LABORATORY FOR
COMPUTER SCIENCE

MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

MIT/LCS/TR-577

# I-95

# THE INFORMATION MARKET

Tennenhouse, Abelson, Chan, Clark, Dertouzos, Hawley,
Gifford, Lippman, Liskov, Malone, Micali,
Szolovits, Ward, Zue (MIT),
Kung, Seltzer (Harvard),
Kleinrock (UCLA),
Broderson (UC Berkeley),
Myers, Reddy (CMU),
McNamara, Cheney, Enein, Graham, Unkauf, Welch (Raytheon),
Greif, Landry (Lotus),
Morse, Fuller, Griffin, Ladin, Lampson, Reid (Digital),
Anthony, Casey (Continental),
Johnston, Skrzypczak (NYNEX),
Lavenberg, Jaffe, Peled (IBM)

August 1993

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

# I-95

# THE INFORMATION MARKET

Tennenhouse, Abelson, Chan, Clark, Dertouzos, Hawley, Gifford, Lippman,
Liskov, Malone, Micali, Szolovits, Ward, Zue (MIT),
Kung, Seltzer (Harvard),
Kleinrock (UCLA),
Broderson (UC Berkeley),
Myers, Reddy (CMU),
McNamara, Cheney, Enein, Graham, Unkauf, Welch (Raytheon),
Greif, Landry (Lotus),
Morse, Fuller, Griffin, Ladin, Lampson, Reid (Digital),
Anthony, Casey (Continental),
Johnston, Skrzypczak (NYNEX),
Lavenberg, Jaffe, Peled (IBM)

Laboratory for Computer Science
Cambridge, Massachusetts 1993

# Foreword

This technical report describes I-95, a proposed collaborative effort concerned with the development of a national information infrastructure. Although the jury is still out concerning the future of the I-95 project itself, I believe that the timeliness of the topic and the importance of the architectural issues developed in the proposal are deserving of dissemination within the research community. Our colleagues in industry, government and academia, are invited to read this TR as a Request for Comments -- an attempt to initiate a national dialogue on the architecture and services of the emerging national (and international) information infrastructure.

Michael Dertouzos, the Director of the Laboratory for Computer Science, has been the principal instigator and one of the driving forces behind this proposal. In addition to their siginificant contributions as co-authors, Hal Abelson, Michael Dertouzos and Peter Szolovits have been the yeomen of the editing process. I-95 has benefited from the personal attention and support of many senior members of the Institute, including the President, Charles Vest, the Provost, Mark Wrighton, and Joel Moses, the Dean of Engineering. I would also like to acknowledge the efforts of our administrative and support staff -- Martin Harris, Paul Powell, Anne Wailes, Jennifer Watton and Tanie York.

Finally, although the text of this report is largely based on the original I-95 proposal, I take personal responsibility (and must apologize to my colleagues) for any errors or omissions that have been introduced during the editing process.

Cambridge, MA
August 1993

David Tennenhouse

# I-95

# THE INFORMATION MARKET

## ABSTRACT

I-95 will facilitate the free-market purchase, sale and exchange of information services. Driven by an alliance of universities—MIT, UC Berkeley, CMU, Harvard and UCLA—and companies—DEC, Raytheon, Continental Cablevision, NYNEX, IBM and Lotus— I-95 will coordinate an extensive network of interoperable telephony, cable, fiber, and wireless technologies; and will stimulate the growth of information services for individuals and for large and small businesses.

The uniqueness of the I-95 effort is rooted in our:
- commitment to an **open information market,** rather than a centralized set of brokered services;
- strict reliance on a crisp **infrastructure architecture** that extends far beyond network connectivity to include **modular shared services** that facilitate the rapid development and interoperation of new applications;
- technical activities, which are focused on the **critical technologies** needed to turn the architecture into reality;
- establishment of integrated **market development teams** that will catalyze new infrastructure uses in health care, software products and services, small business support, and other areas; and
- **instrumentation/evaluation** process for measuring performance and assessing impact.

# 1. Introduction

Lost in the current frenzy of extending high bandwidth computer networks to homes and businesses is the sober recognition that information highways as they are currently planned, provide only connectivity or, at most, brokered services. But connectivity does not constitute an adequate infrastructure for an information market if there is no coherent architecture to ensure the easy development of new applications. And brokered information services cannot be scaled broadly, since they cater by intent to the wishes of a few information suppliers. When a cable company offers a service for the home rental of video on demand, it does not generally put in place the infrastructure capabilities that will also allow individuals to sell substantive information services to others, because it does not perceive this as its immediate business interest. Yet, that capability is essential to a free information market.

For an information infrastructure to be truly effective, it must support a free market for information services, regardless of their origin, geographic extent or quantitative reach. This means that the information infrastructure must be endowed with characteristics that support and encourage its diverse users to easily purchase, sell or exchange an unlimited range of traditional and new information services, under their own initiative and for their own purposes.

The uniqueness of the I-95 effort is rooted in our:

- commitment to an **open information market**, rather than a centralized set of brokered services;
- strict reliance on a crisp **infrastructure architecture** that extends far beyond network connectivity to include **modular shared services** that facilitate the rapid development and interoperation of new applications;
- technical activities, which are focused on the **critical technologies** needed to turn the architecture into reality;
- establishment of integrated **market development teams** that will catalyze new infrastructure uses in health care, software products and services, small business support, and other areas; and
- **instrumentation/evaluation** process for measuring performance and assessing impact.

I-95 is an information infrastructure designed to facilitate the free-market purchase, sale and exchange of information services. Driven by an alliance of universities—MIT, UC Berkeley, CMU, Harvard and UCLA—and companies—DEC, Raytheon, Continental Cablevision, NYNEX, IBM and Lotus—I-95 will coordinate an extensive network of interoperable telephony, cable, fiber, and wireless technologies; and will stimulate the growth of information services for individuals and for large and small businesses. In addition to the founding partners listed above, the I-95 alliance may be expanded to include one or more of ATT, Hewlett Packard, Motorola and others with whom discussions have extended beyond the publication date of this report.

## 1.1 Infrastructure Architecture

The requirements of an information market can be met if the infrastructure adheres to an architectural coherence that extends well beyond the simple connectivity protocols of today's networks into conventions and shared services tailored to doing business in the electronic marketplace. These conventions must naturally and effectively serve a community that is large enough to make the development of infrastructure uses that depend on them easy and economically attractive.

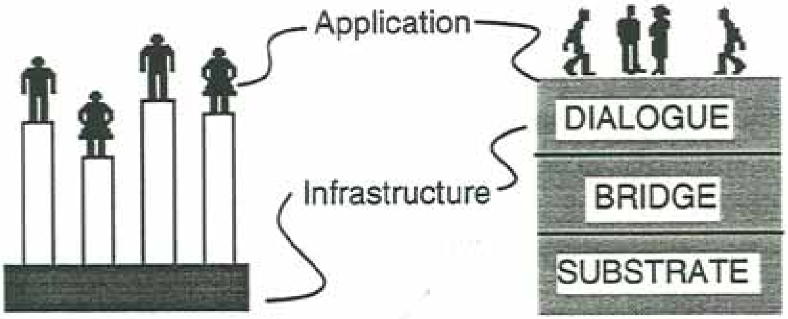

TRADITIONAL "THIN"
NETWORK INFRASTRUCTURE
Applications call for large
development efforts and do not
share much with each other

PROPOSED INFRASTRUCTURE
Applications involve less development effort and share services with
each other at several levels

In I-95, we propose a <u>three-level architecture:</u> At the lowest level, there is the connectivity <u>substrate</u> that delivers packets of undifferentiated bits by matching the needs of various infrastructure services to the available services offered by telephony, video, and wireless transport. For example, if users of a group conference service of I-95 require certain video capabilities, or a designated level of security, the substrate will strive to match these needs with some available transport service and engage the closest available alternative to which the conference service can gracefully degrade. Besides its technical functions of matching properties such as bandwidth, delay, security and reliability, the substrate is the level at which I-95 presses for a broad reach and an unhindered flow of traffic across substrate provider boundaries so as to stimulate competition and safeguard individual freedoms.

At the highest level, there is the <u>dialogue</u>, dedicated to human uses of the infrastructure and to the software agents that provide application-level services to users. Discourse at this level involves <u>scripts</u> that represent the information transactions and information flows that users need to carry out activities such as advertising and discovery, selection and negotiation, binding commitment and settlement. This is also the level to which outside services such as the SABRE airline system are connected. In short, the dialogue is where users and their computers interact on the basis of shared human concepts. *Book this flight. Send these x-Rays to Clinic x with highest privacy. What computer technician jobs are available in Boston? Send me software product y...* are examples of dialogue-level script fragments. Since computers can also interact with one another at this level, scripts may be automated—an infrastructure activity that is essential to human productivity enhancement.

Scripts are coupled to the user via suitable <u>user-interface systems and conventions.</u> This permits scripts and user interfaces to be used across a variety of different infrastructure settings. For example, browsing through a catalogue of products can be controlled by typing, by speech, or even by an automated software agent that negotiates on behalf of its human sponsor. Alternatively, a speech understanding script could be tailored to couple people to dialogue services that book flights, or that interact with a medical or jobs database.

At the intermediate level, between the dialogues and the connectivity substrate, there is the <u>bridge</u>—a set of infrastructure services from which dialogue-level applications are constructed. Bridge services, such as authentication, are used in implementing other higher-level bridge services, such as electronic conferencing. In fact, the entire I-95 infrastructure may be viewed as a rich collection of interconnected bridge services, the highest level of which form the dialogue, with its special provisions for people, and the lowest level of which form the substrate, with its special provisions that couple the infrastructure to existing network protocols. We expect that as

I-95 grows, so will its bridge services, advancing in the process the dialogue boundary to higher more useful human levels of interaction. The bridge includes a kernel, called the Information Mesh, that enables bridge services to be located and linked together.

We contrast the potential utility of the above infrastructure architecture with the fragility, rigidity and often chaotic nature of present network uses, where application-level software is built directly upon the connectivity substrate, in isolation from other applications.

## 1.2 Technical Activities

To make I-95 a reality, we plan to carry out the following technical activities:

**Specify dialogue level languages and libraries of common business practices.**
We will explore several possible languages and will decide whether there will be a single scripting language, or a variety of languages open to continuing evolution. We also plan to assemble a library of processes depicting common business practices to serve as the basis for implementing scripts. The library will be available on-line as a resource for I-95 participants.

**Construct a coherent Personal Information Environment**
The steadily increasing array of personal computational devices—timepieces, cellular phones, portable computers, beepers, and calculators—duplicate functions and interfaces and are not interoperable. Many of the applications emerging on I-95 will require increasingly sophisticated personal interfaces that combine speech, hearing, viewing displays as well as looking at the surrounding world, and tactile feedback. To bring some order to this personal system explosion we intend to develop a personal information environment—a combination cellular transmitter/personal computer worn on one's belt, a BodyNet connecting this unit via wireless low-power LAN to eyeglasses with integrated phones and microphone that can superpose computer-displayed images with real images. The idea is to provide the simplest interfaces where they are needed, i.e., where the fingers, eyes and ears are located, and to economize by having only one display. The possibilities of using such an environment are endless, since they cover the full spectrum of human-computer interaction. We will develop the architecture of such a device, prototype it and use it in a set of applications, to be determined.

**Implement authoring tools for scalable customizable interfaces**
Consistent with the I-95 interface isolation principle, this activity will provide authoring tools for the construction of personalized multiple-modality user interfaces, through specification of the dialogue elements and their interrelation.

**Experiment with human language-based interfaces**
This activity generalizes the speech understanding technologies developed at MIT to a level that can couple directly to different dialogue scripts, thereby bringing the power and ubiquity of interactive dialogue to the user. Human language extends beyond speech to gestures, writing and "mousing". We intend to explore the multi-modal aspects of such human language systems, where the multiple modalities strive to convey the same concept.

**Design essential mesh conventions and key bridge services.**
This minimalist kernel focuses on the key capabilities of the bridge level that permit universal access in a system capable of scaling to global proportions. We will publish a specification of mesh names to be used in I-95, and we will implement a prototype service for name resolution and bridge object location. We will also provide a prototype object catalog and a mesh navigation tool.

**Implement critical security and authentication bridge services**
I-95 will provide powerful mechanisms for controlling information access, for mixing public and private information, and for authenticating messages. Beyond providing services for privacy and authentication, we will explore the use of electronic money through tokens that cannot be easily duplicated, can be verified to be genuine and, when they pass from one party to another, are not possessed by both. We will experiment with uses of Fair Cryptosystems where multiple trustees are involved and with digital signatures with time constraints for bidding services. We also propose to address protection of systems attached to the I-95 substrate via a physical device that we call a security firewall, which we shall also prototype.

**Implement storage and access Bridge-Level Services**
The amount of information that will grow on the I-95 infrastructure will be on the order of millions to billions of information objects. In such a setting, finding the information one needs, being able to store information reliably and securely and operating efficiently with stored information become key requirements. Examples of what we plan to implement in this area include (1) a bridge service that allows users to browse, locate, and search for different types of information such as text, digitized video and audio, through use of a Semantic File System, a Content Routing System and a Structured Video System—all of which we shall prototype, (2) a large scale video server that can store many hours of full motion video files that are accessible by hundreds of clients, and (3) a bridge service that supports the secure, efficient and persistent storage of information

**Stimulate development of the I-95 connectivity substrate**
Through the I-95 substrate (1) we shall strive to make Ethernet-like access ubiquitous throughout the region by 1995 (Partners Continental and DEC have jointly committed to deploy the requisite access technologies); (2) DEC will develop a beta test substrate with local access at 6 Mbps cable, 1.5 Mbps T1, and 56/64 Kbps and (3) through its second-generation Ethernet CATV modem will support the I-95 basic interface; (4) Continental will make its regional fiber optic network, which passes most of the cable head ends in the region, available to I-95 toward a Sonet/ATM facility operating at 155 Mbps; (5) Raytheon will provide access to its private fiber optic network, which covers some 100 miles of Eastern Massachusetts; (6) the I-95 substrate will be supported by a distributed control center operated by Continental, Digital, NYNEX, Raytheon, and Teleport Boston (7) Raytheon and Digital will develop ATM network components; (8) Raytheon will develop a wireless system that supports a 20-100 Mbps shared access channel and allows individual users to burst at rates up to 6 Mbps; (9) Digital and Lincoln Laboratory are working on third generation wireless networking technology, which will also be tested on the I-95 substrate; (10) U.C. Berkeley will contribute a portable wireless pad that supports video, pen input, audio and text and graphics; (11) MIT will prototype an Integrated Services Packet Network, that supports controlled qualities of service; and (12) we will specify architecture and engineering guidelines that facilitate the continuous renewal of the substrate.

## 1.3 Market Development Teams

*This is not a test!* The I-95 network will be a major operational network that serves a major population center -- and the architecture to support an open-market economy must be tested against real applications. We will establish market development teams that will actively tailor the I-95 technology to specific markets and aggressively stimulate its adoption by user communities. They will encourage industries to take advantage of the I-95 architecture, help people who want to do this, provide the architecture team with stress-testing and feedback, and implement scripts for particular business applications. The initial teams will focus on the following areas:

**Software Products and Services:** The Boston software industry is active and ripe for electronic interactions. This market development team will pursue software companies, professionals, customers, and analysts to design network applications that (1) help buyers and suppliers of software find each other, negotiate deals, and place orders electronically; and (2) distribute the products.

**Health Care Products and Services:** The health care market development team will focus on the interactions among key players such as patients and their families, doctors, HMOs, testing laboratories and pharmaceutical companies. We will begin this effort by pursuing a full video and multimedia approach for patients requiring extended home care.

**Small Businesses:** I-95 will reach small businesses that have traditionally been slow to modernize operations but that contribute greatly to the region's business climate by supporting the large corporations that encircle them. The Small Business market development team will define a minimal "infoport" made up entirely of hardware and software components that can be acquired at local retail outlets and connects easily at the I-95 dialogue level.

**Other Markets:** In addition to specific application areas we will also make several capabilities available to all users of the network, including the general public. The market development team pursuing these applications will stimulate (1) the matching of employers with people seeking jobs; (2) want ads in a variety of sectors; (3) educational activities, especially training and courses at a distance; and (4) cultural and personal development activities.

## 1.4 Instrumentation and Evaluation

UCLA will develop a Center for I-95 Substrate Measurement, Evaluation, and Design that will provide traffic measurements, patterns and trends as well as data and tools that facilitate the modeling, evaluation and design of the substrate. MIT economists will assess the productivity impact of I-95 through business performance measures and estimates, and MIT social scientists will assess the social impact of I-95. Such vigilance will be most useful in the evolution of the I-95 information infrastructure, which will be a major new medium of social discourse.

## 1.5 Regional to National Aspects

The I-95 architecture has been shaped by the need to facilitate national coordination while at the same time encouraging rapid and vigorous innovation. We are concerned that the major obstacle to the emergence of a truly national information infrastructure will be architectural, rather than physical. The capacity of substrate-level services will steadily improve as carriers roll out their competing services. However, the haphazard deployment of diverse architectures and user interface paradigms within the individual regions would necessitate a period of consolidation and retrenchment. To avoid this, we will actively encourage the coordination of regional activities and we will work closely with the regional alliances, both on architectural issues and on the exchange of software and services.

Through gateways to the existing Internet, I-95 subscribers will be immediately linked to its large reservoir of information services and, by extension, to the other regional information infrastructures across the nation. Assuming that such a combined infrastructure retains architectural coherence, its uses will be more significant than those of a single regional infrastructure, in exactly the same sense that a telephone system spanning the nation is more useful that one spanning only part of the nation.

# 2. Approach

## 2.1 An open information market requires architectural clarity

After a thirty-year evolution from the ArpaNet to the Internet to the Information Highway, computer-based mass communication has become the defining technology in the transition from the 20th to the 21st century. Plans to extend high-bandwidth computer networks to homes and business are unfolding at a frenzied rate, backed by cable companies, telephone companies, publishers, schools, libraries, shopping services, and governments all eager to participate in the Information Market that these information highways will engender.

Lost in the clamor, however, is the sober recognition that information highways as they are currently planned, provide only connectivity. Connectivity does not, in and of itself, constitute an adequate infrastructure for a free-market economy. In particular:

- An information market requires secure foundations that can support robust and scalable control of information access and sharing. The current Internet foundations, put in place largely with the original ArpaNet, are inadequate for serious commercial use. Today, it is worrisome and ironic that, as Internet functionality increasingly expands, companies are increasingly hiding their internal networks behind gateways that reduce their Internet connections to little more than mail delivery services. This trend flies in the face of the needs of an information market that requires transmission of sensitive information such as sealed bids and patient medical records and flexible access to databases that mix proprietary and non-proprietary information.
- An information market requires modular elements from which high-level services can be constructed in a mix-and-match way. Lacking these, substantial information services become complex, monolithic systems, and constructing them becomes an expensive proposition in which only a few players can participate. A national marketing chain might develop an interactive shopping service, but the corner supermarket could not adapt the catalogue browser from that service to advertise its goods in the neighborhood, nor adapt the payment processing module to automate its own billing. Nor is there room for a third-party information broker to provide automated comparative shopping. As we rush to construct information highways, we should consider whether we are erecting an infrastructure that permits millions of individuals to use their vehicles on the highways for their own purposes and at their own discretion—or whether the highways will be restricted to a few large agencies that broker busses and trucks to meet everyone's transportation needs. As straightforward as this point should be, especially in a free market society, it is nevertheless a source of great confusion among many of today's major players as they try to position themselves for increased revenues through new services.
- The essential point is that an information market requires architectural coherence in the information infrastructure, well beyond simple connectivity protocols such as TCP, FTP, and SMTP. This coherence includes not only protocols for higher-level services, but also conventions for doing business in the electronic marketplace—such as how to pay for goods and services or how to access catalogues of information. These conventions must be agreed upon by a community large enough to make the development of computational services that depend on them economically feasible and attractive.

The orientation, sophistication, and ability of the emerging national information infrastructure to meet the nation's changing requirements have yet to be determined. The issues are complex and their resolution will require the synthesis of technology, management, and policy considerations, even while commercial products and services are rapidly being rolled out. The national information infrastructure is now being forged, and its shape will be determined by the alliances that actively participate in the smithing process.

I-95 is an alliance of universities, computer companies, and communications companies that will assemble an extensive operational network in the greater Boston regional area (I-95 by 1995), including home and small business LAN (Ethernet, Token Ring) service and a variety of ATM and wireless services. We will use this network as a testbed for refining architectural coherence, for field-testing the crucial protocols and modules from which high-level market-oriented services are assembled, and for catalyzing an open and active information marketplace encompassing thousands of individuals and businesses.

## 2.2 I-95 is based on a three-level architecture

I-95 identifies a three-level architecture as a framework for protocols, conventions and services:

- At the lowest level, there is the connectivity substrate. This is the basic connectivity service that delivers packets of undifferentiated bits. The protocols are essentially already fixed by standard commercial use, although we hope to influence these to provide better functionality through quality-of-service and security specifications.

- At the intermediate level, there is the bridge—the mesh of object repositories, security, and information sharing services from which dialogue level applications are constructed. Indeed, the fragility and rigidity of present information market applications derives largely from the fact that current applications are built directly upon the connectivity substrate, without the well-defined intermediate levels of abstraction that the bridge provides. I-95 will specify the basic bridge-level protocols, implement some of the essential bridge elements and services, and establish the foundation for other elements and services.

- At the highest level, there are the application dialogues conducted among human users of the network and software agents that provide application level services. Our goal here is to provide a framework that permits applications to be easily assembled from modules that perform market-related functions. The key idea is to formalize these functions as scripts that describe the activities of users in the information market.

Section 3 describes this architecture in detail. Here we indicate the major I-95 challenges for work at each level:

- At the dialogue level we will construct a language of scripts and we will implement modules that permit individual script acts to be automated or supported by software agents. Crucial to the architecture of the dialogue level is the stipulation that the acts themselves are isolated from the user interface by suitable interface protocols and so can scale across a variety of network bandwidths and user-interface modalities. (For example, browsing through a price catalogue can be controlled by typing, by speech, by gesture, or by an automated software agent that works on behalf of its human sponsor.) To support this commitment to multiple modalities, we will also build a coherent personal information environment whose capacity for multimedia input and output goes far beyond the personal digital assistants now appearing on the market.

- At the bridge level we must develop the framework for an entire mesh of information objects and the services they support. One critical service here is security, and we will implement such services as time-limited digital signatures, shared secret information, electronic money and messages that cannot be repudiated. We will also explore systems aspects of network security, with the objective of providing security firewalls in the network in an application-independent manner. Another service we will implement is automatic content-based indexing of databases to support content-addressable file systems. One key to the bridge level is that it supplies typed data objects in a universal location-independent namespace. We will specify the

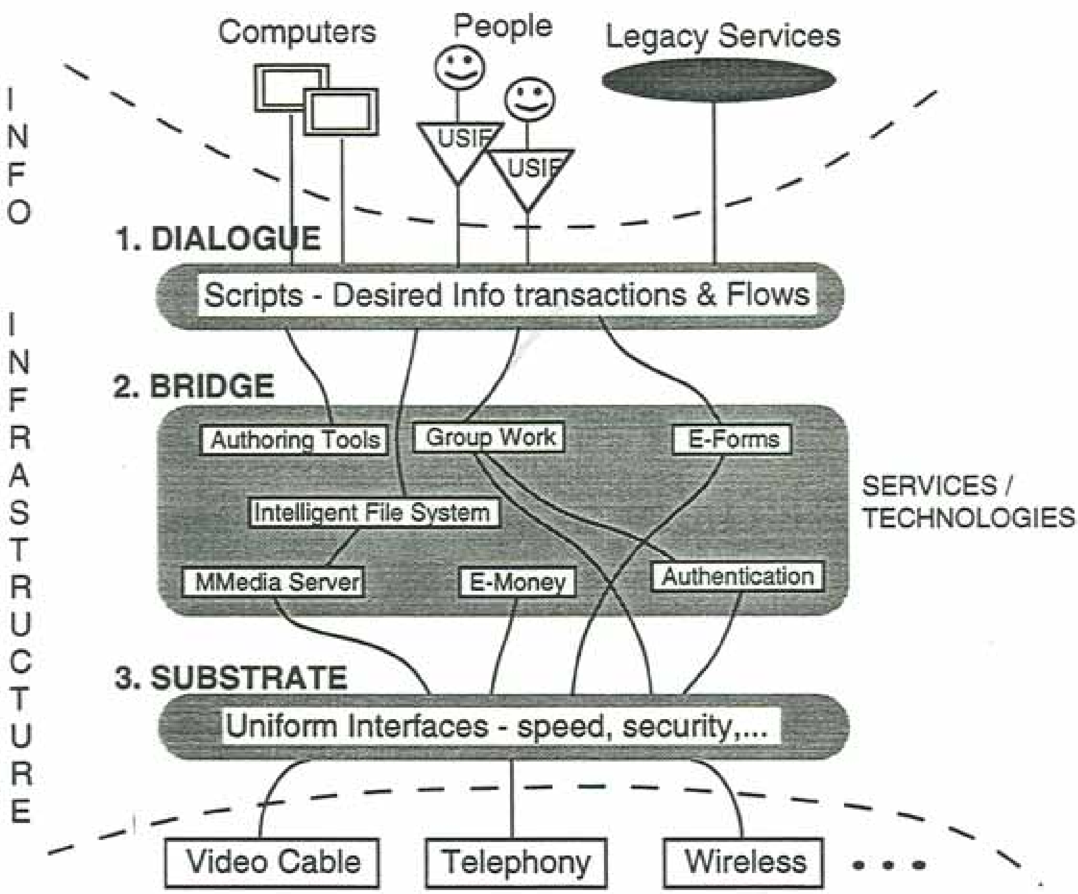

protocols to support this, and present prototype implementations of the basic mechanisms. This kernel implementation, which we call the Information Mesh, is currently under development at LCS.

- At the substrate level, we recognize that the creation of a national information infrastructure has yet to achieve anything resembling a critical mass of connected users. Substantial long distance transmission facilities have been deployed, but local connectivity is still largely restricted to low-speed / fixed-bandwidth access. I-95 will catalyze rapid commercialization of variable-bandwidth local-access technology and make such access technology ubiquitously available. We will catalyze a real substrate, not a limited prototype, pilot, or field trial. The substrate will make a range of connectivity options available, not just to large organizations, but to the vast majority of the small businesses and homes within the region. Although some locations may be restricted as to their access options, the bulk of the population will be offered the opportunity to subscribe to Ethernet-quality service at homes and offices. The I-95 principals will work together to develop advanced wired and wireless access technologies and to aggressively promote subscription and market development. Although I-95 may arrange bulk purchases from the service providers, the access services will be operated on a commercial basis, outside of the funding envelope of the project.

Two catalysts of competition within the computer industry, have been open architecture and short product cycles. These catalysts have direct infrastructure analogs: open interconnect and renewable deployment. Open interconnect will bring the benefits of open architecture to the full range of I-95 products and services, ensuring open markets that encourage vigorous innovation.. Renewable deployment is an approach to providing communications infrastructure that permits the ongoing and frequent enhancement of rapidly evolving components.

# 3. Architectural Framework

## 3.1 Overview

The I-95 architecture provides a systematic organization for a large, multi-partner effort; identifies a set of key ideas according to which other, more detailed design decisions are to be made; and defines base-level capabilities that are commonly used by all and that can be extended to provide additional functionality. A successful architectural design uses a sparse set of mechanisms to provide a rich collection of capabilities, by carefully identifying what is central and primitive and by providing simple, powerful abstraction and extension mechanisms that make it easy to build upon the core mechanisms.

In section 2, we introduced the three major levels of our architecture. The *substrate* supports the exchange of undifferentiated strings of bits. At the intermediate *bridge* level, structured information is exchanged among a set of services. These elements are realized by the *Information Mesh*, which provides naming and descriptions of access methods for all information. We will provide architectural support for mobility and security of information, and key examples of structured information objects, such as electronic money. Finally, the uppermost *dialogue* level models interactions in the information market at the level of human concepts. This level encompasses scripts for performing common generic types of activities, such as searching for information, and languages for expressing specific activities that may be defined as parts of applications.

We expect to gain experience with all aspects of the I-95 architecture, by studying its use by our customers, and through the activities of the market development teams. We will refine the architecture of the connectivity substrate to accommodate the insertion of advanced technologies, including ATM switches and wireless systems. Bridge level information services will be defined and built both in response to the recognized need for services such as authentication and binding commitments, and in response to additional emerging requirements. At the dialogue level, we will enrich the architecture by identifying common marketplace activities and by defining languages of *acts* that will make it possible to express those activities in a common form that is meaningful to human participants. A key requirement of the dialogue level is that modules, once implemented, become new components available as new building blocks.

The development of dialogue and bridge services will occur over time. Early applications will be built directly on the substrate, much as they are today. However, insertion of higher-level architectural capabilities, analysis of the needs of early applications (including existing applications such as Lotus Notes), and deliberate attempts to learn abstractions by the market-development teams will swell the stock of bridge and dialogue level services, and will make it more attractive for second-generation applications to use such capabilities. We will "seed" higher levels of the architecture with important services and proselytize their use.

As higher level services are added to the architecture, competing approaches to any particular service will emerge. For example, at the bridge level there may be several different methods of transferring money (*e-money*) via the network: one may simply extend existing methods of processing credit or debit charges; another may extend interbank transfer protocols to non-bank participants; a third may be based on cryptographic methods of providing authenticated but irreproducible tokens that take the place of cash. Rather than trying to pick ahead of time which is the "right" form of e-money, the I-95 architecture will support several alternatives. Market forces will facilitate convergence on a small number of variants, each surviving because of some specific desirable feature. For example, cryptography-based tokens can easily be made non-traceable, just as cash is today, and this may make them particularly suitable for spending related to activities in which privacy is a significant concern. Debit transactions may be more appropriate in financial transactions that must inherently be traceable and verifiable, and where the pre-existence of sufficient funds to cover the transaction is critical, such as in certain securities transfers. As services become widely adopted, they help to define a *thick* architecture, which then provides not only the basic services needed to build others, but also a large variety of specific, proven and accepted methods to accomplish higher-level tasks.

## 3.2 The Dialogue Level

At the highest architectural level, our main goal is to modularize normally monolithic applications into reusable component facilities. Many advantages accrue to success in such an endeavor. New applications become easier to build as a library of components is identified and implemented. One can improve the operation of applications without rebuilding them completely, by replacing their individual components by more powerful variants. For example, an improved textual search facility could improve a range of existing applications, from those supporting library access to processing orders for electronic parts. Similarly, a language translation module could be incorporated into an existing electronic mail application to facilitate communication with correspondents who speak only foreign languages.

The dialogue level of the architecture is based on the following commitments:

- The ordinary activities of both human and computer-based agents follow a *script*, which represents a sequence of *steps*. These steps correspond either to primitive conceptual operations at the user's level, called *acts*, (e.g., sending a specific message to another specific user) or to another script, whose steps in turn consist of more primitive scripts and acts. Acts will themselves be complex operations at some lower level of the architecture, but appear unitary at the dialogue level. There must be one or more scripting languages that permit the specification of scripts in terms of their components, and that support typical programming language capabilities such as abstraction and flow of control.

- All scripts have an interface that allows them to be used not only as parts of dialogue scripts, but also by programs at the bridge level of the architecture. This permits the incorporation of activities that begin as dialogue level scripts into lower levels of the architecture.

- Information content must be separate from the mode of interaction of the user-interface. A given set of information may be conveyed via text, speech, graphics, or video depending on the interface, but still correspond to the same communication act. Input may arrive from keyboard, pen, microphone, hand-gesture, or eye-tracker, and the script that desires the input must not be restricted to only one input mode. In addition, any "user interface" act can be carried out equivalently by a computer program, supporting the construction of computer-based agents.

- The dialogue elements that are the inputs and outputs of acts are themselves dialogue level manifestations of objects at the bridge level.

- Coordination among multiple (human or computer) agents is not a primitive capability, but is provided through the scripts followed by those agents. As a corollary, scripts that perform cooperative work must be designed to work in the absence of a global state of knowledge for all agents. Certain communication acts employ bridge level facilities, such as guaranteed delivery, to ensure reliable coordination.

### 3.2.1 Dialogue example

The key concepts are best illustrated through an example, for which we use a next-generation version of the Pegasus system under development at MIT. Pegasus supports the booking of airline reservations through the American Airlines Sabre system. Interactions are conducted by spoken language—the system performs speech understanding and speech generation.

The parties to a dialogue in Pegasus are: the user, a dispatcher, which is a relatively primitive form of software-based agent; and the Sabre system itself. The user's interactions are mediated by a spoken language user environment that isolates this particular I/O modality (speech) from other elements in the system. Similarly, the Sabre reservation system is encapsulated in a legacy environment, thereby insulating the dispatcher from the arcane details of this external system.

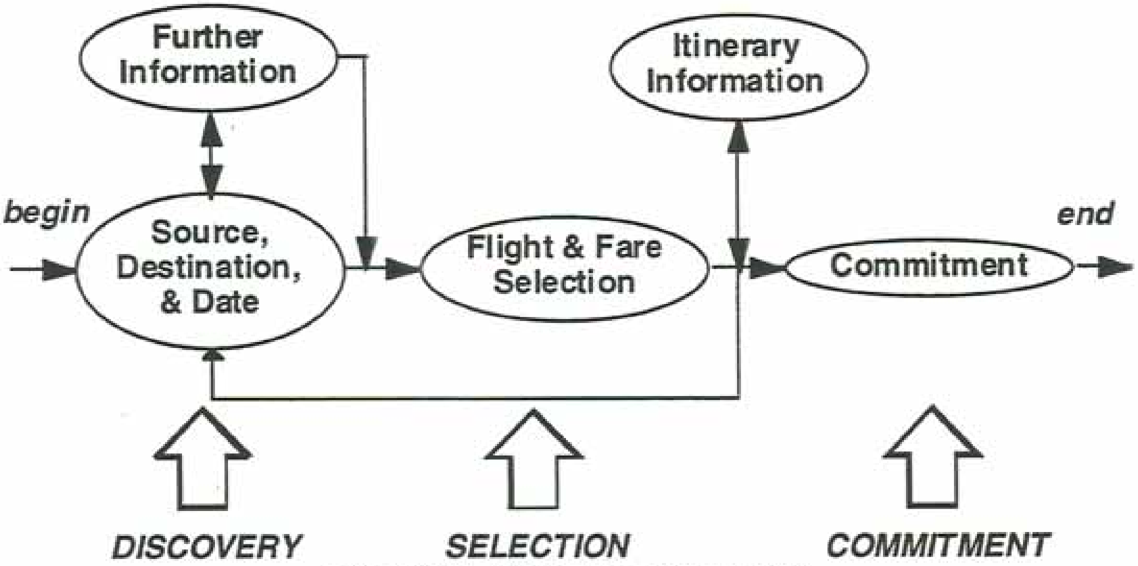

Figure 2-1: Airline reservation script

During the course of a transaction dialogue elements are exchanged in accordance with the reservation script illustrated in Figure 2-1. This script is not hypothetical—it represents the logic embedded in a working system that has been extensively tested. The script is divided into steps such as discovery, selection, and commitment, which may be reusable within the context of other transactions that are not necessarily related to airline travel.

Steps within scripts are themselves aggregates of other, more detailed scripts and acts. For example, "Flight & Fare Selection" involves sub-components such as generating lists of available possibilities, imposing constraints to eliminate untenable flight schedules, and ranking based on criteria such as cost and convenience. It will often be useful to define more abstract scripts. For example, though details vary, the selection of a vendor in a generic purchasing script may share the same overall structure.

The creation of such a sophisticated architectural level in a widely used information system is highly ambitious, and will undoubtedly call for successive refinement. Nevertheless, we believe that I-95 provides a suitable starting point to explore the use of a dialogue level, and that the emergence of the information-based marketplace signals the time at which this challenge must be addressed. We plan to define a script language that can stitch together applications from sophisticated components that are in turn built of bridge level services. Given such a language, and an appropriate understanding of the market-specific processes, it should be possible to develop software that will generalize the Pegasus approach to a wide range of applications.

There is a great richness of scripts to be created. Many people assume that there is a single, basic process for buying and selling things. In fact, there are important variations in how market processes work can work, and one of the important contributions of this project will be the development of market-making software that can be easily adapted to support many different kinds of market processes. Markets differ, for instance, in which party initiates the transaction. In some cases, sellers initiate the process by advertising the availability of their products. In other cases, buyers initiate the process by announcing their desire for a product or service. Other important characteristics of markets include whether product descriptions are standardized and whether quality ratings of products are available. The information market must also support people who wish to collaborate. Although it is tempting to define global coordination services into the architecture of I-95, our inclination is not to do so. Instead, we will rely on explicit scripts for coordination to be used by cooperating agents, supported by primitive communication acts implemented on the bridge-level substrate.

### 3.2.2 Dialogue Activities, Scripts and Acts

To automate a large part of the software creation process, it will be necessary to develop a collection of generic dialogue scripts that perform different activities—i.e., that support generic phases that recur across a variety of business processes. Examples of such activities are:

- Advertising and Discovery: These are activities in which both buyers and sellers identify and find information about each other. For instance registration services help users register themselves and any products or services they offer or might wish to buy. Discovery services entail the equivalent of white and yellow pages.

- Selection and Negotiation: In these activities, buyers and suppliers evaluate each other, agree on prices or other terms, and place orders. We are particularly intrigued by the possibility that parts of these processes may be automated. For instance, "buyers' agents" might automatically screen large numbers of product offerings to identify the most promising alternatives for their human "owners" to consider.

- Binding Commitment: This activity allows two or more independent interests to negotiate a transaction whose terms are verifiable by each involved party.

- Fulfillment: In some cases, i.e., with "information" products such as software, movies, etc., the delivery of a product or service can occur electronically over the network. In other cases, the infrastructure will be used to schedule and track the delivery of physical products.

- Settlement: A particularly important part of an information-based economy will be the billing and payment for services. The benefits of an information market will be substantially greater when money becomes an integrated part of the infrastructure, i.e., when it is directly supported by services operating at the bridge level.

Many scripts (e.g., "Selection", above) will be defined in terms of generic steps. Other scripts will accomplish more specific activities that are specializations of the more general ones (e.g., "Vendor Selection"), and these more specific scripts will normally consist of more specific steps (e.g., "List All Vendors" rather than "List all possibilities"). The script language must therefore provide capabilities to specify constraints and relations among the steps of a generic script which will be applied to their instantiations in more specific scripts. These are the kinds of capabilities found in some object-oriented programming languages, and will need to be supported in the bridge-level implementation of the objects that represent scripts.

### 3.2.3 Dialogue Elements

Dialogue elements capture the information exchanged by participants in an ongoing dialogue. Here we are concerned with communication at a high level, where the abstract representation of the information is potentially meaningful to human users of the system. During the transition to an information-based economy, we will repeatedly confront the question of how to formally represent the questions and statements that are exchanged as part of the dialogue among peer agents, i.e., users and their applications. We are concerned here with the logical form of statements and questions, not with the particular user environments (speech, gesture, etc.) where they are recognized and rendered. We distinguish between two options, the *text* option and the *context* option.

The text option assumes that the information stored within the system is not structured except for very low-level conventions—e.g., that all text be recognizable as text and encoded in ASCII. Many knowledge sources in I-95 will be textual, at the start, and there will be many opportunities within I-95 to build useful and elegant systems that exploit such information, even though the functionality will be restricted to text-based searching such as keyword and word-stem search.

The context option assumes that some higher-level organizational principles are applied. In the longer term, the value of pre-defined structure will become more and more apparent. However, to be accepted in the workplace, it is important that the introduction of new dialogue elements follow a progression in which users gradually accommodate themselves to more sophisticated tools. The following paragraphs describe a progression of formalizations:

- <u>Electronic Templates (E-Forms):</u> Flat-file database management programs, such as FileMaker, support many simple applications quite adequately. This is a popular means of both entering and searching for information. Although both the input and retrieval languages are quite weak, many people develop good intuitions for using them.

- <u>Relational Databases:</u> These provide a clean and powerful extension of the flat-file database, in which multiple values are easy to store systematically, and virtual views can be assembled as needed to answer queries that may not have been anticipated in the original design.

- <u>Frame-Based Knowledge Representations:</u> The simplest knowledge representation ideas focus on capturing taxonomic knowledge and on structured relationships among complex objects. The structured relationships support notions such as default values for unknown properties of objects and inheritance of general characteristics from supertypes. One observation concerning such KR systems is that focusing on what are the "verbs" (or "actions") and what are the "nouns" (or "objects") being represented is a useful point of departure. Pegasus uses this type of dialogue element to support communication between the user and the dispatcher.

- <u>Logical Formalisms:</u> The desire for greater expressiveness argues for representations that can say anything that is sayable. Many formalisms based on various logics have been proposed to support this desire. They typically trade off expressiveness for decreased tractability or computability. The AI community is in the course of developing a variety of such formalisms.

### 3.2.4 User Environments

We have emphasized as an architectural principle that the mode of interaction between human user and computer is orthogonal to the information-content of the communication. This means that different methods of interaction can all coexist and make use of the same dialogue scripts and services. Therefore, a user moving among different environments can retain a coherent view of the underlying capabilities provided by the system. Higher quality user environments will empower the user, and we intend to pursue aggressively the development of interfaces based on human language and coherent personal environments.

<u>Human Language-Based User Environments:</u> One of the critical research issues that our project must address is how to develop an interface mechanism that will enable users to efficiently access, process, and manipulate a vast amount of information. We believe part of the answer lies in an interface based on human language, since it is the most pervasive, natural, efficient, and flexible means of communication for humans. In developing such an interface, we will be guided by several principles. First, the interface will permit a large number of input and output modalities; depending on the task, users should be able to choose among keyboard, pen, microphone, or hand, to convey their wishes and thoughts, and to receive the information in text,

speech, graphics and video forms. Second, these modalities will be unified by a set of language-based principles of communication, such as dialogue (e.g., turn-taking, discourse) referencing (e.g., pointing, use of pronouns), and speech acts (e.g., issuing commands, asking for clarifications). Third, a great deal of emphasis will be placed on modeling the users, both in terms of their needs and behavior.

Coherent Personal Environments: Consider the tangle of the global digital infrastructure—Internets, future video and entertainment services, and the forest of media appliances, and personal gadgets. At the moment, a hi-tech yuppie may be carrying a pager, a cell phone, a digital wristwatch (with built-in weather station), a "wizard" diary or "digital assistant," possibly a laptop computer, not to mention a "walkman," a "watchman," a camera or video camera, and remote controls for car alarms and home stereo systems. The person is festooned with computers—and none of them talk to each other. The consequent level of redundancy is appalling. Count the displays: there may be 5 or 6, with keyboards and clocks scattered throughout. There is no coherent architecture that makes it as easy to "snap together" personal appliances as it is to plug bigger machines into the rest of the connectivity substrate. What is called for is a modular set of components, and a lingua franca so that two-way conversational intelligence can be built easily into any appliance or adornment. MIDI (for digital music instruments), SMPTE and VISCA (for video and media devices), and "Hayes modem-speak" are all crude instances of this trend.

## 3.3 The Bridge

The Bridge level of abstraction, which sits above the connectivity and storage services offered by present day networks and operating systems, will facilitate the development of a new generation of applications. The Information Bridge will be of long-lasting national significance, beyond the geographic scope and duration of I-95. It is possible to construct applications without an information bridge, and many successful applications have been constructed directly on top of transport protocols such as TCP. However, these applications can be used only in isolation. Without a common service framework, they cannot interoperate, they cannot easily exchange data, and they cannot naturally become a part of our larger overall objective, an information-based market economy. For these larger goals, a common set of services, protocols and conventions is critical. In the following sections, we describe the information mesh that constitutes the kernel of the bridge level interface. The description of specific bridge level representations and services that must be acquired and/or developed within the project appears in the activities section of this report.

### 3.3.1 The Information Mesh

Just as the Internet protocols name hosts, to permit transport between them, our protocols must name bridge level objects so that they can be linked together, and so that their location can be determined on demand, even though they may have moved in the physical infrastructure. There are two key aspects to our mesh interface. First, bridge objects must have an *identifier* by which they can be named and located. Second, they must describe, in an agreed manner, the *access methods* that can be used to interact with them. We need to allow for variability in access methods, but the key requirement is that there be a uniform way to ascertain what method the bridge object supports. A uniform first step is necessary to begin the interaction There is almost an exact analogy here with the Internet protocol, which carries only two important facts in its fields: the address of the destination and the type of the protocol in the packet.

Some existing distributed applications provide a form of object name. However, the name spaces in use today are far too limited to scale, either in space or time. For example, the name space of the Gopher system uses Internet domain names and UNIX file system path names to identify objects. This means that an object cannot move within a single file system hierarchy, let alone to another machine.

A uniform name space does not mean that we demand a single location service. Quite the contrary, we must expect that in the life of the Information Mesh, both the physical instantiation and the underlying approach to storage and name resolution will change. Thus, our mesh name space has both the concept of a unique id, and the concept of hints, which provide the practical means for location. A key objective of the mesh is that information can be moved, both physically and within lower level logical organizations such as file systems. Mobility is critical, both to deal with the short term dynamics of mobile computing and the long term dynamics of machine upgrades and changes in information stewardship, e.g. from creator to archivist.

Objects stored in the mesh will support a wide range of access methods. At the dialogue level, every object should have some means to permit users to gain access to its contents. However, at a lower level, there will be a range of information representations. Some may be simple file-like behaviors, that just return their contents in a standard format, such as JPEG. Others may use existing access protocols, for example a data base with SQL. And, not surprisingly, we expect that our own teams and others, will invent new data representations and access protocols. The objective at the mesh level is to provide a common means of interrogating an object to determine what interactions it will accept, both at the dialogue level (the dialogue scripts and elements with which it will interact) and at the low level (the base representation of its visible parts).

Security: The mesh should enforce well-defined and clearly articulated domains of accessibility to every piece of structured information; moreover, effective controls over accessibility must be supported by basic, mesh-wide guarantees. The current tendency to isolationism by many companies, related to the absence of such provisions on existing networks, constitutes a serious deterrent to the growth of an information-based economy. Our approach is to define an architecture for security firewalls, whose presence in the communication is assured at the mesh level, and which provide security assurance at the dialogue level. The firewalls ensure that only valid dialogue elements consistent with the intended script are passed across the firewall. In this way, they operate much as mail gateways today, but in an application-independent manner.

Cost Visibility: Means may be provided by which potentially expensive (e.g., time-consuming) operations may be forewarned by preliminary cost estimates. A corollary principle is that, notwithstanding the technological isolation implied by scalability and ubiquity requirements, technological parameters are program visible. The intent here is to guarantee that software can adapt to technological limitations in an intelligent way—e.g. by changing defaults to accommodate them. This objective requires support at all three levels, dialogue, bridge and substrate. However, it is at the bridge level, where links among objects are implemented using mesh names, that basic navigation is implemented. Thus, the bridge must be a core part of performance visibility.

## 3.4 The Connectivity Substrate

The connectivity substrate supports the exchange and storage of undifferentiated strings of bits between peer processes and/or threads of program execution. We distinguish this level for the very pragmatic reason that it roughly corresponds to the program interfaces supported by modern operating systems. The industry is converging on a small number of de-facto standards , e.g., NT, OSF, etc. Though they differ as to their detailed specifications, their interfaces supporting communication, storage (i.e., file access), and transduction (i.e., I/O) loosely correspond to the functionality of our abstraction. A similar degree of convergence has been achieved amongst the protocols supporting these interfaces, e.g., TCP/IP, NetWare, NFS, AFS, etc. At the substrate level, I-95 has two primary goals: to catalyze immediate and ubiquitous access; and to pioneer the architectural refinements and technologies essential to its continued success.

### 3.4.1 Achieving Critical Mass

To achieve widespread connectivity, we must address two issues: the provision of local access, and the interconnection of access service providers.

Ubiquitous Access: Ubiquitous access to IP-like connectivity will be crucial to the success of the regional networks. We put great emphasis on variable bandwidth access because of the bursty demands of modern software-based systems. These bursts directly reflect the work habits and expectations of our users. When they drop video file icons into folders they rightfully expect the system to respond to the burst of activity that is initiated. Although the average bandwidth requirement may be quite low, an access arrangement supporting multi-megabit per second bursts yields a qualitatively different experience than one supporting multi-kilobit per second access. Large organizations have long benefited from variable bandwidth connectivity services, such as the Internet, but this has not been true for smaller organizations and/or individuals. Service providers have been reluctant to deploy technology in the absence of demonstrated demand and the availability of suitable customer premises equipment.

I-95 resolves this deadlock. Our access interface will be based on LAN standards, such as Ethernet, for which customers can acquire adapters at neighborhood vendors, such as CompUSA. The large installed base of platforms supporting these standards, and the commercial availability of interfaces, directly address the carrier concerns. The separation of our access interface, initially Ethernet, from the underlying distribution technology is an important feature of our substrate architecture. Competing carriers can support the same interface and, at the same time, can adopt the distribution technologies that are best suited to their existing plants.

Unfortunately, the widespread availability of customer equipment is only part of the access problem. Large organizations negotiate bulk purchases that justify the custom engineering of their access arrangements. However, the high costs associated with such customization create a high barrier to entry for smaller purchasers. To provide them with affordable access, the service provider must be prepared to provision entire streets, thereby incurring risk and capital expense, with no return in advance of subscription. In addition, small customer concerns are exacerbated by the need for variable bandwidth access. If a dedicated access medium is used, then statistical sharing cannot be leveraged—the access medium must be provisioned in accordance with the maximum, rather than the average, rate of each customer.

In the case of large multi-tenant buildings, I-95 service providers could speculatively install a shared facility of moderate bandwidth and upgrade its capacity as subscription increases. In this case the local distribution system within the building may partition the users so as to ensure the integrity of the shared facility. Analogous approaches could be used to support variable bandwidth access from houses and medium sized buildings. For example, a cable television provider might initially dedicate a single 6 MHz television channel to I-95 access. As subscription ramps up, they can increase their capacity in a number of ways, for example: additional spectrum, i.e., adding channels; geographic re-use, i.e., partitioning the neighborhoods into cells; or the adoption of more efficient modulation techniques.

This shared approach facilitates broad geographic coverage, while minimizing the initial outlay that must precede widespread subscription. Although this initial substrate will be incapable of absorbing a large volume of users, the ramp-up in subscriptions will establish a proven revenue flow that can be used to finance its growth. Our approach is the modern day equivalent of the party-line, which broke the same deadlock in the early days of telephony. However, packet technology introduces a qualitative improvement: whereas party line access was an all or nothing proposition, at least for the duration of a neighbor's phone call, the fine granularity of packet multiplexing largely obviates this source of neighborhood tension.

Open Interconnect: There is a long-standing tradition of ensuring access to infrastructure services through the institutions of common carriage and universal service. Open interconnect, the unhindered flow of traffic across service provider boundaries, is a logical successor to common carriage. It is a policy approach that lends itself to self-regulation (as opposed to external regulation). It will stimulate competition where possible, and safeguard individual freedoms where economies of scale and scope lead to natural monopolies. Open interconnect is a key tenet of I-95 and we are committed to working together to identify and implement mechanisms that will ensure the smooth flow of data across organizational boundaries. Although this may appear to be a relatively straightforward issue, long ago dealt with by the Internet community, it is complicated by the large number of players. While NYNEX serves the entire region through a single subsidiary (New England Telephone), every town is served by its own cable operator, or MSO. Although Continental owns or operates a large fraction of the region's MSOs, many are owned by other companies. I-95's interconnection of these MSOs, and their further connection to NYNEX and the Internet will pose many interesting problems, especially in the areas of billing, management, security, and provisioning.

### 3.4.2 Advanced Technologies and Architecture Refinement

Variable Bandwidth Wireless Access: I-95 will pursue the development of wireless transfer modes that are consistent with our variable bandwidth access philosophy, yet compatible with existing services such as telephony. We are determined to head off a repetition of the unhappy experience of digital wireline technology, in which the fixed bandwidth transfer mode of traditional telephony is only now being reconciled with variable bandwidth data services.

Enriching the Connectivity Interface: The networks of today offer essentially one quality of service, which is "best effort" delivery, with no assurance as to the security, throughput or delay that will actually be achieved. While this service has been successful in supporting a wide range of applications, new applications such as real time video and audio may require more controlled qualities of service, especially with respect to delay, bandwidth and loss characteristics.

Renewable Substrate: Present approaches to infrastructure renewal are incompatible with the rate of technology development. Halfway through a twenty-year deployment period, the technology being newly deployed is many generations obsolete. Not only is this style of deployment inappropriate during the present period of rapid innovation, it is wholly incompatible with our approach to local access, in which carriers will initially provision for sparse subscription and rely on the rapid diffusion of new technologies, such as improved modulation techniques, to cope with increasing subscriber density. We will explore the design of renewable substrates whose modularity and topology are tailored to the frequent enhancement of the components that are evolving most rapidly. For example, a large component such as a monolithic one-hundred-thousand-line switch, may pose a barrier to renewal that could be averted through greater modularity. Inappropriately positioned modularity (e.g., the embedding of thousands of small components in difficult to reach places) is also problematic. Finally, there are areas, such as compression, where speed-ups arising from new algorithms are as dramatic as those arising from VLSI refinement. In these cases, the renewability of software-based implementations may prove to be superior over the long term even if they are not immediately cost effective.

# 4. Technology Research & Development Activities

I-95 is a huge and complex project. Developing dialogue level languages, experimenting with novel user interfaces, prototyping a large repertoire of bridge services, stimulating market activities in several sectors and ensuring that a suitable substrate is in place represent a multi-year activity that will extend well beyond the scope of this project. Accordingly, we have chosen to focus on the critical elements required to make the project viable, and on some unique technologies at all three levels of the architecture where we do not see other commercial or research activities filling the gaps. These activities are described next.

### Activity 1: Build languages for dialogue-level scripts, libraries of common business practices, and demonstrate critical dialogue acts and services

The central motivation for the I-95 architecture is to create modules that support the mix-and-match implementation of information market activities, captured as a collection of acts arranged according to a script. We will explore several possible languages and will decide whether there will be a single scripting language, or a variety of languages open to continuing evolution.

We will assemble a library of processes depicting common business practices to serve as the basis for implementing scripts. We will also make the library itself available on-line as a resource for I-95 participants. Our efforts in assembling this library draw upon work underway at the MIT Sloan School of Management. The library includes examples of how different organizations (or networks of organizations) perform similar processes. It also catalogues the relative advantages and disadvantages of the alternative processes, and includes references to software that supports these processes. As an on-line service the library will help companies (a) redesign their existing organizational processes, (b) invent new organizational processes that take advantage of I-95 technology, and (c) automatically generate software to support organizational processes.

A key element of this work is a novel approach for representing business processes. This approach draws upon ideas from computer science and coordination theory. Its primary advantage is that it can explicitly represent similarities and differences among related processes, so that one can easily find or generate sensible alternatives to a given business process. The motivation for using these approaches stems from the complex nature of economic activities which typically need to be implemented with many variations. For instance markets differ in the characteristics of the "auction" through which buyers and sellers select each other. In a familiar "English" auction, bids are public and increase until the highest bidder buys the item. In a "Dutch" auction, on the other hand, a "clock" ticks down prices until someone indicates a willingness to buy. In a "second price" auction, sealed bids are submitted and the highest bidder gets the item, but pays only the amount of the second highest bid. To support many different kinds of markets, our negotiation software will be able to flexibly combine many of these different parameters: standardized vs. customized products, sealed vs. public bids, rolling bids vs. fixed time, and so forth. An intriguing question for us is whether the availability of new technologies will make some of these market forms which are rare today, more common.

### Activity 2: Construct a coherent Personal Information Environment

One principle of the I-95 architecture is that the dialogue-level information transactions and flows are decoupled from the user interfaces through which they are accessed. This helps assure that the services we develop will scale across a wide range of interconnect bandwidth and input-output modalities. For example, the act of browsing a database should be possible from a low-speed ASCII terminal, from a voice-only mobile telephone, or from a "vision-only" gesture interface. We will construct an advanced mobile user interface that we call the personal information environment.(P.I.E) in order to test the commitment to interface independence, to stress the wireless component of the I-95 connectivity substrate, and to stimulate new applications on the I-95 Information Infrastructure.

Imagine this P.I.E. in the following somewhat futuristic terms:

> Assume that a central P.I.E. component, the one you wear all the time, is a cellular computer on your belt. You pick up a Nikon camera. The device in your belt buckle "handshakes" with it. The camera emits an identifier ("I'm a Nikon model X, serial number Y"). If the belt hasn't met this camera yet, it requests a dump of a command set. Now your personal machine knows how to instruct the camera, and can therefore match it to you: a voice command ("aperture priority mode on") can be accurately matched because the lexicon of all commands within earshot is tiny. The command is transmitted to the camera. You take a picture: the camera sends a blip of data ("picture #17, f3.5, .250 sec, fill flash") that is caught by the belt, time-stamped, tagged with latitude/longitude/altitude and compass orientation that comes from the weather station in your watch, and logged for later use. The picture data ricochets off your belt and, through a cellular link, into the net.

We will develop the architecture and prototype some of the devices to realize the above scenario, starting from more realistic alternatives. The applications we envision build upon a variety of consumer and household appliances. In particular, we'll focus on four fundamental components:

- BodyNet—the local network that connects the devices you wear. This is similar to a wireless Ethernet, operates at extremely low power, is robust and secure within the radius of your person, and is capable of transmitting high-bandwidth data of sounds and images.

- Local Com cells—These are the communication/computing module built into the devices that communicate through the BodyNet. They will be designed to be fully scalable These elements need a communication component, and a processor to map the outside language to internal device-specific controls.

- BodyTalk—like MIDI for synthesizers, the language spoken by the local com cells in these devices must be simple enough to use in situations like diagnosing a furnace, but scalable to permit interaction with complex machines, like automobiles.

- Magic glasses—The P.I.E. user needs one display channel, one audio channel, and a small device or two (like a button/pointer built into a ring). "Magic Glasses" are not a head mounted goggle-and-helmet display in the Air Force tradition, but rather, a high-quality color display built into glasses that are comfortable and even stylish. Sound i/o will be integrated in the "temple" of the glasses, and should be no more obtrusive than a pair of earrings. Existing industrial alternatives (including "eyephone" displays as well as new component technologies, like deformable mirror devices) will be evaluated.

**Activity 3: Implement authoring tools for scalable customizable interfaces**
Consistent with the I-95 interface isolation principle, this activity will provide authoring tools for the construction of multiple-modality user interfaces. As an example of work in this area, we will provide application designers with tools with which they can specify dialogue elements at an abstract level. Given this specification, the user's personalized interface environment will interpret the dialogue elements in accordance with the user's current preferences. This approach has been successfully demonstrated at CMU for small-scale interfaces.

It is well known that a single interface will not allow everyone to work efficiently. That is why most of today's software provides customization mechanisms. For example, many Macintosh applications provide menu-based customization mechanisms, and advanced applications in areas such as CAD and Scientific Visualization provide embedded programming languages for customizations. We will follow the same personalization mindset in this work, striving to make available a rich repertoire of different-style interfaces to suit different taste—all constructible and accessible with relative user ease.

**Activity 4: Human language-based interfaces**

This activity is rooted in the technologies developed for the MIT/LCS Pegasus system. These include: (1) a speech recognition technology that converts continuous speech utterances by any speaker, without training, to a wide range of possible phoneme combinations that could conceivably correspond to the utterance (extensive auditory, phonemic and other experiential knowledge is used to make this mapping); (2) a dictionary of accepted words, along with their variant pronunciations; (3) a natural language back end that performs additional filtering based on natural language parsing and related rules; and (4) a domain-specific context filter that further restricts the acceptable utterances based on the domain of discourse.

The human language-based interfaces of I-95 will go a long way toward making the infrastructure accessible by people who are not computer experts. Starting from the rich and proven base of a fully functioning airline reservation spoken language system, we plan to generalize its structure so that it may apply to a wide range of different domains. This will be done within the dialogue level, striving to match the applicable dialogue level scripts to the corresponding human language level. We will then explore the customization of this generic spoken language structure for use in specific market sectors, such as travel, labor, software publishing and health care. In addition, we intend to explore multimodal aspects of human language interaction, i.e. the use of speech, typing, and mousing to convey the same concept—an approach that results in considerably more effective communication than any one of these modalities taken by itself.

**Activity 5: Design of essential mesh conventions and key bridge services.**

At the core of the bridge level is the Information Mesh, a set of conventions for object naming and typing. This minimalist kernel focuses on those key capabilities which permit universal access in a system which is capable of scaling to global proportions and extent in time.

Based on existing MIT/LCS work on the Information Mesh, we will develop a specification of mesh names to be used in I-95, and we will implement a prototype service for name resolution — an object location service. While the mesh architecture explicitly permits a number of location services to coexist, we will prototype one or more such services for use in our concept demonstrations. Presently, we envision a location service that runs locally on a client machine, and another more robust service for public objects.

We will also develop an approach for object definition, which permits the type, or role of objects to be specified in a standard way. We will provide a prototype object catalog, where object definitions can be stored and retrieved. Using this catalog, one can select an object type, and request an object server to instantiate an object of that type. The object will be created with the proper interfaces, and the user or client can automatically activate the correct tools to interact with the object using the correct access method. For our demonstrations, the preferred access method will be via a dialogue script.

In order to complete a prototype demonstration of mesh capabilities, we will develop a prototype server which can store objects and provides an execution environment for bridge services and other active elements. This first prototype will not stress operational features such as availability or scalability, but it will provide enough performance to support experiments by friendly users.

Linked by the mesh are a set of bridge services, which realize the key building blocks of the dialogue level scripts. We will identity and implement a set of such services, in order to validate the script concept, permit a range of demonstrations, and encourage the development of other operational service elements and of higher level applications.

In activity 6, we identify the security related services of user authentication, digital signature and electronic money. Here we provide some additional examples that illustrate what a bridge service is, and how it fits into the overall architecture. Identification of a more complete list of required bridge services is an early objective of the research.

- Mesh navigation tool: This tool, which relies on the base mesh services of naming and typing, helps a user navigate the mesh, examining and visualizing mesh objects. It can be used for demonstrations, debugging and simple services.

- Simple indexing tool: The mesh concept does not constrain how objects are indexed, except to define the low level names which any index should return. Complex data structures can be assembled and used to search for objects. However, a simple service, perhaps one based on the matching of text strings or simple attribute tags, should be sufficient to support the early development of applications that index and organize mesh objects.

**Activity 6: Implement critical security and authentication bridge-level services**
I-95 should provide powerful and trusted mechanisms for controlling information access, for mixing public and private information, and for authenticating messages. Some of these capabilities will be based largely on the known Kerberos authentication system, public-key cryptosystems and digital signature technology. But I-95 requires more flexible and sophisticated security mechanisms, discussed next.

The benefits of an information market will be substantially greater when money has its own bridge level representation. We call this future token electronic money and envision it as having some of the properties of real money, namely, that it cannot be easily duplicated, that it can be verified as being genuine and that when it passes from one party to another it is not possessed by both, as is normal information. Privacy of electronic money transactions will be very important, and we shall prototype methods to assure that electronic money used in payment for day-to-day transactions is as difficult to trace as cash.

We also plan to experiment with uses of Fair Cryptosystems. This technology permits a participant to choose a secret value and commit to that value by sending a special message to each member of a group of trustees. These messages are special in two ways. First, no minority of trustees can compute the value in question from the information in their own hands. Second, once the right time comes, the information in the hands of the majority of the trustees must be sufficient to reconstruct the value, without the further cooperation of the participant.

One application of fair cryptosystems is to digital signatures with time constraints. Ordinary digital signatures enable a buyer to commit to a given purchase. However, they are not sufficient in more complex situations, like an electronic secret bid. In such cases it is necessary for a prospective buyer to commit to a given price, but be able to keep this value secret until a predetermined time. At that time, all committed values must be revealed, so that the highest bidder can be distinguished. Having all participants digitally sign their bids, and relying on them to later decrypt their own values is not a satisfactory solution—all committed values must become public at the pre-determined time, independent of the availability or the wishes of the participants.

Our proposed security work also contains a more systems oriented component, which addresses not secure transmission of objects but protection of systems attached to the I-95 substrate. Today we see concerns for system security driving institutions to isolate systems behind mail relays. These heighten the assurance that attacks at the packet level cannot be launched against unprotected systems, but restrict the available network services to electronic mail. Relays will prevent the deployment of the information marketplace. Our solution to this problem is a device we call a security firewall. This is a physical device in the network, through which traffic is routed when passing in or out of a security domain. Once in place, the firewall examines the network traffic at a high level, and assures that the communication is well formed. The key to this approach is the dialogue level, which provides an application independent set of conventions and dialogue elements to which all communication conforms. The firewall can insure that only well-formed dialogue elements are forwarded, and of those only the ones that are expected at this point in the script being performed.

A service can thus be assured, without any explicit security assurance on the local host, that the only packets which reach it contain actual, well-formed dialogue utterances from the expected script. Low level attacks against the network transport services or the operating system are thus prevented except in the special case that they masquerade as properly formed messages, which is a very powerful restriction on the range of attacks. The local service is thus left with only the problem of assuring that the application itself is constructed properly. We will prototype a security firewall. We will select some programming environment that is as secure as possible, and build our system on this environment. The firewall will obtain the definition of the valid script for a service, and use this to configure itself to protect that service. We will put this firewall in place, and invite attacks to validate and evolve the concept.

### Activity 7: Implement bridge level storage and access services

The amount of information that will grow on the I-95 infrastructure is likely to be of the order of millions to billions of information objects. In such a rich setting, finding the information one needs, being able to store information reliably and securely and operating efficiently with stored information become key requirements. Here are some of the issues we will address:

*Resource discovery—finding information*

Within I-95, we will prototype bridge services that allows users to browse, locate, and search for different types of multimedia information such as text, digitized video and audio. Our demonstration of these services will include a dialogue level browser that helps the user to formulate queries, discover resources, and retrieve relevant information. This browser will make use of three bridge-level components—a Semantic File System, a Content Routing System and a Structured Video System.

The Semantic File System is an information storage system that provides associative access to the system's contents, automatically extracting attributes from files by using type-specific transducers. Semantic File Systems may be implemented as conservative extensions to existing tree-structured file system protocols or by protocols that are designed specifically for content-based access. The former approach has allowed semantic file systems to be seamlessly integrated into existing file system tools like directory editors or regular expression search commands. A gateway to existing WAIS servers provides access to an array of information providers.

The Content Routing System provides query-based associative access to distributed information systems. Queries are automatically forwarded to servers that contain relevant information. Since there will be so many objects on I-95, under-constrained queries can produce large result sets and extraordinary processing costs. To deal with this scaling problem we use content labels to permit users to learn about available resources and to quickly formulate queries with adequate discriminatory power. The content routing system is organized as a network of servers that present a single query-based image of their distributed information systems.

The Structured Video System provide a framework for indexing, editing, storing and accessing digital video on existing workstations and networks. It organizes the video data in a nested logical structure of video segments. It provides the structure information (shots, scenes, sequences), the temporal ordering of segments, descriptive information (description of segments; close captioning text), and presentation information for video data. This rich representation defines an adequate data structure for associative access to video segments based on attributes extracted by the Semantic File System.

The browser, which is the dialogue level peer of the end user, issues queries to the Content Routing System which forwards them to the relevant Semantic File Systems. Browsing through content labels allows discovery of interesting information providers. When the search space has been sufficiently narrowed, the user can formulate an object search query. Once objects of interest are found, they are presented to the user. The Structured Video System handles the details associated with the playback of discovered video objects.

*Video capture and storage*
We plan to make a large scale video server available to all I-95 providers and consumers of information. The server will be available to store full motion video files and other types of files, accessible by hundreds of clients. In addition, Digital has developed a low cost audio/video module for PC-compatible platforms. Using these components, we will investigate the scaling of servers to the support of large user populations and the substrate demands arising from pervasive video traffic.

*Persistent storage—reliably and securely*
The bridge level must support secure sharing of its objects; must provide mechanisms for users to find objects of interest; must ensure that objects persist and are available when needed; and must provide good performance.

Secure sharing requires objects whose behavior can be relied on, and mechanisms that allow users to control who can use objects. The first will be provided by having strongly-typed, abstract and encapsulated objects within an extensible type system (i.e. objects are accessible only by calling their operations). Control of sharing will be linked to the operations of objects, allowing a finer grain of control than can be achieved by read/write sharing.

Since the mesh universe will be very large, mechanisms that make it easy to find objects of interest are crucial if the system is to be usable. There will be a low-level naming mechanism that allows objects to be accessed quickly, but these names are not intended for people to use. People need to find objects using very high level names related to the task they are doing. Two approaches to providing high-level names will be supported. The first is a query mechanism that supports efficient searching of large collections containing objects of varying, abstract type; queries can be based on abstract properties of objects (e.g., the occurrence of a pattern in a video image). The second is a 'whiteboard' approach in which users post requests for services and suppliers post advertisements.

Performance will be achieved by using object caching and prefetching. Caching permits objects that are likely to be used soon to be retained in a location close to the user (e.g., at the user's workstation, or at a cable company's head end). Prefetching allows objects to be fetched to the cache in advance of their use. The two together allow most user accesses to be performed locally.

*Storing and retrieving information efficiently*
I-95 servers will store large amounts of information and should be able to retrieve and store data on disk efficiently. The performance of processors, however, is increasing faster than that of disks. Therefore most servers and applications are likely to become I/O bound. Fortunately, the effect of the I/O bottleneck can be reduced by making more effective use of the full disk transfer rate. In most file systems the effective transfer rate for small files is only a small percentage of the raw transfer rate, as much of the time spent in accessing disks is used to perform seeks.

We will pursue a new approach to solving the I/O bottleneck that we call Logical Disk (LD). LD provides an abstract interface to the disk that is based on logical block-numbers and lists of blocks. The LD interface is powerful enough to simplify the efficient implementation of many file systems, yet abstract enough to allow for different implementations of the interface. LD defines an interface in which the file system is responsible for file management and LD is responsible for disk management.

Separating these two concerns has three advantages. First, it leads to a file system structure that makes files systems easier to develop, maintain, and modify. Second, it makes file systems more flexible. Different file systems can use the same LD without having to rediscover the most efficient disk layout scheme. Similarly, different implementations of LD tailored for different access patterns can coexist, allowing each file system to choose the best one for its needs or even use several implementations. Finally, it can improve overall file system performance, as demonstrated by MIT's log-structured LD prototype.

**Activity 8: Stimulate development of the I-95 connectivity substrate**

The I-95 substrate will be a real network, providing access to all subscribers within reach of the participating service providers. Initially, we will focus on the provision of Ethernet-like access. However, as the project progresses we will introduce new services and technologies. The following paragraphs outline specific substrate activities.

*Local Access and Open Interconnection*

- I-95 by 1995: The I-95 partners will work towards the realistic goal of making Ethernet-like access ubiquitous throughout the region by 1995. A key step towards this end has already been realized — during the formation of the I-95 alliance Continental has committed itself to the aggressive deployment of the requisite access technology. Although local access will be offered on a commercial basis, we will jointly plan, test, and operate key substrate components. In particular, Digital has committed the resources of its Networks Systems Laboratory to this project. Finally, the I-95 market development teams will actively support substrate users and will work with market-specific trade groups to facilitate substrate subscription.

- Beta Substrate: Continental, Digital, MIT, Nynex, and the Lincoln Laboratory will develop a beta test substrate to be built using cable-TV and telephone physical facilities, multi-protocol routers, and various access devices. The substrate will consist of a backbone of mesh-connected multi-protocol routers that identify externally visible Points of Interconnect (POI). Each POI will house the devices needed to provide network access to the subscribers it services. These access devices will be multi-protocol routers that are optimized for cost-of-access. Three types of local access will be provided: 6 Mbps cable, 1.5 Mbps T1, and 56/64 Kbps. In every case there will be a demarcation router placed at the client premises. This device will be the formal demarcation between the substrate and its clients. Approximately 100 client access drops will be constructed. The majority of these will be cable access; some will be 56 Kbps telephone links; and a dozen or so will be T1 carriers.

- Ethernet Cable Access: Digital's second-generation Ethernet CATV modem will support the I-95 *basic* interface. Digital has considerable experience in the construction of access networks based on Ethernet CATV technology and will bring this expertise to bear on the design and construction of the I-95 substrate.

- Backbone Facilities: Continental will make its regional fiber optic network, which passes most of the cable head ends in the region, available to I-95. At present, we envision a Sonet/ATM facility operating at 155 Mbps. In addition, Raytheon will provide access to its private fiber optic network, which covers over 100 miles of Eastern Massachusetts.

- Network Operation and Management: The I-95 substrate will be supported by a distributed control center operated by Continental, Digital, NYNEX, Raytheon, and Teleport Boston.

### *Advanced Technologies*

ATM Interconnect: Raytheon and Digital are developing ATM network components, including adapters and switches that will be used within the I-95 backbone. Raytheon will analyze I-95's traffic requirements and expressly configure its nodes for backbone operation. They will leverage expertise developed in their military programs to ensure the robustness and security of these components. These Raytheon nodes, which will be based on technology developed for the GPALS and JPST programs, will be inserted in the beta test substrate. Interoperability between Digital's and Raytheon's systems will be demonstrated.

- Advanced Cable Access: In the early phases of I-95 the demarcation routers will be constructed by bundling existing commercial components. Digital's Network Systems Laboratory (NSL) will prototype a custom demarcation router designed to be maintainable, and trouble-free.

- Wireless Transfer Modes: Raytheon will develop a system that supports a 20-100 Mbps shared access channel and allows individual users to burst at rates up to 6 Mbps. This rate is roughly equivalent to the burst capacity of the initial cable-based offering. The system will be based on an intelligent hub that arbitrates channel access and provides a gateway to the wired network. Digital, MIT and the Lincoln Laboratory are also working on third generation wireless networking technology, which will be tested on the I-95 substrate. U.C. Berkeley is prototyping a portable wireless pad which supports video, pen input, audio and text and graphics. It is basically the I/O port into the backbone networks and is designed to be very inexpensive, light weight, have low power consumption and be the size of a standard notepad. They will provide these to the project and work with the other proponents on their integration. Finally, the Lincoln Laboratory will provide wide area connectivity via its LES9 satellite.

### *Architectural Refinement*

Enriching the Connectivity Interface: MIT has developed a model for an Integrated Services Packet Network, that supports controlled qualities of service. We will work with the substrate service providers to support these new services, and to explore the range of applications that are enabled. Our goal is both to validate our model of the service interface, and to promote the commercialization of multi-service packet networks.

Renewable Substrates: Software-based implementations, especially of processes previously supportable only in hardware, may be an important tool in planning for substrate renewal. One area of particular interest involves the complex channel coding, or modulation, techniques being developed for digital television transmission. Steady progress is being made in channel coding efficiency, i.e., the number of bits/hertz achievable in a particular environment. We will experiment with software-based channel coding schemes that will allow carriers to rapidly upgrade plants to take advantage of new coding techniques. Technology originally developed for remote sensing can be used to sample entire frequency bands for digital analysis and/or transmission. Although a software-based scheme will be computationally expensive and impractical today, the rate of processor improvement suggests that the time is ripe for a proof of concept experiment, even if it is restricted to ultra low bit rates by present day processing platforms. On a broader front, we will develop plans for substrate evolution and, working with the carriers, we will publish architecture and engineering guidelines for the design of renewable substrates.

# 5. I-95 Market Development Activities

I-95 will accelerate the transition towards an information based market economy. Although the infrastructure and services to be developed within I-95 are intended to be quite general, it is important to focus on particular markets so that we can achieve *critical mass* within specific communities of interest. This activity will be driven by the market development units, each of which will be an integrated team that focuses on a specific market or group of overlapping markets. I-95 users will be encouraged to participate in a number of these markets, thereby stressing the generality of our user environments and dialogue level abstractions.

The market development teams will work with each user community to identify network-based business processes that are tailored to their market sector. During the course of the project the market development teams will steer their users through a progression of increasingly sophisticated services. Initially, the teams will assemble market-specific tool kits based on existing products. Working closely with the technology development teams they will upgrade these market-specific kits, taking increasing advantage of the models, services and interfaces under development within the project.

The sections below describe the markets upon which we intend to focus our initial efforts. During the first months of the project, we will work closely with user communities in the region to refine this selection of markets.

## 5.1 Software Products and Services

There are over 1600 software companies in Massachusetts, and the software industry seems particularly suited to electronic interactions. Participants in this industry include: software companies, software professionals, software customers, and software industry analysts. I-95 will help buyers and suppliers of software find each other, negotiate deals, and place orders electronically. An obvious (but not necessarily trivial) step is direct fulfillment, i.e., the electronic distribution of the products themselves—software programs, updates and documentation. Sophisticated support for the discovery of suppliers might even enable a market for infrastructure components, such as bridge level services. The Massachusetts Software Council has agreed to work with us on the development of this market, and Lotus, the largest software developer in the area, will explore the adaptation of the I-95 architecture to specific software products. During the course of the project, the scope of this effort may be extended to the broader publishing industry, a market sector that NYNEX is already studying.

## 5.2 Health Care Products and Services

The timing of health care reform is particularly fortuitous, since the I-95 infrastructure will come on-line just in time to support the new applications and services that will be required. NYNEX and Digital will investigate Community Health Information Networks that link providers, employers, payers and patients. Services under consideration include: clinical data access, medical records, claims processing and supplies ordering.

Home health care is an area in which both Digital and IBM are presently experimenting. For example, IBM and the New England Medical Center (NEMC) are investigating the use of full-motion video and other multimedia techniques for patients requiring extended care at home. The initial study focuses on children with leukemia, who have many needs that are met by family members at home. These include patient monitoring and active participation in patient care. Audio and video significantly enhance the quality of information in such areas as home care training, symptom evaluation and emotional support.

I-95 offers significant financial benefits over existing alternatives. At present, each family is loaned an $8000 system containing 2 hours of video data, which is approximately 15% of the available information. With I-95, we can loan patients a $1500 system that uses the substrate to access the answers to all likely questions, resulting in greater family and provider satisfaction. The price decrease also makes this application useful for almost all chronic medical conditions, not only life-threatening and expensive conditions such as leukemia. We believe that this approach will eventually be funded by medical insurance, since health outcomes are improved and hospital costs are avoided. This application will be used by many non-computer-users, operating at times when they are tired, under stress, and distracted. This results in special requirements on the user interface and on the underlying transport mechanism.

## 5.3 Small Business

I-95's reach must extend to small businesses, some of which have been slow to modernize operations either because they are not high-tech, or because their commerce is in physical deliverables. They are the support structure for the large corporations that encircle them. Examples include the nearby printer, the office supply store, the messenger service, the garage, the florist and the local restaurants. The small business team, led by the MIT Media Lab, Lotus, and Raytheon, will search out specific markets, where a limited amount of support can facilitate the transition to an information based economy.

Small businesses epitomize situations where the benefits of being on-line diffuse throughout the community, yet the costs of interconnection are acute. The company directly pays the costs of terminals and network access but the whole environment gains the benefit of more fluid operations. For them the connection between an investment in communications and the bottom line is vague. However, these businesses are in no small measure responsible for the definition of the business climate of the region. Implicit is a requirement that the infrastructure contains access mechanisms that allow cheap entry. The small business development unit will define a minimal "infoport" made up entirely of hardware and software components that can be acquired at local retail outlets such as CompUSA. The customer's equipment should cost no more than $5000 and communication costs should be competitive with those of existing appliances such as fax and telephone.

## 5.4 Federal Business

No business community is larger than the combined agencies of the Federal Government and its broad range of suppliers and "clients". The government's growing need for cost containment and accountability has led to increased paper-based reporting. The I-95 approach can greatly reduce the energy, cost, and time consumed by the current business processes of this community. Raytheon, which will leverage expertise developed in the Government/Industry sponsored CALS (Computer-aided Acquisition and Logistic Support) Program, will adapt the I-95 architecture to such tasks as compound format messaging, information search and retrieval, electronic document receipt/acknowledgment, annotation and acceptance.

I-95 will provide the common infrastructure now missing from the CALS model. Instead of each contractor developing its own access and delivery strategies and implementations, all users will exploit a rich set of generically available services. This approach will foster a third-party market for low-cost turnkey systems, allowing smaller enterprises to fully participate in the federal sectors of the information marketplace. Service can be extended into many areas of federal involvement including: environmental data bases and reporting; employment services such as job banks, job description and resumé databases; smart postal services or alternatives; and other opportunities related to the GSA, Energy, Transportation, Law Enforcement and Agriculture.

## 5.5 Mass Markets

In addition to specific applications for the markets identified above, we will also make several capabilities available to all users of the network, including the general public. Because of their wide appeal, these mass markets will stress the generality of our approach, e.g., testing the hypothesis that user skills and interface environments developed within a specific market can be applied to adjacent and/or overlapping market sectors.

### 5.5.1 Labor Markets

As the current corporate restructurings suggest, efficiently matching workers to a continually changing array of jobs is likely to be a significant social and economic problem for some time to come. As information technology is applied to the labor markets, regions and nations with I-95 like infrastructure will be more attractive for both companies and employees.

At a minimum, this market will require shared databases of available jobs and resumés of people available for jobs. More elaborate services can include automatic screening tools (e.g., to find plausible candidates for a job), on-line reference checking (or pre-stored recommendations from previous clients), video-interviewing, and so forth. This activity should co-opt existing service providers, such as head-hunters, who might collaborate on the establishment of a multiple listing service (MLS) patterned on that of the real estate community.

One possibility that overlaps nicely with the software market unit is to focus on the labor market for software professionals. An extreme example of what these capabilities could support is the formation of "overnight companies"—rapidly assembled teams of people who perform short projects and then disband. Such efforts are already common in some industries (e.g., film-making and construction), and I-95's network capabilities make them feasible in other situations.

### 5.5.2 Want ads

As preliminary services of this type on the Internet already indicate, there is a latent demand for market-making services of the type currently provided by newspaper want-ads: apartments for rent, used furniture and appliances, real estate, etc. With little customization to particular types of markets, general tools such as our dialogue level browser, can replace a laborious manual search with a simple database query.

### 5.5.3 Education and Training

NYNEX will experiment with I-95 delivery of educational services. Specific opportunities to be investigated include:

- Remote Classrooms: Using video conferencing and shared whiteboard technology, teachers can lead classes attended by students at geographically dispersed sites.

- Lesson Planners: Teachers will be able to consult content experts who assist in the location of educational materials that can readily be incorporated within lesson plans.

- Portfolio Assessment: This tool will support the collaborative review of student portfolios.

# 6. Instrumentation And Evaluation Activities

## 6.1 Substrate Measurement And Analysis

UCLA will develop a center for I-95 substrate measurement, evaluation and design. The center will serve the I-95 community by tracking traffic patterns and trends. It will evaluate the overall performance of the substrate and provide directions for growth.

The I-95 substrate will be an interconnection of networks based on different technologies (wireless, wired backbone, LANs, satellite channels, etc.), covering different geographies, and serving diverse needs of different communities. An important responsibility of the UCLA center will be to discover correlations between traffic patterns of users within the same market as well as across markets. The center will be equipped with tools to: monitor network and user behavior; collect measurements; derive correlations; predict performance; and plan growth and reconfiguration. In addition, the center will make measurement data and tools available to other groups (users or network infrastructure providers) who may wish to develop their own performance projections and reconfiguration plans.

The center will be a shared facility that allows a diverse community to interact, to share ideas, and to document their findings in a "live and active" repository of data, tools, system models, and reported results. The work of earlier developers will be captured and re-used so that. the effort that goes into its formation is not lost when a project ends. The UCLA center will be a national resource and we hope to work with other regional alliances on its establishment and operation.

## 6.2 Productivity, Management and Social Impact

Much of the value created in a modern economy results from the economies of specialization and scale enabled by efficient markets. The value that I-95 can create by improving search, reducing transaction costs and enabling new coordination mechanism is likely to be very substantial. We will evaluate our productivity improvements through traditional indirect measures, such as output and productivity, profits and sales, and consumer surplus. However, since I-95 will create new markets and radically redefine many existing markets, some of the most profound changes it enables may escape the traditional measures. This is especially true of changes in the nature of products and services and in organizational structure.

The changes enabled by an information infrastructure include improved "fit" of products and services to customer needs, increased timeliness, reduced error rates, reduced inventories, greater product variety and increased innovation. For instance, inventory levels and the frequency of "stock-outs" are likely to be reduced, and, with the cooperation of participating firms, these benefits can be directly estimated.

The full benefits of I-95 will only be achieved if people adjust and firms re-configure themselves to adapt to the new environment. For instance, firms increasingly rely on market coordination (e.g. outsourcing) as information technology reduces the cost of using the market. I-95 will enable flexible approaches to production (e.g. "agile manufacturing") and the rapid coordination of teams of experts to meet particular needs (e.g. "virtual corporations"). Careful measurement of the organizational changes associated with the introduction of I-95 will help firms prepare for the changing competitive environment.