M.I.T. Laboratory for Computer Science

Request for Comments No. 242
March 10, 1983

Inter-Organizational Networking: Stringing Wires Across Administrative Boundaries

by Deborah L. Estrin

NOTE: This paper will be presented at the Eleventh Annual Telecommunications Policy Research Conference in Annapolis, Maryland, April 1983.

1. Introduction

Inter-corporate networks have been implemented in many industries; including transportation, travel/airline, insurance, banking, and grocery. In addition, as local area networks are installed at more sites, questions are commonly raised about the possibility of interconnecting them. Technically, it is becoming straightforward to accomplish these interconnections. However, when the links cross administrative boundaries, either within an organization or between quite distinct organizations, administrations often pose requirements for control of the use of the links; or they belatedly discover the demand for such controls as a consequence of unfortunate experiences. Organizational-policy problems that arise include: controlling where data is stored and processed; cost accounting; addressing; control of transit; and security, i.e., authorization of use and authentication of users. Public-policy issues include the impact on, and reflection of, changing inter-corporate relationships; implications for non-participatory firms; relative benefits for large and small firms; and the standards setting process. The importance of such policy issues will increase as computer systems that process external communications are interconnected and integrated with the systems that each organization uses internally. Both technical and economic factors determine the extent to which available communication-protocol mechanisms can be relied upon to enforce policy requirements, although existing, state-of-the-art, inter-networking technology typically lacks such mechanisms altogether.

WORKING PAPER – Please do not reproduce without the author's permission and do not cite in other publications.

In the first half of our paper we introduce features of inter-organizational network interconnections that should evoke interest or concern on the part of organizational and/or public policy makers. We describe three categories of network application in terms of the range of capabilities made available to external users and the degree of control over external usage available to the organization. Finally, we briefly describe particular examples of inter-organizational network interconnection.

In the second half of our paper we expand on the organizational-policy issues identified in the previous examples. We discuss the relationship of organizational-policy concerns to the type of administrative boundary(ies) crossed and the role of technical mechanisms in addressing these concerns. We conclude by speculating on a number of public-policy questions that inter-corporate network interconnection will pose.

2. Assumptions and characteristics of inter-organizational networks

Our discussion pertains to organizations that make use of their own computing facilities for processing of external transactions and communications. Organizations that also employ internal networks to support electronic communications among a range of internal computing devices and resources (e.g. shared databases, computational resources, text processing, file and record management, inventory control, transaction processing, etc.) have added economic and technical incentives to integrate external-transaction and internal-processing facilities.

On the one hand, we assume that the interconnecting organizations desire to facilitate some operation that crosses an administrative boundary so as to achieve increased efficiency, cost savings, etc. We also assume that the organizations do not want to permit unrestricted access or integration with all internal facilities; i.e., they want to discriminate between internal and external users. On the other, we assume that the organizations place a significant but definitely finite value on security; in a word, we assume a commercial as opposed to military environment.

The interconnection of the organizations' facilities (i.e., machines or networks) need not manifest itself in the installation of a physical wire or switch, but only in an agreed-upon protocol for transferring data to a designated destination and interpreting the data at the remote location. We

refer to the mechanism whereby data or information is passed between organizations as a *gateway*[1].

The fundamental difference between computer-communication networks that operate across administrative boundaries and more traditional inter-organizational communication modes such as voice telephone, telex, and postal mail, is that a user in one organization can cause some event to occur automatically within the domain of another organization, without any human intervention or auditing. We will view the degree to which external capabilities are integrated, in an automated fashion, with internal capabilities as a measure of the affinity between interconnecting organizations.

3. Categories of network application

In this section we identify three categories of network application, *electronic mail, file transfer, and remote login.* These categories significantly differ from one another in the range of capabilities made available to external users and the degree of control over external-usage available to each organization.

The control objective in all three categories is to check that a user requesting access to a remote host or other facility is indeed authorized to do whatever it is that is being requested[2]. This might require, for example, identifying and authenticating who the user is, if the user is authorized to use the facilities, and if the particular use intended is appropriate.

3.1. Electronic mail

The forwarding of electronic mail among computers or networks can be viewed as a very restricted form of transferring files. In most mail transfer schemes, files containing the message texts can only be *sent to*, not *removed from*, the destination computer. In addition, the user who is originating the message does not specify the name of a file on the destination computer into which the text is to be written, rather the mail is automatically sent to a specific mailbox file corresponding

---

[1]The term *gateway* more traditionally refers to a device that interconnects networks at both the physical and higher communication-protocol levels.

[2]The term *host* refers to a shared computing resource with which users communicate via a terminal (e.g., VDT) over a direct wired connection or over a remote telephone, or some other media, connection. A host or facility is remote to a user if it does not belong to that user's organization.

to the intended recipient of the message. Finally, it is not even possible to overwrite an existing file on a remote system, rather mail is *appended* to mailbox files.

As a result of their relative restrictiveness, many electronic mail systems do not require that users in one organization identify themselves to the computer system of a remote organization in order to send messages to users on that remote system. Because the capabilities of the sending users are so restricted, unauthorized usage is less threatening.

Nevertheless, there are a number of interesting problems that mail-forwarding connections raise and these problems are aggravated by the absence of an identification or login procedure. In particular, the restrictiveness of a mail-channel connection can be significantly weakened by the manner in which persons *within* an organization use their internal, electronic-mail system. For example, internal users can implement simple programs that "execute" the contents of received messages, thereby turning a passive, non-interactive mail channel into an active one. Such a mail channel might be used as a general file-transfer channel if a message from an user contains a machine-readable command to send a file located within the domain of the message recipient to the message originator. This facility may be quite appropriate and useful within the organization but unacceptable when accessible by external users. Similarly, an authorized network-user may take on the role of a mail-forwarder for otherwise unauthorized network-users. On most computer systems that support electronic mail any authorized user can automatically place messages from other users within his/her own message envelope, thereby causing the message to appear to be from that authorized user [2]. Although the authorized user's organization might not want to incur the cost of transporting such messages, there are no mechanisms to detect such traffic without reading the contents of the message[3]. Automatic execution and forwarding of external mail imply that an incoming message may not be subject to human judgment before the receiving system "acts on" it. Therefore, these capabilities intensify the organizations' need to *authenticate* the identity of the message originator.

A second problem regards mailing lists, which broadcast messages to a roster of destinations all over the network(s). Although mailing lists, in and of themselves, do not constitute an abusive use

---

[3] In most cases, examining the contents of messages is neither an ethically nor technically attractive solution.

of mail facilities, they nevertheless can use up significant network resources and generate significant off-network charges if such mail is forwarded onto tariffed networks. Unless mailing-list, message files are specially tagged, there is no way to detect or differentiate this traffic.

Even in the absence of the more complex uses described above, interconnected organizations whose gateway only accommodates electronic mail will require some subset of the following policy controls. The particular requirements will depend upon the nature of the organizations interconnected and of the boundaries spanned:

- ~ Protection of valuable resources such as traffic handling and mail storage; including the control of transit.
- ~ "Appropriate" usage of mail facilities by authorized users, as determined by the operating organization.
- ~ Protection from maliciousness, in the form of large amounts of traffic or junk mail generated onto the network, i.e., ability to reject incoming traffic.
- ~ Accounting mechanisms that charge the originator, recipient, or forwarder, as appropriate.
- ~ Controlling or accounting for the flow of outgoing mail onto tariffed networks.
- ~ Ability to indicate class-of-service and routes to be avoided.

3.2. File transfer

File transfer capability allows users to send files to and copy files from remote systems. Each system is responsible for restricting access to those files that its internal users do not want copied or overwritten. In contrast to typical mail gateways, most file transfer services do require that users from one organization identify themselves to the computer system of a remote organization before requesting files from or sending files to that remote system[4]. If such identification is required and the remote-system's security mechanisms are adequate, then no unauthorized traffic can be generated onto the network or off of the network. But, if the remote-system's security is not sufficient, the potential loss from unauthorized file transfer is much greater than with a mail-only gateway because file transfer provides users with a means of *extracting* information (in the form of files) from a remote system; similarly, it provides them with a means of overwriting information.

---

[4]The most common technique of identifying oneself to a system, i.e., *logging in*, is by providing a valid user account name and corresponding secret password.

If it is necessary to restrict authorized, external users from gaining access to proprietary, private, or otherwise valued, internal information, secure login procedures may not be adequate in and of themselves. Since many systems in operation today do not adequately protect the files of authorized users from one another, and robust system security can not be retrofitted, additional security-mechanisms may be required in the gateways between some organizations.

3.3. Remote login

A gateway that supports remote login communications is both more and less restrictive than a mail-only gateway. On the one hand it is more restrictive because it is only equivalent to a dial-up telephone line and unlike a mail-gateway, it provides no capability to generate traffic onto the network unless users can identify themselves to the satisfaction of the security system of the remote host. On the other hand, a remote login capability exposes many more of the remote-system resources to external users than does a mail capability.

As in the case of file transfer, even if the remote-system's security mechanisms are adequate to keep *unauthorized* users from gaining access, the existence of a larger and more diverse community of *authorized* users intensifies the need for file protection and data-base security mechanisms. In contrast to file transfer, remote login capability may avail all[5] system resources to the remote user, with the exception of those resources that are explicitly protected by the system security.

The policy requirements for remote login and file transfer via a network or internetwork are to a large extent the same as the security requirements of the host system. Nevertheless, connection to a larger and more heterogeneous community of external users is problematic because an organization may never have thought through its *internal* policies adequately.

4. Three arenas of inter-organizational networking activity

In this section we briefly describe particular network interconnections in terms of the applications supported and the problems posed. We selected the examples from three arenas: research-center/university communications; industry-wide communications; and customer-supplier communications.

---

[5]For example, computational, operating system, and application packages are all accessible, in addition to the file system.

4.1. Research-center/university communications

Of the three arenas, research centers, universities, and government institutions currently engage in the most sophisticated inter-organizational networking activity in terms of:

- ~ The degree of integration of those systems that handle external communications with systems that the organizations use internally,
- ~ The type of communication application supported (see section 3), and
- ~ The intensity of usage.

As a result, of the three arenas discussed, this arena presents organizations with a unique set of policy issues.

- ~ The Arpanet: [4] "The Arpanet is an operational, resource sharing, host-to-host network linking a wide variety of computers at DoD facilities and non-DoD research centers..."[6]. It began as an experimental project in 1969 to advance the state-of-the-art in computer networking. Participating organizations include universities, government agencies, and commercial contractors (e.g., Xerox, Honeywell, and Digital; but neither IBM nor Bell Labs). In addition to providing an infrastructure for communication research, the network supports technology transfer and diffusion among the research communities. The cost of Arpanet usage typically is covered by DoD research contracts. The Arpanet is used for mail, file transfer, and remote login.
- ~ CSNET:[7] Recently the National Science Foundation organized CSNET to interconnect computer science research groups in universities, industry, and government to provide a means for collaborative research. CSNET is a logical network which operates over a combination of telephone lines, Telenet[8], and Arpanet facilities. Unlike the Arpanet, the cost of access and usage is not covered by the coordinating organization, NSF, but by each participating research center. As might be expected, research centers have been slow to join CSNET because of a poor sense of what it is they are purchasing. As a result, facilities are currently limited to mail-transfer and do not accommodate file transfer or remote login. Mail is automatically forwarded between CSNET and Arpanet users.
- ~ BITNET:[9] Bitnet is a network composed of university computer-center computers for messaging, file transfer, and remote login; most of the computers are IBM models. Some computer companies, such as IBM, are also participants. Mail and other files are

---

[6] Network Information Center, *Arpanet Directory*, Technical Report NIC 49000, Defense Communications Agency,, March 1982.

[7] L. Landweber, M. Solomon, "Use of Multiple Networks in CSNET", In *Compcon*. IEEE Computer Society, Spring 1982.

[8] Telenet is a commercial packet-switched communication service.

[9] I. Fuchs, "BITNET -- Because it's time", *Perspectives in Computing*, volume 2, number 2, April 1982. This journal is published by IBM.

passed from one machine to the next-closest, participating machine, via public telephone lines, until the the files arrive at their intended destinations (as specified by the sending machine).

- USENET:[10] Bell Laboratories developed a communication protocol to allow file and mail transfer among any two computers that run a particular operating system, Unix™. Based on this protocol, an ad hoc network, USENET, has developed into a nation-wide network of Unix machines via which users can send text to one another, to electronic bulletin boards, and to users on the ARPANET, CSNET, and BITNET (via those machines that are connected to both Usenet and one of the other three networks). The messages and files are sent over telephone lines on a rather unstructured basis from one Unix machine to the next closest machine, which in turn automatically forwards the text on to other nearby machines, and eventually the text (which contains a destination address) reaches its intended destination. Unlike the previous three networks there is no central body that manages the operation of this network. Participating organizations now include a mix of university and computer-company computers, including a large number of Bell laboratory machines and IBM. The cost of access is limited to the telephone call to the nearest USENET neighbor.

- University-Corporate links: University-corporate links supporting joint or sponsored research comprise another variety of network interconnection that appears in this arena; examples include: MIT-Symbolics-LMI; IBM-MIT; UMaine-IBM (Bitnet), etc. In most cases the motivation is technology transfer, and improved industry-university relations. Typically these connections provide the participating firm with internal-user status on the university's system or with access to the university's internal network, and thereby to a number of university resources.

One policy concern that arises in this arena is the use of publicly funded and supported network facilities by private users, for personal or company profit. For example, a company that achieves an indirect connection to more than one site on the Arpanet (e.g. via a legitimate Arpanet site such as M.I.T. or via forwarding mechanisms from Usenet to Arpanet), could use it as a transit network across the country from one of its connected sites to the other. Indirectly connected organizations may also have access to Arpanet information and computational resources that are inadequately secured against external usage because attempted access was only foreseen via the legitimate Arpanet community.

If an organization (A) that is connected to the Arpanet also connects to a non-Arpanet-connected organization (B), how accountable is organization A for unauthorized traffic sent by B via A's Arpanet connection? Similarly, connections between universities and commercial firms raise questions as to the liability of the university for unauthorized access to proprietary information within the private firm's facilities. One response might be to require that an organization gain the

---

[10] M. Horton, *How to Read the Network News*, Technical Report, Bell Telephone Laboratories, 1981.

approval of every other organization to which it is connected before connecting to some new organization; for example, M.I.T. would have to acquire DoD's approval for every desired network interconnection with a non-Arpanet organization. The inability to establish connections via pairwise agreements (i.e., between the two interconnecting organizations) that is implied by this scenario is both undesirable and infeasible. Technical mechanisms are needed to isolate the networks of various organizations in a way that does not impair their intended functions.

The problems of excessive mailing-list traffic, private mail-forwarders, etc. (section 3.1) are experienced *daily* in the combined Arpanet-CSNET-USENET-BITNET environment. Some users in commercial firms accuse university users of engaging in inappropriate discussions over the shared USENET resource; on the other hand, some university users accuse commercial users of misusing government resources by forwarding mail from one USENET site to another USENET site via Arpanet facilities.

4.2. Industry-wide communications

Current inter-corporate interconnection activity for the most part involves the interconnection of facilities that are dedicated to the single purpose of processing external communications. In particular, the interconnected computers are for the most part isolated from other internal computing and communication facilities that the individual organizations employ. As a result, existing inter-corporate networks do not provide a pathway via which information not intended for external use might be accessed. But if, as is predicted [3], organizations increasingly employ distributed data processing technology to support numerous administrative and business functions, they will increasingly integrate their internal applications of computing and communication technologies with one another and similarly with their processing of external communications. Given such integration, the resulting inter-corporate interconnections may begin to encounter the complex problems of the Arpanet community discussed above.

~ Electronic Data Interchange:[11] An American National Standards Institute committee[12] is currently defining standards for Electronic Data Interchange. This effort is not concerned with the construction of a network in any physical sense but rather the definition of a common standard and protocol that will allow transfer and automatic

---

[11] T. Jones, "Paving the way for universal document interchange", *Data Communications*:123-131, July, 1982.

[12] ANSI X.12 committee.

interpretation of highly-structured textual information. The transportation industry developed the document format on which X.12 and a number of other industries' standards are based. The elimination of costly rekeying of data, increased speed, and reduced errors due to rekeying are among the motivations cited for these standardization efforts. The ANSI committee estimates a potential savings of $10 billion/year with 50% of inter-corporate business transactions implemented electronically. Similarly, the grocery industry estimated a $300 million/year savings with 50% of their inter-corporate transactions implemented electronically; $85 million of this savings is attributed to rekeying cost savings alone. The standards will be used to transmit purchase orders, invoices, and other business and administrative information among firms of a common industry. These messages will be sent to and received by a company's computer and, depending upon the extent of that company's internal automation, the contents of the message may be automatically interpreted and acted upon by the company's inventory system.

The standards will not provide for interactive communications between a user in one company and the facilities of another company. Moreover, the very structured nature of the text will deter the facilities from being used for anything but the intended application. Unlike the mail systems discussed in section 3.1, the standards do not provide for forwarding of messages from one company to another via the facilities of a third company; messages are sent directly from the originator to the source.

~ Insurance:[13] The insurance industry, under the direction of the Insurance Institute for Research, is planning the construction of a nationwide data communications network that will link computers at thousands of agencies, insurance companies, and service organizations. The industry hopes to increase the efficiency and speed of issuing insurance policies by using direct computer communications rather than postal mail delivery and by eliminating the need for rekeying of document information. The insurance industry project will entail designing standard message formats in addition to the means of transmitting these messages. Because of the large number of small independent agents (approximately 60,000) the technology must be very simple and inexpensive (thus, for example, the feature of variable-length fields in the forms specified by the ANSI was rejected as being too costly). Such consideration of smaller users is consistent with the fact that most of the pressure for establishing industry-wide standards has come from the independent agents who must deal with multiple insurance companies and cannot afford the expense of multiple interfaces. The network will accommodate transfer of files and messages regarding policy requests, status information and other administrative business forms.

~ Airlines:[14] The international airline-reservation network, developed and operated by SITA[15], is one of the oldest inter-organizational data networks[16], interconnecting the reservation systems of the majority of airlines, worldwide. It supports real-time exchange of messages in addition to batch transfer of commercial, technical, and

---

[13] *RFP for Development of an Intelligent Data Communications Facility*, Insurance Institute for Research, Inc., White Plains, New York, March 31, 1982.

[14] "Interconnect net navigates airline reservations", *Data Communications*, September 1982, pp. 99-107.

[15] Societe Internationale de Telecommunications Aeronautiques.

[16] Planning began in 1949 for use of radio-telegraph circuits.

administrative information. The airline industry is traditionally an advanced user of computer technology. As we might therefore expect, many airlines have integrated the processing of external communications with their internal reservation-systems. Although we have no data from which to conclude the presence or absence of policy problems concerning these inter-organizational connections, over the past year some controversy has arisen over the operation of the reservation systems themselves. A number of smaller airlines that make use of the reservation systems operated by larger airlines accused the larger airlines of giving themselves preferential treatment in the processing of reservations[17]. This example suggests the sort of difficulties that can arise when computing facilities are shared across organizational boundaries between competing organizations.

The relative simplicity of current technology constrains the interconnections described above so that they resemble postal or telephone connections more so than computer-network connections. The participating firms will find these interconnections of greater organizational-policy concern once a wider range of functions are automatically invoked and integrated across organizational boundaries. With increasing automation within the organizations, firms will find it ever more economical to interconnect their facilities, both internally and externally. Then a more complex set of policy issues, such as those that the Arpanet community experiences, are likely to arise. The absence of a governing body (such as DoD in the case of the Arpanet), and the competitive relationships between the participating organizations, may make the resolution of these problems, especially problems of liability and shared costs, quite difficult. All of these networks are of most interest from a public policy perspective because of evidence they provide of increased inter-corporate activity.

4.3. Customer-supplier communications

Another variety of inter-corporate communications is the connection of facilities belonging to customers with suppliers of industrial products and equipment. In a number of industries suppliers are providing their customers with access to specialized order/entry systems. In some cases, the customer's inventory system is directly connected to the order/entry system of the supplier so that orders are automatically submitted when stock levels in a particular item reaches a specified minimum. If a terminal is used to access the supplier's system directly then the function supported is strictly remote login. If on the other hand, the customer communicates with the supplier's computer via its own internal computer system, then file transfer of invoices and transfer of

---

[17] R. Witkin, "Airline Computers Generate Fare Rift", *New York Times*: Business Section, March 1 1983.

administrative messages from supplier to customer can also be supported. Customers and suppliers cite the motivations for such facilities as: cost-saving, efficiency, level of service, and, in the case of suppliers, solidification of the relationship with customers[18].

- ~ American Hospital Supply Corporation:[19] AHSC provides its customers with access to a system that supports inventory, purchase order, and automatic invoicing of customers. In addition, AHSC sells a turn-key inventory system that is integrated with this order/entry system. Using the order/entry system, 58% of the business input is done by the customer. As a result, no AHSC employee handles any information related to a routine customer order until the order arrives at the warehouse for shipping.

- ~ Raytheon[20] A similar project interconnects Raytheon Co. with its electronic-component suppliers. At Raytheon, 50% of the selling price of many small electronics parts is the price of purchasing materiele. Depending upon the supplier's facilities, Raytheon accesses the supplier's order/entry system via a terminal and telephone lines or, preferably, a direct connection is established between Raytheon's inventory system and the supplier's order/entry system; the latter eliminates additional rekeying effort on the part of the Raytheon customer. In addition, some of Raytheon's suppliers are in turn connected to the facilities of their distributors and will permit Raytheon customers to inspect distributors' inventories via the supplier's facilities. As their use of computer aided design and manufacturing equipment increases, Raytheon foresees adapting this technology to provide their customers with a means of inputting product requirements into the design process.

In the above examples, despite the provision of a remote login connection to the machine of another organization the specific function that is meant to be accessed externally is largely isolated from any other internal resources of the supplying organization. As a result, security issues can be minimized. Similarly since there is no network or shared communication resource, there is no concern about transit. As the use of external systems and networks increases and is integrated with internal functions, both of these aspects of the current arrangements are likely to change in ways that intensify the organizational-policy concerns. For example, if a supplier's internal system *automatically* interprets and acts on customers' orders, without subjecting them to human inspection, then the need to verify and authenticate the source of each order will be intensified. The issue of greatest public-policy concern is the tendency of these arrangements to increase the rigidity of customer-supplier relationships.

---

[18] "Electronic order entry called wave of the future", *Wholesalers World*:1,14, April 1, 1982.

[19] Edward Doerhoefer, AHSC; *Inter-corporate networking: supplier-customer networks*, M.I.T. Laboratory for Computer Science seminar, Cambridge, MA, November 17, 1982.

[20] Arthur Casavant, Raytheon Company; *Inter-corporate networking: supplier-customer networks*, M.I.T. Laboratory for Computer Science seminar, Cambridge, MA, November 17, 1982.

5. Organizational-policy issues

The nature of the organizations being interconnected and their relationship to one another determine the nature of the policy requirements for internetwork access controls. Nevertheless we can identify characteristics and concerns that commonly arise. In this section we focus on the organizational-policy issues that arise across *budget* and *proprietary* boundaries.

5.1. Budget boundaries

When networks belonging to different organizations are connected, problems arise with regard to budget boundaries. The transmission and handling of messages must be billed to the users, or groups of users, and payment must be made to any intermediate organizations that transmitted the communications. Some of the accounting functions required among interconnected data networks are equivalent to those used in the existing voice, record, and postal communication networks, while others are unique to data networks. In particular, interconnecting organizations will desire mechanisms to address the following accounting requirements:

- ~ Allocating payment to all the networks across which a message travels to its destination.
- ~ Charging the recipient for incoming communications when the application is such that it is not appropriate for the originator to incur all or any charges.
- ~ Allowing the recipient to reject unwanted communications without/before incurring the transmission and storage costs that are associated with accepting and handling messages on the recipient's network.
- ~ Allowing inter-networked organizations to maintain dissimilar internal accounting practices; for example, usage-sensitive and non-usage-sensitive.
- ~ Charging for non-communication, network services such as storage, printing, data, computation, etc.

Mechanisms will be required to allow internal networks to charge for transmission and handling of external traffic. This traffic may be destined for a user on the internal network or for a user on a third network, to which the internal network is connected[21]. It might not be practical for each network to keep track of the originator of each message that is transmitted over the network and then bill that user or host directly, particularly on packet-switched networks that send data as

---

[21] The latter instance was referred to earlier as third-party routing.

individual *packets*[22] of data as opposed to groups of packets corresponding to a single message or file [5][23].

If we can assume that the organizations' internal networks are equipped with mechanisms for allocating charges to internal hosts and users, then each network could treat each external gateway as another host and thereby charge the neighboring network for any traffic that arrives via that gateway[24]. The costs of the communications would eventually be passed back to the originator's internal network and then to the originator. If, as is often the case, the internal network is operated by a central authority that has had no previous need or mechanism for internal accounting and charging of internal network resources to individual hosts on the network, no internal accounting mechanisms can be brought to bear and mechanisms must be implemented in the gateways for recording the communication flows into the network. In fact, charging internal users or hosts for internal-network usage may be *inappropriate* if an organization desires to encourage local communications. If a considerable amount of traffic flows among a group of networks (electronic mail carriers, industry-wide networks, etc.), accounting might be handled on a statistical basis, where charges are based on a sampling of the traffic flowing between networks, thereby reducing the amount of monitoring needed in the gateway.

Dissimilarities among the internal accounting practices of the the interconnected organizations

---

[22]The term *packet* refers to a group of binary digits, including data and control information, which is switched as a composite whole.

[23]The former is referred to as a *datagram* network and the latter as a *virtual circuit* network. Telephone-billing methods are not necessarily applicable to packet-switched networks because of the large number of packets associated with any single data-communications session, e.g., a file or host-to-terminal communication. Accounting is easier if packets are relayed between networks a message-at-a-time, as opposed to a packet-at-a-time, because the message originator is easier to identify. On the other hand, message-at-a-time relaying can degrade the performance of the system. In addition, the international voice and record carriers have significant experience in this area but because the number of carriers involved is generally restricted to one per country, i.e., with only a few exceptions telecommunications is a monopoly-provided service, the solutions might not be applicable to data networks where there will be a great number of independent private and public networks in many countries; excepting those countries that forbid private networking.

[24]An additional complication arises when the recipient, not the originator, should be charged for the communications, e.g., some mailing lists. In some cases the originator can pay the communications charges and bill the recipient for services rendered. When this is not feasible, a mechanism similar to collect telephone calls, or pay-on-delivery postal service is desired. In addition to imposing complex accounting requirements on the gateways, such mechanisms do not prevent the originator and transit networks from incurring the cost of transmitting the message should the recipient choose not to accept the charges.

might dictate a need for special control mechanisms at network boundaries. If one of the inter-networked organizations does not charge its users on a usage-sensitive basis (e.g., Arpanet) then access to external services that are charged to the organization on a usage-sensitive basis (e.g., Telenet, TWX, Telex) need to be restricted in some manner. That is, either the external service-charge must be passed back to the user, requiring an accounting system at the junction of the two networks (e.g., long-distance phone calls), or there must be an access-checking mechanism at the junction (e.g., a list of authorized users which is matched with the originator field on the electronic envelope, or a password procedure)[25].

Although the term *inter*-organizational is repeatedly used in this paper, controls and restrictions are in many cases desired between networks belonging to different departments or divisions of the same organization. The primary difference in the case of inter-departmental networking is that there exists a central authority to which both departments report, whereas no such central authority exists in the case of most inter-organizational networking efforts.

5.2. Proprietary boundaries

When organizations interconnect their private networks, concerns arise regarding protection of proprietary information and resources from being viewed, copied, or altered. A significant difference between interconnection across proprietary and military-security boundaries is that unlike the military, the owners of private networks typically are not willing to expend large amounts of money on security mechanisms (either for networks or single machines).

Protection of valuable resources entails restriction of access to transmission, computational, and information database resources to authorized users; it might also entail deterrence of inappropriate usage by authorized users.

Transmission resources can be abused through third-party routing, inappropriate usage, and uncontrolled incoming traffic. Third party routing [1], refers to the situation in which a network is used as a transit path between two other networks. The third-party network therefore bears the cost

---

[25] One example of a budget boundary that is complicated by the dissimilarities in the internal accounting practices of the two organizations is the Telenet gateway on the Arpanet. Currently the Arpanet allows any authorized user of an authorized Arpanet host to generate mail on Arpanet facilities, but only a restricted subset of users, those belonging to a restricted set of hosts, are permitted to send mail via the Telenet gateway.

of transmission for the other networks. In those cases in which the third party is not affiliated with either the sender or receiver, this arrangement is not appropriate unless the third party has agreed to act as a transit network; even so, issues of accounting, dependence, liability, and competition with common carriers are raised. Transmission capacity and other facilities required to carry the extra traffic are valuable resources which many organizations will want to protect via control of usage.

In an attempt to protect access to valuable transmission resources, incoming traffic from external sites might be prevented from traveling over private, enhanced[26] facilities. Unfortunately, such action can impair the ability of in-house network users to engage in valuable communications with external sites.[27] One solution might be to accept only traffic that is destined for one of the networks to which the gateway is directly connected, i.e., prevent all transit. Unfortunately, such action might defeat one of the motivations for interconnection, namely, economies of scale that result from the formation of an internetwork whose physical transmission links are comprised of the private network facilities of the participating organizations. In other words, one of the features of interconnection might be the ability to accommodate transit; the issue then remains how to provide each network owner with the capability to set limits upon, or in some other way account for and control, the transit usage (see section 5.1).

Protection of computational resources and information contained in data bases are important issues both within a single computer system and within a single network. The need for protection of these resources is aggravated when the number of users that might have access to the resources is enlarged, particularly when the additional users are not part of a single community, i.e., not

---

[26] *Enhanced* refers to lower cost, higher performance, higher security, etc.

[27] Autovon, the voice network used by the defense community, can be used by persons within the physical community (i.e., certain private exchanges) to contact persons outside, via connection to the public switched network. But incoming calls that are not from other Autovon users, are not routed over Autovon facilities and use only the public network. Therefore, although the necessary connectivity is achieved with the public, access to the special network facilities are restricted to outgoing calls of authorized users or to calls within the Autovon community. Autovon thereby solves the policy problem of restricting access to enhanced facilities to authorized users while still allowing authorized users to connect to unauthorized users via enhanced facilities. This solution is not immediately extendable to private data networks because in order to restrict external, public, messages from being transmitted via private-network facilities, each host that supports mailboxes must have a direct connection to the public network so that incoming mail may come to the host directly, and not via private facilities which connect the particular host to a shared gateway. The difference between the Autovon voice and the data network case is that all Autovon-user complexes (analogous to hosts on the private data-network) are in fact directly connected to the public switched network in addition to Autovon facilities, whereas one of the explicit roles of private data networks is to provide shared gateways to such public facilities.

controlled by a common authority. In general, protection of system resources is best left to the system itself. In particular, a *network* should not be expected to increase the security of a *system*, but at the same time it should not lower that security. One approach to controlling access to valuable resources that are not proprietary is to charge for them (see section 5.1). Otherwise, the solution to this problem lies in the domain of a single organization's system security.

Malicious behavior can take the form of unauthorized viewing, destruction or damage to valuable resources, or denial of service. The first two examples can be addressed in a similar manner as protection of valuable resources. On the other hand, denial of service through such things as jamming gateways or flooding communications paths requires a rather different approach. In many communities, these threats are too expensive to address given the relative value and sensitivity of the information and resources at stake and the high cost of protection measures. Because of the greater diversity of the user-community, inter-organizational networks must address problems of malicious behavior more seriously than do strictly internal networks. In response to maliciousness, interconnecting organizations may resort to increased auditing and monitoring, and in extreme cases, the ability to shut down the gateway.

While many of the control requirements discussed are also applicable to protection of users' privacy, some of the methods for protecting valuable resources have the unintended consequence of decreasing the level of privacy in a system; e.g., the increased monitoring and accounting of communication flows through gateways. We would emphasize the distinction between restricting access to network facilities to authorized users, and restricting usage of the facilities by authorized users to appropriate applications; the first is a matter of control of carriage, the second, control of content. Whereas the control of carriage might lend itself to a priori control policies and mechanisms such as examination of the source and destination fields on an electronic envelope, a priori control of content impinges on the privacy of users because it entails examination of message contents[28]. A posteriori content-regulation, i.e. policy statements as to appropriate usage of facilities accompanied by strict enforcement should abuse be detected, also entails examination of message contents in order to detect abuses (although some abusive traffic can be detected without monitoring traffic contents via the complaints of subscribers who receive such mail). Different

---

[28]The term *a priori* control refers to prevention of the abuse through advanced detection; *a posteriori* control refers to enforcement of policies once an abuse is detected.

organizations will undoubtedly adopt different postures as to the ethics of examining message contents, but given that any one among a group of interconnected organizations regards it as unethical, examination of contents is unlikely.

5.3. Role of technical mechanisms

The current technological basis for providing policy control between networks is almost completely non-existent[29]. Today, whenever a packet of data arrives at the boundary between organizations it is difficult for any person or program to discover its purpose, since that purpose is buried in layers of protocols, and this packet may be only one of many that are part of a single activity (e.g., a file transfer or host-to-terminal communications stream). Present approaches fall into one of three categories, none of which provides both satisfactory function and satisfactory control:

1. Allow the packet to cross, and depend on the end points (i.e., host systems or users) to initiate only communications that meet policy constraints. This technique fails, for example, if network B finds that it can be used as a transit network between stations on network A and stations on network C. In such a case, network B gets no chance to exert any policy control.

2. Require that all protocols terminate at each gateway between networks. For every application, place a program at the gateway to act as a monitor and relay. Since the protocol is terminated, the underlying purpose of the connection is visible to the monitor, which can more easily enforce policy constraints. This approach is analogous to making a telephone call in which each party can talk only to an intermediate operator, who relays the conversation. While acceptable for some applications, delay and loss of special features cripple other applications.

3. Do not permit the connection in the first place. This approach provides conservative control, but is rather devastating from an application point of view. Given the fear of the alternatives it is probably the most widespread technique used today.

Technical mechanisms can provide varying degrees of support for: access control, authentication

---

[29]It is worth noting that mechanisms of access control, accounting, and billing on networks are at a rather primitive stage of development relative to most other aspects of networking technology. We attribute some of this to the origins of the bulk of the networking experience in this country; namely the Arpanet where, access was for the most part encouraged through minimal access procedures and no billing of users; network development and operation was paid for by sponsoring agencies. Similarly, relatively little research effort in internal area networks has been dedicated to the development of accounting or access-control mechanisms because of the limited geographical and organizational span of most internal area networks.

of user and destination, proper addressing of remote users, and screening forwarded mail. But, all such mechanisms, once tempered by economic constraints, will require that the organizations establish contractual agreements. Only in military, and perhaps in the sensitive financial arenas, will the extremely high cost be paid to eliminate dependence on contracts.

6. Public-policy issues

We find that the public policy issues are for the most part related to the *phenomenon* of increasing inter-corporate communications as opposed to the characteristics of the interconnections themselves. As a result, the public-policy maker is likely to be more interested in the occurrence of such interconnections than in the details of their design or management.

The changing nature of inter-corporate communications raises the following public-policy questions:

- ~ To what extent has inter-corporate communication activity increased? Does inter-corporate networking just represent a shift in technology or does it signal a shift in the character of corporate relationships and modes of industrial operation?
- ~ What are the implications for public telecommunication and postal systems of increased emphasis and reliance on enhanced, private, communication facilities? Will interconnected private networks pose significant competition to common carriers?
- ~ The standardization process is pivitol to implementation of industry-wide networks and technical trade-offs reflect economic and policy decisions; what are the appropriate standardization procedures and practices to achieve efficient and fair results?
- ~ National policies regarding the use of private telecommunications services differ among countries; how are nations to enforce national policies regarding the use of private facilities such as electronic mail and still permit interconnection of networks that are located in different countries?
- ~ What are the implications for non-participatory firms in an industry where networking activity is significant; what are the implications for small business in particular?
- ~ What are the anti-competitive or anti-trust implications of these interconnections?
    - * Will the customer-supplier arrangements discussed increase the rigidity of customer-supplier relationships and mobility?
    - * Is it a matter of public policy whether industry-wide networks practice non-discriminatory access to all members of the industry? Should public policy be concerned with cost barriers to participation?
    - * Should the government monitor or audit the extent of interaction among competing firms in the interest of anti-trust enforcement? If so, how?

~ What will be the employment impact of eliminated data-entry requirements ?

7. Summary

The fundamental difference between computer-communication networks that operate across administrative boundaries and more traditional inter-organizational communication modes is that a user in one organization can cause some event to occur automatically within the domain of another organization, without any human intervention or auditing. We discussed three categories of network application, each of which provides a different range of capabilities to external users and a different degree of control over external-usage to an organization.

Of the three arenas of inter-organizational networking activity that we described, the research-center/university arena employs the most sophisticated technology and applications. We attributed the relative intensity of organizational-policy problems encountered in this arena to the degree of integration of each participating organization's internal facilities with its external-communication facilities. Similarly, we speculated that the absence of such integration in existing industry-wide and supplier-customer communication arenas partially accounts for the rarity with which organizational-policy problems have been encountered.

We generalized on the organizational-policy issues identified in the previous examples and described the concerns that might commonly arise across two types of administrative boundary, budget and proprietary; for example, cost accounting, security, and control of transit. We then commented on the difficulty of employing existing technical mechanisms to implement policy controls in gateways.

We concluded by speculating on a number of public-policy questions that inter-corporate network interconnection will pose; in particular the implications for competition and/or cooperation among firms.

Acknowledgments:

The author gratefully acknowledges the contributions of V. Cerf, P. McClure, J. Saltzer, and M. Sirbu to the ideas presented here. W. Lazarus provided valuable comments on the final draft of this paper.

This research was supported in part by IBM through discretionary funding made available to the M.I.T. Laboratory for Computer Science; and by the Defense Advanced Research Projects Agency of the United States Department of Defense, monitored by the Office of Naval Research under contract number N00014-75-C-0661.

References

[1]
Cerf, V., Kirstein, P.
Issues in Packet-Network Interconnection.
Proceedings of the IEEE 66(11):1386-1408, November, 1978.

[2]
Cohen, D., Postel, J.
Internet Mail Forwarding.
In *COMPCON*. IEEE, Winter, 1982.

[3]
Guiliano, V.
The Mechanization of Office Work.
Scientific American :149-164, September, 1982.

[4]
Kahn, R.
Resource-sharing computer communication network.
Proceedings of the IEEE 60(11):1397-1407, November, 1972.

[5]
Roberts, L.
The Evolution of Packet Switching.
Proceedings of the IEEE 66(11):1307-1313, November, 1978.