Network Behavior and Quality of Service in Emerging Applications [1]

Bharat Bhargava
Department of Computer Sciences
Purdue University, W. Lafayette, IN 47907, USA
bb@cs.purdue.edu

Abstract

The performance of network and communication software is a major concern for making the emerging applications in a distributed environment a success. Emerging applications that we consider in this paper are transaction processing (for financial institutions or electronic commerce), digital library (including web search), video conferencing, and finally stock trading. The quality of service in each case can generically be measured by response time, throughput, reliability, timeliness, accuracy, and precision. We will present experimental data that gives an idea of communication behavior and how it impacts the quality of service in each application. Finally some ideas for dealing with anomalies such as adaptability will be proposed. We are conducting a series of experiments that will lead in the development of policies for adaptability at the application, system, and network layer to meet the quality of service requirements. Next we study the impact of network constraints in determining the quality of service that can be guaranteed to the user. Based on these experiments, we identify guidelines and expertise that will allow the applications and network to meet the quality of service requirements at all layers.

1 Basic Concepts of Quality of Service Control

Quality of service control requires an understanding of the quantitative parameters at application, system, and network layers. The interplay among these parameters needs to be evaluated by building an application such as video conferencing on top of a distributed system. The measurements and experiences will provide rules for controlling the QoS parameters for both application and network.

1.1 Quality of Service Layering

There are several efforts which try to formalize the QoS requirements for distributed multimedia systems [23]. There is no global agreement on the precise meaning and formal definition of QoS. We partition the QoS parameters into two subsets, namely application-dependent parameters and application-independent parameters. The former are high-level QoS parameters and the latter are low-level QoS parameters. There is a strong correlation between the parameters in these two subsets. For instance, frame rate is strongly related to the throughput in the network layer. We classify QoS parameters into three layers: user parameters, application parameters, and system parameters as shown in Table 1. The system parameters can be classified into two categories: network and operating system parameters and device parameters.

Application parameters describe requirements for application services and are specified in terms of media quality and media relations. Media quality includes source/destination characteristics such as media data unit rate and transmission characteristics such as response time. Media relations specify relationships among media, such as media conversation, interstream synchronization, and intrastream synchronization.

[1]This research is supported by NSF under grant number NCR-9405931 and a NASA fellowship

QoS Layer	QoS Parameters
Application	Frame Rate
	Frame Size/Resolution
	Color Depth
	Response Time
	Presentation Quality
System	Buffer Size
	Process Priority
	Time Quantum
Network	Bandwidth
	Throughput
	Bit Error Rate
	End-to-End Delay
	Delay Jitter
	Peak Duration
Device	Frame Grabbing Frequency

Table 1: Quality of Service Layers

System parameters describe communication and operating system requirements that are needed by application QoS. These parameters are specified in quantitative and qualitative terms. Quantitative criteria are those that can be evaluated in terms of concrete measures, such as bits per second, number of errors, task processing time, and data unit size. Qualitative criteria specify expected services, such as interstream synchronization, ordered delivery of data, error recovery mechanisms, and scheduling mechanisms. Specific parameters can be connected with expected services. For example, interstream synchronization can be defined by an acceptable skew relative to another stream or virtual clock.

Network parameters are specified in terms of network load and network performance. Network load refers to ongoing traffic requirements such as interarrival time. Network performance describes the requirements that must be guaranteed, such as bandwidth, end-to-end delay, and jitter. The network services depend on a traffic model (arrival of connection requests) and perform according to traffic parameters such as peak data rate or burst length. Hence, calculated traffic parameters are dependent on network parameters and are specified in a traffic contract. Device parameters typically specify timing and throughput demands for media data units.

2 Distributed Transaction Processing in an Internetwork

The communication software has been deemed as a crucial component of Distributed Transaction Processing (DTP) software[21, 20]. Various components of a DTP system rely on this for high-performance and acceptable QoS. Atomic commit protocols require the exchange of control messages among the participating sites for ensuring the global commit/abort of an active transaction. Replication control protocols require the transmission of both the data items and the control (e.g., votes) messages for maintaining the consistency of replicated copies. Some distributed deadlock detection algorithms exchange information periodically to determine and break the cycles in distributed "wait-for" graph. Monitoring services such as surveillance controller [14] sends *I_am_alive* messages for node/link failure detection, to all the other sites. The control messages although small (usually ranging in tens of bytes), are frequently exchanged. For some components, the frequency of these exchanges will determine their effectiveness. For example, in the deadlock detection component, the more often the messages are exchanged, the more prompt will be detection of a cycle. Overall, the DTP demands high performance and reliable communication. These requirements become more crucial and difficult to satisfy in wide-area network environment.

The performance of the communication software is largely dependent on the underlying communication media. Networks with different technologies and characteristics have been merged by the internetwork connections. Thus, the communication network spanning large number of geographically dispersed hosts will vary in speed, reliability, and processing capability. The range of these parameters across networks is growing [24]. For example, a distributed system spanning both ATM and Ethernet networks has bandwidth variations between 145 Mb/s to 10 Mb/s.

We have experimented with the transaction processing on the Internet using different protocols for concurrency and atomicity. We have attempted to understand the impact of multi-programming levels on these protocols in this "new" environment. We report the results of the performance evaluation of these protocols in the WAN environment.

2.1 Communication in the Internet and its Impact on Distributed Transaction Processing

To study the performance of DTP in WAN environment, the need for understanding the communication latencies and packet losses in an internetwork is essential. We have conducted the experiments to study the performance of message delivery in the Internet. A complete illustration of the experiments can be found in a technical report [7].

The performance characteristics between LAN and WAN are significantly different:

1. LANs are usually fast and very reliable. The communication delay and failure pattern are uniform and small for all hosts. In the Internet, error rates are higher, communication delays large and non-uniform [24]. The links between two sites in the Internet can range from 56KB/second (lease line) to 45MB/second (T3 link). The number of hops (gateways) between two sites can vary from a few to a very large number.

2. A LAN does not require routing. Even in multi-LAN environment, routing is simple and almost fixed. In the Internet, source routing is dynamic. The routes between sites can change from time to time [18].

3. A LAN consists of a few hosts (e.g., around hundred maximum) under the same administrative unit. Internet connects hundreds of thousands of hosts all across the world, spanning several organizational boundaries. Thus, it is subject to conflicting administrative policies and different usage characteristics.

2.1.1 Communication Performance

Based on our understanding of the Internet, we identified the three factors that are most important in assessing the performance of message delivery: the physical connection, the size of message, and the cross-traffic. The physical connection between two sites includes the distance, the type of links, and the number of hops (gateways). It determines the lower bounds on the communication latency for a message delivery. To establish the connection, we studied the performance across different sites on the Internet. We examined the effect of the size of the message transmitted to its transmission time. The delays and losses due to cross-traffic were quantified but a useful relationship with the communication performance could not be established due to uncontrolled and dynamically varying usage of Internet by various foreign agents. Also, there is no easy way to determine the cross-traffic on a particular link. One way to circumvent the problem is to determine the traffic pattern at large, for the Internet. We suspected that it would be a function of the working hours, as the Internet usage determined the traffic pattern. To verify this, we examined how the time of the day and day of the week affects the message delivery performance in the Internet.

To summarize, we have conducted measurements in three dimensions: the *time* dimension by periodically repeating the experiments, the *site* dimension by repeating experiments with different sites, and the *size* dimension by varying the message sizes. We are interested in two performance measures: the round-trip time of a message and the message loss rate. In our DTP model, round trip time is the time for a site to send a request message to another site and receive a reply message back. Message is said to be lost when the transport service of the Internet fails to deliver the message in time. This is a important parameter for us, since a lost message not only blocks or aborts the transaction but also increases the contention for the shared data, such as the indices.

Our experiments involved over 2000 sites and 500 networks in the United States. We probed the Internet with ICMP and UDP messages periodically and collected the data [7]. Based on these measurements, we can make the following observations.

- We observed that there is a large variation in parameters such as communication delay and message loss. The variations exist in two dimensions: along the time axis and across the networks.
- We observed that the time of day has strong influence on the message delivery. The message loss rate is much higher in the noon working hours, and much lower in the early mornings. The round-trip time for a message, on the other hand, does not have a strong correlation with the time of the day, except for the hourly peeks. This we believe is caused by the hourly jobs scheduled to run on gateways.
- We observed that the message delivery has an unbalanced performance across the wide area networks, although most of the hosts reported within 400ms round-trip. The *"clustering"* effect in the Internet is also observed. The communication between a site and many different sites on another local network has similar performance, which can be represented by any host on that network. Therefore, the latency between two networks can be used to estimate the communication delay between two hosts in these two networks.
- Finally, we observed that for small messages that can fit in an IP data gram without fragmentation, there is an approximate linear correlation between the transit time and the size of a message. However, the message loss is not affected by the size.

2.1.2 Impact on Distributed Transaction Processing

The performance analysis of communication in the Internet, reported in the previous section has a significant impact on distributed transaction processing on the Internet.

The time to deliver a transaction message in the WAN is a number of magnitude longer than in a LAN. While it takes only a few milliseconds to deliver a message in a LAN [1], on the Internet it is several hundreds of milliseconds to send a message across the continent [13]. This means that a transaction stays longer in the system, implying the larger lock holding time for data items, if two phase locking is used for concurrency control. This leads to increased contention to the database, affecting the throughput adversely.

The already difficult problem of finding a "good" value for *timeout* in LAN is further aggravated in WAN environment. Timeout is used in DTP systems to trigger special treatment for the transactions that cannot be finished in time. The timeout value usually equals a constant multiplied by the number of read/write operations in the LAN environment. In a WAN, this flat timeout rate is not adequate. As the CPU and disk I/O performance improves, most of the time spent for a transaction is in the waiting for the messages to be delivered. Thus, the timeout value for a transaction must be dependent on the number of remote messages and their destinations.

Autonomous control over LAN allows modification to the communication software improving the performance of DTP. [9] discusses many changes, such as physical multicasting, light weight protocols, etc that can be affected. Physical multicasting is not supported by all WANs. Direct control passing or memory mapping may not have a significant impact, because the message delivery latencies may cause a performance bottleneck. Unless dedicated links or special networks are adopted, one can not do much to the shared public WAN such as the Internet. The performance of message delivery is determined by traffic and various other factors beyond the designer's control. Therefore, the focus of improving communication has to be shifted toward reducing the number of messages exchanged in DTP.

The mechanism to handle message loss in DTP has to be changed. In a WAN, message loss rate is much higher. The percentage of message loss is usually 5%, sometimes as high as 30% [7]. Frequent transaction abort and restart caused by message loss will drastically degrade the overall performance of DTP. Transport protocols that have higher degree of reliability should be considered.

DTP algorithms must be able to adapt to the high variations in parameters such as communication delay and the message loss to different sites. For example, the quorum consensus replication control algorithms should consider the dynamic performance of each links. Such site-to-site estimated performance data are stored in a matrix structure, called cost matrix or weighted adjunct matrix. However, the values are not pre-defined and fixed but time varying, and cannot be specified as a function of geographic location of the sites. In consequence, the algorithms (such as distributed query optimization) that use static cost matrix are no longer adequate. The communication system need to collect the performance data periodically to update these cost matrices.

Surveillance facilities have helped in early detection of site and link failures and repairs by exporting an up/down vector to the DTP algorithms [14]. In the WAN environment, modeling the communication as an up/down vector is not sufficient. Early detection of changes in communication performance, such as latency and message loss rate must also be considered. This will improve the performance of adaptable DTP algorithms such as the quorum consensus replication control protocol.

3 Digital Libraries

Digital libraries provide online access to a vast number of distributed text and multimedia information sources in an integrated manner. Providing global access to digitized information which is flexible, comprehensive, and has easy-to-use functionalities at a reasonable cost has become possible with developments in areas such as databases, communications, multimedia and distributed information systems. Digital libraries encompass the technology of storing and accessing data, processing, retrieval, compilation and display of data, data mining of large information repositories such as video, audio libraries, management and effective use of multimedia databases, intelligent retrieval, user interfaces and networking. Digital library data includes texts, figures, photographs, sound, video, films, slides etc. Digital library applications basically store information in electronic format and manipulate large collections of these materials effectively.

Digital libraries typically deal with enormous quantities of data. The National Aeronautic Space Agency (NASA) has multiple terabytes of earth and space science in its archives. NASA is going to launch the Earth Observing System (EOS), which will collect a terabyte a day. Video-on-Demand systems have thousands of video clippings. Almost every organization has repositories of old versions of software and business related data. The CORE project, an electronic library of Chemistry journal articles deals with 80 Gbytes of page images [11]. The University of California CD-ROM information system in 1995 consisted of 135 Gbytes of data [17]. The ACM digital library, functional since July 1997, provides access to about 9,000 full text articles and several tables of content pages and bibliographic references.

3.1 Digital Libraries in a Distributed Environment

Digital libraries are distributed over national and international networks and their infrastructure is inherently distributed [4]. Existing repositories are distributed, and the data needs to be shared by many users. Information processing is distributed, and particularly, the user's queries can be so complicated that the process of information retrieval requires multiple rounds of message exchange between users and the various servers of the information system. These factors result in communication behavior being one of the most important parameter for providing QoS. Along with other components it contributes to the cost of providing digital library services. To keep the cost reasonable a digital library designer has to be aware of the communication overheads and the possible solutions to reduce these overheads.

In a wide area environment, the anomalies (failures, load on the network, message traffic) affect the communication of data. The multiple media of digital library data introduce further complexity since each media has its own communication requirements. The current network technology does not provide the bandwidth required to transmit gigabytes of digital library objects. The cost of access in the context of communication and networking is the response time required to access digital library data. A digital library user might have to wait for several minutes to receive the data due to bandwidth limitations.

We study communication in a distributed digital library at the information systems layer. The underlying information transfer mechanisms can be information protocols such as Z39.50 or HTTP.

In Table 2 we give a few estimates of the size of digital library data objects to give an idea of the size of packets in a digital library application. The figures do not represent an average or generalized size of data items of a particular media, but a sample of possible data item sizes.

Media	Size (Mbytes)
Text (An encyclopedia section)	0.1
Image (A NASA image)	0.8
Video (Uncompressed)	48000
Audio (2 minute speech)	1

Table 2: Examples of Digital Library Data Item Sizes

3.2 Communication in Digital Libraries

Figure 1 illustrates the round trip times in a LAN and MAN. The files under observation range from 6 Kbytes to 496Kbytes. In a LAN, the round trip times range from 722.84ms to 1316.82ms. In a MAN, the round trip times range from 749.41ms to 2738.63ms. We can make two observations here. The difference between a LAN and a MAN for a file size of 6 Kbytes is only 26.57ms. On the other hand, the difference in round trip times for file size 496Kbytes is 1421.81ms. The second observation is that the difference in round trip times in a LAN environment between files of sizes 6 Kbytes and 496 Kbytes is only 593.98ms. The same difference in a MAN environment is 1989.22ms.

Figure 2 illustrates the round trip times in a WAN. When compared to LAN and MAN the round trip times rise sharply as file size increases and as the number of the hops increases. The difference between the largest size and the smallest size is as high as 23811.691ms.

3.2.1 Text Data Retrieval

Text data transmission has to be lossless. Every character is important and random loss of some bytes will result in messages which appear scrambled. Lossless compression techniques will result in compression ratios of 2:1 which will reduce the transmission time by 50%.

Digital library applications can use semantic knowledge of the data to reduce communication time. For example, a document can be represented hierarchically in varying degrees of detail. As shown in If the user has specified to the system the level of detail which is sufficient for her, the system can choose the representation which will reduce communication time and yet satisfy the application.

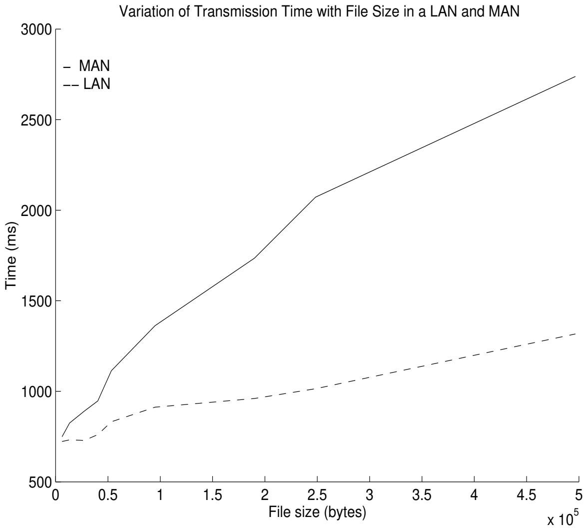

Figure 1: Variation of Transmission Time with File Size in a LAN and MAN

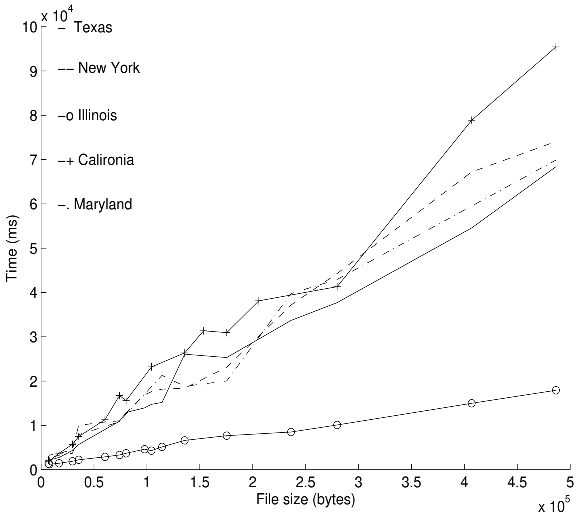

Figure 2: Variation of Transmission Time with File size in a WAN

3.2.2 Video Data Retrieval

Video data can be treated as a stream of images. The techniques described above for the efficient transmission of images apply to video data. Since video data is continuous, there are some issues specific to video data which are addressed in this section. The approach developed in our laboratory is based dynamic adaptability of the quality of video transmission to the bandwidth conditions.

Adaptable Transmission of Video Data

Video transmission applications have to maintain a constant frame rate. The current TV frame rate is about 30 frames per second. The variation in available bandwidth does not allow this frame rate to be maintained without reducing the amount of data by trading off some aspects of video quality. We have identified four aspects of video quality that can be changed to adjust to the available bandwidth:

- Color Depth Compression: Color video can be compared to gray-scale video to reduce the size of the data since gray-scale pixels require fewer pixels to encode than color pixels.
- Frame Resolution Reduction: Replacing every 2x2 matrix of pixels by one pixel reduces the size of the video frame by a factor of 4. The image is reconstructed at the receiver to keep the physical size of the frame unchanged. Since the resolution reduction process is lossy the receiver gets a frame which is an approximation of the original.
- Frame Resizing: The frame size is changed to reduce the size of the data. For instance, reducing the frame size from 640x480 to 320x240 reduces the bandwidth requirement to 25% of the original.
- Codec Schemes: Different coding schemes have different compression ratios. Typically, schemes with high compression ratios require more time to compress but the smaller compressed frames can be transmitted more quickly. If the bandwidth available is extremely limited it might be worth while to reduce the communication time at the cost of computation (during compression) time.

Our research group has conducted several detailed experiments to test the feasibility of the above ideas and have come up with a framework to determine the parameters of quality that should be used for video transmission. The framework allows the quality of video transmission to adapt according to the bandwidth available. For further details the reader is referred to [15].

3.3 Communication Bottlenecks

In this section we will identify the different factors which contribute to the communication overheads. The categorization and understanding of the different factors will help us in finding solutions to overcome the bottleneck.

Communication delays are caused by the following factors:

- Physical capacity limitation of the communication link: The different physical media have an inherent capacity limitation.
- Technology: Some of the currently available network technologies are circuit switching, packet switching, frame relay, cell relay, FDDI, and ATM). Some of the technology specific factors which influence communication delay are: cell setup delay, packet formation time, total lost frames/packets/cells, load balancing/load sharing limitations, masking, filtering, and forwarding rates, error detection and correction efforts, level of redundancy etc.
- Number of hops: The number of hops between the sender and receiver gives a rough estimate of the network distance between them. Every time the data is forwarded by a router, it is referred to as a hop. At each hop there is delay due to speed of hardware and software interfaces, memory

and buffers, address database look-up, address verification, processing, filtering and forwarding of packets, frames and cells etc.

- Traffic: The physical capacity which is bounded above has to be shared among different applications. The bandwidth allocation scheme determines the network bandwidth allocated to a given application. There are several bandwidth allocation schemes and over a public network such as the Internet they follow a 'fair' policy which ensures that no application is deprived a share of the network bandwidth. Consequently the network bandwidth available for existing applications is reduced when a new application requests bandwidth.
- Buffer limitations: The buffer limitations at the nodes at either end of a communication path and the routers on the communication path also contribute to the communication delay. The buffer might not be able to store all the packets which arrive and hence some packets are dropped. This results in re-transmission (in a lossless protocol such as TCP) and consequently more contention for existing network bandwidth.
- Out-of-sync CPU: CPU speed is much slower than network bandwidth. Packet or frame or cell processing functions such as packet formation, address lookup, instruction execution, buffer filling time and error checking have their speed bounded by the computation power of the CPU.

3.4 Quality of Service for Digital Libraries

Quality of service (QOS) specifications are used by distributed multimedia systems to enable applications and users request a desired level of service. The system attempts to satisfy the specifications, and if that is not possible due to resource availability restrictions, the application can enter into a negotiation with the system. During the negotiation process the QOS specifications are changed so that the system can meet the requirements. The process might continue over several iterations.

An example of QOS parameters can be found in video-conferencing applications. This is a real-time application and needs a guaranteed supply of system resources to sustain a uniform level of performance. Some of the parameters are loss rate, throughput, frame rate, response time and presentation quality. We believe that similar QOS parameters should be associated with a digital library system. This will allow the application or the user to negotiate with the system and arrive at a set of values for the parameters which both satisfy the user and can be supported by the system. The application can trade-off some parameters in exchange for others.

- **Accuracy of Information:** Digital library queries are satisfied not by exact matches as in traditional databases, but by similarity matches. The accuracy of the match between the query and the retrieved data item can be specified as a QOS parameter.
- **Data Comprehensiveness:** Some components of the item can identified as not required to reduce system and communication requirements and response time. For instance, a video clipping can be retrieved without the audio component. A text document can be retrieved without the images. Components with a higher computational and communication overhead such as multimedia data can be ignored if they are not required to satisfy the query.
- **Presentation Quality:** Data items can be presented at different levels of quality. Visual data items (images and video) can be presented at different resolutions, different color depths, different sizes, and different codec schemes. Audio can be presented at different sampling rates. The presentation quality can be a user specified parameter. An application might prefer a lower quality data item if the response time can be lowered. For example, a K-12 student would be satisfied with a lower resolution medical image than a medical student.
- **Response Time:** We define response time as the time between the instant the user submits a data retrieval request to the time the data appears on the screen. This can be several seconds or

even minutes. Several applications would like tradeoff the quality of the data, accuracy of data, precision, and recall in exchange for a lower response time.

The application or user can specify the different parameters desired. Upper and lower bounds can be used to express acceptable situations. From a communication point of view, the goal is to minimize response time and maximize accuracy of information, precision and recall of the data retrieved, presentation quality and comprehensiveness of the data.

4 Video Conferencing

Video conferencing systems (VCS) have become practical in commercial and research institutions because the advances of technologies in networking and multimedia applications [22]. A video conferencing session involves multiple parties, possibly geographically interspersed, which exchange real-time video data. However, anomalies such as site failure and network partitioning affect the effectiveness and utilization of the communication capabilities. Video conferencing systems [22] lack the ability of dynamically adapting themselves to the variations in the system resources such as network bandwidths. In VCS, changes in parameters such as frame sizes, codec schemes, color depths, and frame resolutions can only be made by users *interactively*. They cannot be made *automatically* based on the system measurements of currently available resources. We need to limit the users' burden in keeping the system running in the most suitable mode to current environment and make it possible to provide the best possible service based on the status of the system. Incorporating adaptability [5] into a video conferencing systems minimizes the effects of the variations in system environments on the quality of video conference sessions.

4.1 Adaptability in Video Conferencing System

A video conferencing system should provided some policies and mechanisms to make it adaptable to the anomalies based on the available resources. The advantages of the adaptability schemes for VC system includes:

- **Heterogeneity:** A VCS will adapt to heterogeneous environments. That is, a video conference session can be held on different hardware platforms and different networks.
- **Scalability:** A VCS will adapt itself as more users, more sites join a video conference in progress.
- **Anomaly Management:** A VCS will adapt to anomalies and degrade gracefully when available resources decrease or become unavailable.
- **Resource Management:** A VCS can make efficient use of resources like storage, CPU time, and communication bandwidth.

The basic idea for achieving adaptability for video conferencing system is to trade some aspects of video quality for others. For example, the frame rate decreases as the available network bandwidth drops. Since the smoothness of a video session is sometimes more important than any other aspects of video quality, we may have to maintain a reasonable frame rate during a video conference session even though the network performance becomes degraded. To achieve this, we have to sacrifice some aspects of video quality, such as color or resolution of video frames.

Adaptability can be achieved ether by the user intervention or by the system itself. If user intervention is required, the system accepts inputs from the user and change the level of service or some system parameters according the the new specification. If adaptability is to be done by the system automatically, the distributed control system periodically measures the available resources and supply these parameters to the video conferencing system. Based upon the current parameters, the video conferencing system reconfigure itself in a user-transparent way to provide the best possible service based on some user-specified criteria that must be satisfied.

4.2 Quality of Service for Video Conferencing

Timeliness, Accuracy, Precision (TAP) can together form a good criteria for QoS. Timeliness is defined as "when an event is to occur". Maintaining it means meeting a deadline. Accuracy is defined as "the degree to which the output conforms to the semantics and contexts of the applications". Maintaining it means guaranteeing the correctness of the data. For example, lossy compression algorithms cause the loss of accuracy. Precision is defined as "the quantity of information provided or processed". Maintaining it means maintaining the amount of data being processed or transmitted over the network. For example, the number of frames per session, the number of pixels per frame, and the number of bits per pixel are parameters used to describe the precision of a video conferencing session.

TAP cannot be maintained at the highest level simultaneously during anomalies. We must trade among these attribute values through experimental studies [2]. The policies to trade among these attributes has been developed as follows:

- Maintaining Timeliness when Bandwidth Decreases
  - Reduce frame size (The accuracy is maintained unless the frame size is below a certain value).
  - Reduce frame resolution (Both accuracy and precision are reduced).
  - Dither color frame to black and white.
  - Compress color depth.
  - Switch to a codec scheme that has a higher compression ratio (Side effect: CPU utilization increases. This can be compensated by frame resizing and resolution reduction).
- Maintaining Accuracy when Bandwidth Decreases
  - Switch to a lossless codec scheme with reduced frame size.
  - Dither color frame to black and white.
  - Compress color depth (compress Y and UV no more than 2 bits each).
  - Do not use lossy codec schemes.
  - Do not reduce frame size or resolution by a big factor.
- Maintaining Timeliness when CPU Utilization Increases
  - Switch to a codec scheme that requires less computation (usually with lower compression ratio).
  - Reduce frame size.
  - Dither color frame to black and white.
  - Do not compress color depth.
  - Do not reduce frame resolution.
- Maintaining Accuracy when CPU Utilization Increases
  - Switch to a lossless codec scheme
  - Reduce frame size.
  - Dither color frame to black and white.
  - Do not compress color depth.
  - Do not reduce frame resolution.
  - Do not use lossy codec schemes.

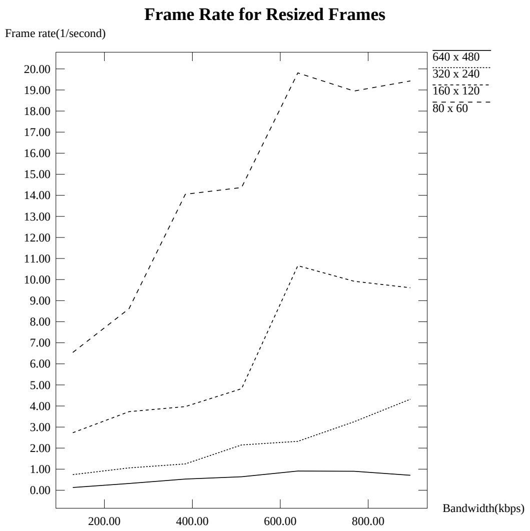

Figure 3: Frame Rates for Resized Frames

4.3 Experimental Setup:

We chose Network Video (NV) [12], with enhancements to incorporate adaptability and recording feature, as our testbed to conduct performance studies.

The platforms for the experiments include a Sun Sparc 10 station and a Sun Sparc 5 station connected in a LAN environment, and a couple of video cameras. The Sun workstations are running Solaris 2.3 operating system.

- Experiments on Resolution Reduction and Frame Resizing

  To see how the frame resizing can be used for adaptability for video conferencing, we conducted some experiments to measure the frame rates for frames of different sizes (Figure 3). For a particular frame size (e.g., 640 x 480), as the available network bandwidth decreases, the corresponding frame rate also decreases. This will result in loss of continuity and smoothness of video presentation at the sender as well as remote sides. In such a situation, the VC system would adapt by changing to smaller frame size to maintain (or even improve) the original frame rate. At present, the system supports only four discrete levels of frame sizes. Thus, the frame rate may be changed (improved) when the system changes to a level of service with less bandwidth requirement. In the future, we plan to provide more levels of frame sizes, which will allow the system to adhere (very closely) to the current operating frame rate while reducing the network bandwidth requirements at the same time.

  To implement the dynamic frame resizing and resolution reduction, we manipulated data before it is encoded and after it is decoded. Though it takes extra CPU time, our results show that this overhead is tolerable. We computed the average time percentages for encoding and decoding in processing one frame. In frame resizing, The combined time percentages for encoding and decoding are more than 50% of the overall processing time of one frame, which means that encoding and decoding are the most expensive parts in video conferencing. When the compression factor value become 2 or 4, the combined time percentages for encoding and decoding drop dramatically; they are no longer the most expensive parts of processing in video conferencing. Instead, the video transmission becomes the most expensive part of processing in video conferencing. Similarly, in resolution reduction, the time percentage for encoding decreases when the frame size is decreased. However, the time percentage for decoding when compression factor is equal to 2 or 4 increases compared to that when compression factor is 1. This is due to the extra computation overhead involved in the resuming the frame size to the original one. But the combined time percentages of encoding and decoding over overall processing time are only slightly bigger than that for original NV.

5 Infrastructure for Communication Experiments

We have developed an extensive environment for conducting experiments in the Raid laboratory at Purdue along with necessary benchmarks and tools. We have conducted scientific experiments in a variety of subjects: communication experiments for distributed applications [26], network communication measurement experiments [25], adaptability experiments for distributed systems [15], etc. with various of communication softwares [8]. The laboratory has a network of Sun workstations running the Raid distributed system [6] and a variety of communication libraries. We have built experimental infrastructures with facilities such as WANCE tool [26], AVC tool [3] and Active Gateway [16].

WANCE is a tool for conducting experiments over WAN with processing control in a LAN environments. It facilitates experimentation in the wide area environments without the needs to obtain accounts on other administrative domains.

AVC tool adds the adaptability to the NV video conferencing tool, originally developed at Xerox PARC [12]. It can dynamically adapt to network bandwidth changes while at the same time maintaining reasonable QoS of the video conference session perceived by users. We have integrated a network bandwidth measurement tool [10] into our AVC tool. The network bandwidth measurement tool can measure the currently available bandwidth of a network connection periodically and send the information to the AVC control module.

In addition, we have built a facility, called active gateway, for experimentation of multimedia communications using Active Networks [16]. We have built this facility on top of our campus IP network environment. Active gateway flexibly supports network programmability, ubiquitous control, dynamic policy enforcement, and traffic control in applications.

6 Electronic Trading

In a new emerging applications of electronic commerce there are transactions involved. Examples of such applications are lending institutions which charge on a per day basis or charge for downloading documents such as the ACM and IEEE digital libraries. Supporting payment by a client brings in the issues of security during the financial transaction.

Security can be enforced by authentication or encryption. Authentication has a communication overhead. It involves a lengthy exchange of information between the client and server such as keys before the secure channel is set up. Encryption has a computational overhead. If encryption is used only for small data messages used in a financial transaction, then the overhead is acceptable. But if huge multimedia data items are encrypted, along with the compression and decompression routines the encryption and decryption routines add a huge overhead to the data retrieval process [19].

Electronic trading is the most exciting but financially appealing applications. In electronic trading, there are several overheads. They are computation of algorithms, particularly encryption; I/O time for database access from various files; and the communication time for servers are involved in executing an order. For example, if the user wants to buy a particular stock, the real time quote has to be provided by the quote.com service. The user has to go the server at her machine to a server at the broker's site and finally to a server that has the actual information. In some cases, there could be as much as twenty message exchanges involved due to the additional need for authentication, reconfirmation, and seeking input from the user.

The process of executing a trade electronically is very similar to the process of trading in person or via phone. In person the additional overhead is in physically going to the broker' office. In trading over phone the overhead is in getting a busy signal or being put on hold. It is however rare that using these two methods, the user can be dissatisfied since a human broker is involved. So the quality of service is acceptable even if the communication overhead is high. In electronic trading the quality of service could be bad even for an experienced computer user. In an article in the magazine Individual Investor in Apr, 98 issue, a user mentioned difficulty installing an Internet Browser, getting access to his account and placing the trade. Basically he did not know the problems with all the servers and the communication was poor. It took him several hours to learn that his order was not executed but was good news since next time, he bought it at a lower price. In my experience to trade on phone, I found the steps as follows. The person dials the phone number of the broker (phone may be busy), the person picking up the phone has to page the broker assigned to the account, the caller is identified (not much problem of security on phone). The customer specifies the stock of interest, asks questions like bid and ask price (detailed questions as volume, high/low of the day may not be always possible without being put on hold again). The customer places the order for the trade and the broker calls back with confirmation of execution. The steps take about 3 to 4 minutes and the return call from broker may take up to 15 minutes or more.

Another way to trade stocks is to use the phone where the whole transaction is completed by punching account number, password, transaction type and other details. This process just send the transaction to the broker who then enters it in the system. The transaction may not execute for 15 to twenty minutes since the whole process is repeated on phone and computer.

The electronic trading over the Internet is the emerging technology and over 20different than computer trading by institutional investors where computer sell/buy programs are triggered based on some criterion. The electronic trading involves distributed processing and communication among several servers. The network latency plays a major role in the response times. First one must open the browser such as Netscape or the Internet explorer. This takes about 15 seconds on a PC. Next the user accesses the broker's home page that takes another 2 seconds. It takes another 30 seconds to go the trading page where one can be secure and login. After logging, the customer may want to get real time quote from a New York Stock Exchange server (such as quote.com service). It has taken 5 to 10 seconds depending on the time of day. After the transaction is entered, the system presents the order back again to the user for confirmation and that takes about 10 seconds. The user finishes the transaction with a confirmation entry. The user can also check the status of the order in 5 seconds. Other requests for holdings in account, price charts, research reports are simple database queries and take 5-10 seconds.

The whole process of ordering a transaction for stock or option trading takes several round trips among the servers on the person computer, the broker's computer and the NYSE/NASDAQ computers. The communication time over the WAN, LAN, and the security mechanisms affect the response time for the customer. Some time is for the display of pages on the screen. If the communication time can be reduced, the transaction can take place in about two minutes and that may be acceptable quality of service. The stock price can fluctuate in this time so for a day trader, or institutional investor, this is too much time. The electronic broker can not succeed unless the communication behavior can be improved via higher bandwidth. Many individual users have only one line coming in home for voice. The multi-media presentation requires better connections and modems. The user may want to watch a financial channel (CNNFN or CNBC) and at the same time talk on the phone and do electronic trading. The new box recently announced by Sprint Corporation that has one line going in the house but many phone services available inside will be a step forward in meeting such a requirement. It is the same thing as one electric connection to the house but inside, we can simultaneously use many appliances and can get charged based on the units used as measured by a meter. The present configuration of having one line for electricity, another for TV, another for Internet going in the house is not very appealing and must be changed. The communications will make or break the success of not only trading but other electronic commerce applications and we are fortunate to see many innovations coming in this area.

Acknowledgment

Melli Annamalai contributed in digital library research. Shunge Li and Sheng-Yih Wang contributed in video conferencing research. Anjali Bhargava contributed in electronic trading applications.

References

[1] Bandula W. Abeysundara and Ahmed E. Kamal. High-speed local area networks and their performance: A survey. *ACM Computing Surveys*, 23(2):221–264, 1991.

[2] Bharat Bhargava. Adaptable video conferencing. Technical report, Purdue University, Department of Computer Sciences, 1997.

[3] Bharat Bhargava, Shunge Li, Shalab Goel, Chunying Xie, and Changsheng Xu. Performance studies for an adaptive video-conferencing system. In *Proceedings of the International Conference on Multimedia Information Systems (MULTIMEDIA 96), New Delhi, India*, pages 106–116. IETE, McGRAW HILL, February 1996.

[4] Bharat Bhargava, Shunge Li, and Jin Huai. Building High Performance Communication Services for Digital Libraries. In *The International Forum on Advances in Digital Library*, Tysons Corner, Virginia, May 1995.

[5] Bharat Bhargava and John Riedl. A model for adaptable systems for transaction processing. *IEEE Transactions on Knowledge and Data Engineering*, 1(4), December 1989.

[6] Bharat Bhargava and John Riedl. The Raid Distributed Database System. *IEEE Transaction on Software Engineering*, 15(6), June 1989.

[7] Bharat Bhargava and Yongguang Zhang. A study of distributed transaction processing in wide area networks. In *Proceedings of the COMAD 95 , Bombay, India*, 1995.

[8] Bharat Bhargava, Yongguang Zhang, and Enrique Mafla. Evolution of a communication system for distributed transaction processing in Raid. *Computing Systems, The Journal of the USENIX Association*, 4(3):277–313, 1991.

[9] Bharat Bhargava, Yongguang Zhang, and Enrique Mafla. Evolution of communication system for distributed transaction processing in Raid. *Computing Systems*, 4(3):277–313, 1991.

[10] Robert L. Carter and Mark E. Crovella. Measuring Bottleneck Link Speed in Packet-Switched Networks. Technical Report BU-CS-96-006, Computer Science Department, Boston University, March 1996.

[11] R. Entlich, L. Garson, M. Lesk, L. Normore, J. Olsen, and S. Weibel. Making a digital library: The chemistry online retrieval experiment. *Communications of the ACM*, 38(4):54, April 1995.

[12] Ron Frederick. Experiences with Real-Time Software Video Compression. In *Proceedings of the Packet Video Workshop*, Portland, Oregon, September 1994.

[13] Richard Golding and Darrel D. E. Long. Accessing replicated data in an internetwork. *International Journal of Computer Simulation*, 1(4):347–372, 1991.

[14] Abdelsalam Helal, Yongguang Zhang, and Bharat Bhargava. Surveillance for controlled performance degradation during failure. In *Proc of the 25th Hawaii Intl Conf on System Sciences*, pages 202–210, January 1992.

[15] Shunge Li. *Quality of Service Control for Distributed Multimedia Systems.* PhD thesis, Department of Computer Science, Purdue University, December 1997.

[16] Shunge Li and Bharat Bhargava. Active Gateway: A Facility for Video Conferencing Traffic Control. In *Proceedings of COMPSAC'97, Washington, D.C.*, pages 308–311. IEEE, August 1997.

[17] D. Merrill, N. Parker, F. Gey, and C. Stuber. The university of california CD-ROM information system. *Communications of the ACM*, 38(4):51, April 1995.

[18] Calton Pu, Frederick Korz, and Robert C. Lehman. An experiment on measuring application performance over the Internet. In *Proceedings of the 1991 ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems, San Diego, CA*, May 1991.

[19] Changgui Shi and Bharat Bhargava. A light-weight MPEG video encryption algorithm. In *Proceedings of the International Conference on Multimedia Information Systems (MULTIMEDIA 97), New Delhi, India.* IETE, January 1998.

[20] Alfred Z. Spector. Communication support in operating systems for distributed transactions. *Networking in Open Systems*, pages 313–324, August 1986.

[21] Liba Svobodova. Communication support for distributed processing: Design and implementation issues. *Networking in Open Systems*, pages 176–192, August 1986.

[22] Ronald J. Vetter. Videoconferencing on the Internet. *IEEE Computer*, 28(1):77–79, January 1995.

[23] Andreas Vogel, Brigitte Kerherve, Gregor von Bochmann, and Jan Gecsei. Distributed multimedia and QOS: A survey. *IEEE Multimedia*, 2(2), 1995.

[24] Larry D. Wittie. Computer networks and distributed systems. *IEEE Computer*, 24(9):67–76, September 1991.

[25] Yongguang Zhang. *Communication Experiments for Distributed Transaction Processing – From LAN to WAN.* PhD thesis, Department of Computer Science, Purdue University, 1994.

[26] Yongguang Zhang and Bharat Bhargava. Wance: A wide area network communication emulation system. In *Proceedings of IEEE Workshop on Advances in Parallel and Distributed Systems (PADS)*, pages 40–45, Princeton, NJ, Oct. 1993. IEEE.