bridges were used to connect individual Ethernet segments and token rings together. By their nature bridges propagate certain traffic (broadcasts) to all stations. This worked well in smaller networks, since broadcasts are normally a very small percentage of total traffic. But in very large networks the number of devices became large enough that the broadcasts started to overload the network.[3]

Intelligent hubs solved the cabling problem; they evolved to support both Ethernet and token ring. Every station could use existing twisted pair cabling. Although many stations still shared a segment or ring, the wiring was much more manageable.

Routing, both in file servers and in external units, solved the second problem of excessive traffic by allowing network managers to segment their LANs. Broadcasts were stopped at the router (especially valuable over low-speed wide area links), and user traffic between LANs was controlled. Routers did this by running very complex protocols which allowed them to filter broadcasts intelligently.

I.4. Switching

Advanced desktop computers can now handle throughput rates significantly higher than Ethernet or token ring provides. This is especially true for servers, as corporate information functions migrate to application servers, much higher data rates are needed to effectively support a large number of users (Figure.4).

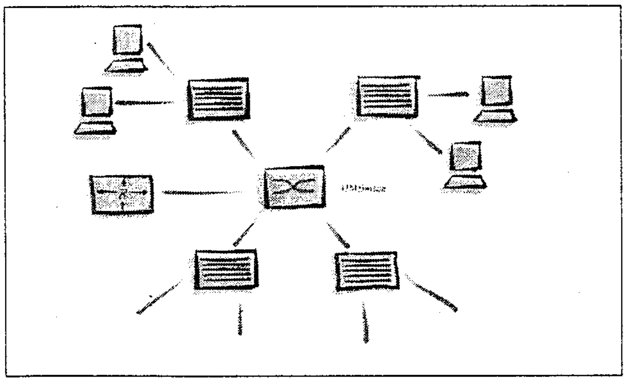

Figure.4. The Switched Network

A new way of providing information to users is taking advantage of this hardware power. In almost every industry and every type of application, data is being presented through images rather than text. The World Wide Web, document imaging, medical radiology, CAD, video training, and pre-press editing are just a few of the applications that are absorbing enormously greater amounts of bandwidth. It takes 8 bytes to transfer the word "airplane". It takes 80,000 bytes to send a picture of an airplane. It takes 8,000,000 bytes to send a simple video sequence of one.

The economic benefits of graphical computing are enormous. Information is understood more quickly and intuitively. Users spend less time learning the application and more time working with information. The use of paper is reduced dramatically or almost eliminated. Work can be shared across a distance which previously required a human to be present. Expertise can be transferred throughout an organization or an industry. Other trends are also overloading existing networks. More and more users are being added to existing networks. Client/server applications are multiplying bandwidth needs.[4] Enterprise application servers and central server farms are moving more traffic outside of local workgroups.

It's becoming obvious that the hub and router technologies developed five years ago were not designed to cope with such a heavy load. A short-term patch is to split up segments and rings. But this solution doesn't scale well to support really high-bandwidth applications, and it's also hard to manage.

A better solution is switching. In a switched network each workstation and server has its own dedicated connection to the network. This means that 20 users, instead of sharing a 10 Mbps Ethernet segment, have 200 Mbps of throughput available to them. And if the switch supports high-speed server connections, such as 100BaseT, ATM or FDDI, data can move rapidly to many workstations throughout a network.[5]

I.5. The Basic Elements of the Switching Revolution - ATM Switching

One key element of the shared-bandwidth LAN has carried over to the switching : Users expect that their suppliers will meet broadly accepted standards. The most important recent standard is Asynchronous Transfer Mode.[6]

Asynchronous Transfer Mode (ATM) has become widely accepted as the standard switching mechanism for future networks. It's being deployed rapidly within local and wide area networking products as well as by wide area public networks. ATM uses small fixed-length cells which keeps latency low rather than variable-length packets. It can combine bursty, information, such as LAN packets, and constant bit rate information, such as voice. Because ATM switching is implemented in hardware, and switches can be meshed (interconnected), it will also scale to support applications which are orders of magnitude larger than today's largest networks support. The question is not whether a transition to ATM will occur, but how quickly. Eventually, networks will be built with a high-speed ATM connection to every desk, but this is unlikely to be widespread in the near future. ATM interface cards are still much more expensive than Ethernet NICs.[7] ATM switch architectures are still evolving; connecting all computers in a campus directly via ATM would create an aggregate rate too high for current switches to effectively support. Many users who want to run ATM at really high data rates, will need to rewire with fiber optic or Category 5 unshielded twisted pair cable. In the next few years most organizations will deploy ATM in an important but delimited way for backbones and for workgroups with special bandwidth needs.[5]

I.6. LAN Switching

LAN switches are the fastest-growing segment of the networking industry. The reasons are simple:

Prices are dropping rapidly, from the 1993 price of $ 1,500 - $2,500 per switched Ethernet port to today's price of $500 per switched Ethernet port. In contrast to routers, LAN switches move data primarily with hardware rather than software. So switches are very fast and provide much greater throughput at a much lower cost. Most LAN switches support lower-speed technologies, such as Ethernet and token ring, and higher speed technologies, such as 100BaseT, FDDI and ATM.[8] Switching's ability to support various LAN technologies immediately increases network throughput without a massive infrastructure upgrade. Users who want to gain the benefits of switched networking today don't have to leap into desktop ATM. Most applications see an enormous increase in throughput by dedicating Ethernet and token ring connections to individual devices, and switching those devices to 100BaseT, FDDI and ATM servers. This means that two kinds of switching are needed. The first type is LAN switching which connects to workstations and servers with Ethernet, token ring and FDDI interfaces. Every device has its own dedicated connection. Connections between two devices attached to the same LAN switch (for example, a workstation and a local file server) are handled locally at very high rates. The second type of switching, in turn, connects those switches to ATM switches. These act as a multi-gigabit backbone fabric. And, using ATM LAN Emulation, a server attached to an ATM switch is able to support workstations attached to a variety of LAN switches.[9]

This optimizes the ATM switch backbone in several ways. Locally switched connections (within a LAN switch) don't have to be translated into and out of ATM cells. Instead, data stays in its native packet form, thereby reducing latency. Local switching reduces the load on the ATM backbone. LAN switches provide a statistical multiplexing effect, smoothing the inputs to the ATM switch fabric; this optimizes the throughput of the backbone network. The combination of LAN switching as the basic access element, and ATM switching as the core fabric, represents a powerful new paradigm. The resulting architecture is easier to manage and provides much better performance than the earlier model of routers and intelligent hubs.

I.7. Virtual LANs

The LAN switch/ATM switch network allows users to be combined into as "virtual LANs," or "VLANs." There are two ways to think about VLANs. A VLAN is a logical, rather than a physical, collection of users. In a router-based network users are identified by their physical location in the network. This is expressed in a network-layer address, which tells the router which physical segment or ring it must send data to. Users in a physical LAN must therefore be in one building or part of a building. A virtual LAN need not be restricted to a floor, building, or even a city. A VLAN is a broadcast domain - all broadcast traffic isolated to a VLAN. Protocols assume that a layer-two ("MAC-layer") broadcast will be seen by all members of group. In a router/hub network this group is the workstations connected to a single segment or ring.[10] In a switched network this group is a virtual LAN. VLANs allow, very large switched networks to be built without the problem of broadcast overloads. They can be a very limited mechanism, simply grouping physical switch ports together, or they can provide a very flexible toolkit, combining existing hubs, routers and FDDI backbones with dedicated switched ports, ATM backbones, wide area networks, and more.

I.8. Routing

In switched networks routing will still be needed, but its role changes. Routers are no longer needed to move all the data around a building or campus. The switching system provides that function at a much higher rate but at a much lower cost. Instead, routing now interconnects VLANs. A VLAN is a broadcast domain, and broadcast domains are interconnected with routing. The routing function in a switched network can be provided in several places. Route servers call be built into LAN switches. This minimizes the amount of equipment in the network and simplifies network management. If modular route servers are available, the degree (and cost) of routing power can be tuned to the needs of the specific application. Route servers built into LAN switches can be either distributed or centralized, depending on the needs of a particular application. Route servers attached to ATM switches minimize management and inter vendor complexity. However, there is an inherent architectural problem.[11] In order to route, an ATM switch has to reassemble cells back into packets and then perform packet-based routing. Most ATM switch vendors have chosen to concentrate on cell switching. The same routers that provided the backbone in the previous hub/switch network can be reconfigured. The router's existing Ethernet, token ring or FDDI ports can be attached to ports on LAN switches, routing between VLANs which span the LAN switch/ATM switch fabric. Or all ATM port on the router can support multiple logical channels, each connected to an ATM Emulated LAN. A third option is to use a trunking protocol operating over FDDI or 100BaseT to achieve this same result.[2]

I.9. Network I The Problem

Varying wiring topologies is a common phenomenon among organizations which originally had small networks servicing individual departments. Eventually companies incorporated those individual networks into the single campus/enterprise networks that are common in today's businesses. Figure 5 shows an example enterprise network which employed Ethernet and Token Ring on individual segments (sometimes on the same floor) and used 100 Mbps FDDI as its backbone technology.

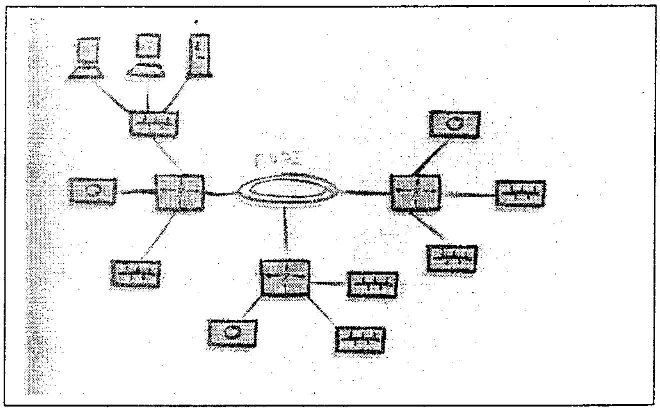

Figure.5. Network 1 - The Problem

Depending on what segment and floor a workstation was on, it could have been running Token Ring at 4 or 16 Mbps or running 10 Mbps Ethernet over twisted pair or coaxial wiring. To solve the problem of integrating these topologies network implementers used traditional bridge/router configurations to weave them into an enterprise network. While this solution worked for a while, the size and bandwidth demands of network applications grew, and consequently so did network problems.

A users began using the connectivity conveniences a network provides, network traffic skyrocketed. They shared large data files on servers, as well as activated peer to peer networking. Often this meant data had to pass through multiple routers to get to its final destination. Each router read the data from its respective segments and forwarded it to another segment on its domain, or released the data to the corporate backbone. Because each router process information in this manner there is a delay while the router processes the packet. This is known as latency. High latency, caused by multiple routers in the data path, can slow a network down dramatically.

In the case of this network, the shared data files were extremely large and low router throughput ground the entire network to a halt while large files passed.

I.11. The Solution

Switching is an inherently versatile technology. When the company who used the network shown in figure 6 decided to upgrade their network, a couple of solutions were presented. They could opt to give every workstation on their network its own dedicated port on a switch or they could replace their current routers with switches and reallocate those routers to other areas of the network such as Internet gateways and WAN gateways. For budget reasons, the organization chose the latter option which left the option open to upgrade to a fully switched solution in the future.

The company realized the benefit of adding switches to its LAN shortly after installing them. The most obvious improvement was in intranetwork traffic. No longer did data have to pass through successive routers which could not handle the demands of a high traffic network. Instead data zipped through the network without slowing down. The switch automatically performs translation so 16 Mbps token rings can transmit data to Ethernet segments or FDDI rings rapidly and effectively.

While most switch ports were used to connect to hubs and MAUs (Media Attachment Unit), some ports were still left for users who needed extra bandwidth.

In addition, the switch solution seamlessly integrated the company's existing 100 Mbps FDDI ring which allowed the company to keep its original backbone in place.

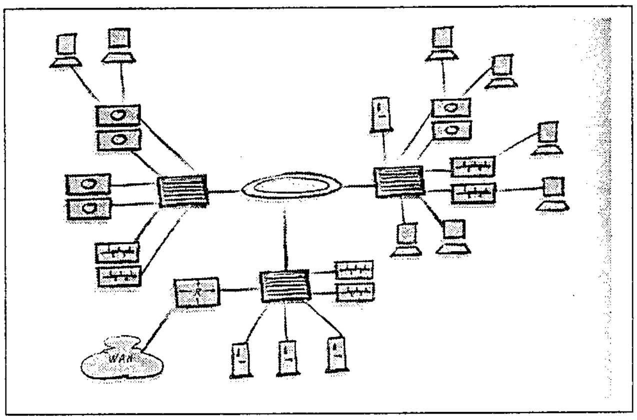

Figure.6. The Solution for Network 1

REFERENCES

[1] Halsall, Fred, "Data Comminications", **Computer Networks and Open Systems**, Harlow, England, Addison-Wesley, 1996.

[2] Mark, A. Miller, **LAN Protocol Handbook**, ISBN 1-55851-099-0.

[3] Worsley, Debra J.; Tokunbo, Ogunfunmi, "Isochronous Ethernet - An ATM Bridge for Multimedia Networking," **IEEE Multimedia**, January-March 1997.

[4] Cypher, David; Shukri, Wakid. "Standardization for ATM and Related B- ISDN Technologies," **StandardView**, 1:1 (40-47), September 1993.

[5] Chaney, Tom; Fingerhut, J.Andrew; Flucke, Margeret; Turner, J.S., "Design of a Gigabit ATM Switch", **Proceedings of Infocom**, 1997.

[6] Kalaycı, Turgay "Neden ATM?", **TBD 13. Ulusal Bilişim Kurultayı Bildiriler Kitabı**, Eylül 1996, ss. 215-220.

[7] Chao, H.J., "An ATM Queue Manager Handling Multiple Delay and Loss Priorities", **IEEE/ACM Trans. On Networking**, vol.3,no.6, Dec.1995.

[8] Lynch, Daniel C.; Rose, Marshall T., **Internet Systems Handbook**, ISBN 0-201-56741-5.

[9] Craig, Hunt, "TCP/IP Network Administration", **A Nutshell Handbook**, ISBN 0-937175-82-X.

[10] Comer, Douglas E.; Stevens, David L., **Internetworking with TCP/IP**, Volumes 1 and 2 ISBN 0-13-468505-9 (V.1), ISBN 0-13-472242-6 (V.2)

[11] Carpenter, Tamra J.; Cosares, Steven; Saniee, Iraj, "Demand Routing and Slotting on Ring Networks," **DIMACS Technical Report 97-02**, January 1997.