A FREE-MARKET EXCHANGE FOR INFORMATION

R. J. Lipton
J. S. Sandberg

CS-TR-298-90

December 1990

A Free-Market Exchange for Information

R.J. Lipton
J.S. Sandberg
Princeton University

Abstract

All members of the high-technology community possess a wide variety of valuable information. Such information ranges from operation procedures for Computer Aided Manufacturing hardware to a 100 line spreadsheet program. Some world class high-technology experts (e.g. Nobel Prize winners or Industry leaders) possess unique information or insights that cannot be independently reproduced. Other members of the high-technology community, familiar with specific subsystems of profitable or critical products, possess information that can only be reproduced at significant expense. Examples of such information include: specific development information about a software module used in a large defense system, the curvature of a mirror deployed in a satellite, or the location in the company database of last month's sales forecast. The value of the information possessed by employees of profitable corporations, researchers in leading laboratories, or administrators in government offices is hard to overestimate. "You can take my factories," boasted Henry Ford, "burn up my buildings, but give me my people and I'll build the businesses back again."

People know the approximate value of a given piece of information. They generally have some informal mechanism for trading (and sometimes selling) the information, although frequently, there is no formally established marketplace for them to sell much of the valuable information they possess. We propose to establish an electronic marketplace for a broad class of high-technology information including such topics as: computer software and hardware, electronic components, high-technology consumer products, systems integration, and data communications systems. More specifically we define and discuss a free market-based Information Exchange for software systems development and maintenance that promises to increase the flow of valuable information through laboratory, company, and international computer networks.

Establishing a worldwide electronic network for purchasing technical information will accelerate the pace of innovation and the degree of adaptation of computing, communications, and information technologies in analysis, design, and manufacturing processes. Perhaps most importantly, a international information market will enhance and motivate the human resource base to efficiently meet the economic and technological needs of the new decade.

1. Introduction

The collective expertise of a corporation's employees is an asset that typically is left off annual balance sheets; yet this informal distributed database is potentially more valuable than any capital equipment or accounts receivable. Employee training, research and development, and market analysis are all corporate activities that introduce and distribute information throughout an organization. Anderson Consulting expects to spend nearly $250 million on employee training this coming year. Corporations typically reinvest up to 10% of their earnings in research and development. Marketing costs are nearly always greater than research and development costs. DEC estimates that the Electronic Document Interchange market, only a small portion of the total information management market, will exceed $2 billion annually by 1991. With such a massive investment in obtaining information it is not surprising that corporate management is constantly hunting for better ways to distribute data and information.

Information flows within a company according to the predictive or reactive capabilities of its employees. If someone identifies information that is widely applicable throughout the company they might predict that writing that information down will be a significant advantage for their colleagues. Books, manuals, and accounting systems are examples of predictive information supplied by colleagues both inside and outside of the company. Predictive information tends to be fairly general and there is a lot of it. New operations procedures, prototype systems design specifications, and current sales & inventory figures tend to be information acquired through employees reacting to specific queries. This reactive information is generally too new and volatile to be completely recorded. Frequently, the people that possess the expertise necessary to competently record the desired reactive information do not, because they "have too many other things to do." This highlights not only that reactive information is rarely recorded, but also that it is difficult to obtain.

As both companies, and the data relevant to their operation, grow more of their information transfer must be through predictive transfer mechanisms. Distribution of responsibilities, management hierarchies, and geographical separation all contribute to the disruption of paths for reactive information flow. If the only person that knows the answer to my question is 3000 miles away it is more difficult to get the required answer than if that same person is in the next office. The more an organization relies on predictive information channels the less specific the information flowing through the organization. It is one of the responsibilities of management to augment the corporate information flow, to and from their organizations, by becoming a source of reactive information.

This information management problem is important in purely information processing contexts such as: computers, networks, software, and artificial intelligence, however, it is equally relevant in all high-technology and systems administration areas. On the McDonnell Douglas Apache Helicopter line a mechanical assembler asks engineering if the buttonhead rivet right next to the bracket, called for in the blueprints, might be better replaced by a flush-head rivet thus preventing chafing and a possible catastrophic system failure. At the

BMW Regensburg assembly plant a systems integration consultant in charge of the network operating system upgrade installation needs to know if each of the 998 plant microcomputers have a spare 16-bit AT slot for a new fiber optic LAN adapter and is the hardware interrupt IRQ14 available on each machine. The cardiology fellow at Cedar Sinai Hospital needs to know how to enlarge an image (exactly what buttons need to be pushed and in what order?) on Toshiba's cardiovascular angiographic system. The key to efficient information processing is to ensure rapid high volume flow of predictive and reactive information throughout the organization as demonstrated in the following examples:

> NASA Hubble Telescope - is a recent example of a spectacular systems administration failure due to retarded reactive information flow within a single project. A $1.5 billion systems development project failed because the critical viewing apparatus was not manufactured to specifications and nobody performed even basic suitability tests.
>
> HPC Grand Challenges - are a set of initiatives, proposed for federal sponsorship, to solve a set of outstanding basic scientific questions including: Human Genome, Superconductivity, and Turbulence. In recent administrative planning sessions for the Human Genome Project senior researchers identified information technology resulting in collaboration between scientists as one of the core problems requiring resolution.
>
> DARPA Megaprogramming - is a component-based software engineering and life-cycle management program run by DARPA/ISTO. They are looking for high capability systems through high capability components in the areas of real-time, trusted, concurrent, heterogeneous, and high performance systems. DARPA considers that the areas of greatest leverage in a successful Megaprogramming effort are exploiting existing infrastructure, developing information repositories, and buying information.

One basic problem that has yet to be successfully addressed in all previously proposed models of systems administration and development is how to increase the exchange, within an organization, of reactive information in the form of: answers to specific questions, software modules, system documentation, or hardware system specifications? For many of the types of data we commonly buy and sell there are individuals or corporations that act as "market makers." Their role is not only to facilitate the exchange of data for money, but, to cause data to be priced fairly. Today, no such market mechanism exists for any but the largest software and systems development projects. I can purchase an AT&T telephone switching system or a Microsoft C compiler, however, I cannot buy a short 2-page description of a software module written by my colleague 18 months ago. Predictive information may be purchased, however, reactive information is rarely for sale. Obtaining significantly faster, smarter, more flexible, highly interconnected information technology depends on being able to readily purchase both predictive and reactive information.

This paper is a proposal to develop a formal free-market structure for data that will increase the flow of technical information through organizations. Our central hypothesis, borrowed from Adam Smith, is that owners of information (just as laborers in classical economic theory) will do what is to their best monetary advantage. Given an organization-wide or international electronic medium (a computer network) through which data and information can be purchased, valuable data will become available when and where it is needed. In Section 2 we will define the fundamental transactions executed through the Exchange. Section 3 presents evidence to support our assertion that an information market will encourage a specific group of corporate employees - software engineers - to invent, innovate, expand, and take risks to further their self-interests thereby improving their organization's productivity. The final section summarizes our accessment of the prospects for an free-market based Information Exchange.

2. The Information Exchange

The proposed Exchange is run by a Board of Directors responsible for enforcing the orderly conduct of business through a variety of predefined transactions. The transactions are designed to be efficient, fair, and legally enforcible mechanism to promote the trade of information and data for financial compensation.

The Exchange operates by collecting, organizing, and displaying offers, in the form a standard contract, requesting specific data or information. Collection and display of the offers are facilitated through and organization-wide computer network. For a single transaction a contractor issues an offer, through the Exchange, containing all the data and information required to fulfill the contract including:

(*1*) A price to be paid upon successful execution of contract,
(*2*) A time period for which the offer is valid,
(*3*) Terms - (exclusive vs. nonexclusive),
(*4*) A question, request for data, or system specification,
(*5*) Restrictions
- Posting Limitations - i.e., within group, organization, country, etc.
- Data or Information Ownership Classification
- Testing Procedures, and

(*6*) Posting Directions - (Keywords).

The parties to a transaction, the contractor and the contractee, as well as all potential contractees must each have accounts in good standing with the Exchange. The funds in the account must be sufficient to cover all outstanding contracts, contracts under review, and offers. All information that passes between the contractor and potential contractees - including the initial offer, all notifications and acknowledgements, and any responses - must be processed and recorded by the Exchange. Specifically, all such information must be collected, logged, and distributed through the Exchange's computer network. Any information passed between a contractor and potential contractee that is not registered with the Exchange cannot be referred to - in any way - as part of the transaction. A transaction is completed in one of two ways depending on the exclusivity terms of the contract.

If the contract is exclusive then executing the transaction requires the contractee to notify the contractor that the offer has been accepted. The contractor's role in the notification process is simply to acknowledge (uniquely) that the contractee is accepted or rejected. A contractor can accept no more than one contractee for a single exclusive offer. Furthermore, a contractee can only be rejected because of a) failure to meet the restrictions posted in the offer or b) the offer has already been accepted. Contractees may appeal initial rejections to the Exchange's Board of Directors. When the contractee has completed assembly of the specified data and information he notifies the contractor, sends the data, and receives an acknowledgement of receipt. Finally, the contractor responds to the exchange within a fixed period of time (24 hours) either that the contractee's response has been accepted or rejected. If the contractor accepts the response then the contractor's account will be debited and the contractee's account credited the price established in the contract.

If the contract is nonexclusive then the execution of the contract is simpler because there is only one acknowledgement required from the contractor. When a potential contractee has assembled the information specified in a nonexclusive offer he may accept the offer. The contractor acknowledges one or more of the contractees who then immediately release the requested information. As in the exclusive terms case, the contractor then has a fixed period of time to either accept or reject the contractees' responses. If the contractor fails to accept at least one of the responses then all relevant transaction data must be submitted for review to the Exchange's Board of Directors. If the contractor accepts one or more responses then the contractor's account will be debited and each contractee's account credited the price established in the contract.

Regardless of the exclusivity terms, if the contractor rejects a contractee's response then the original offer and the response are submitted to the Exchange's Board of Directors who have the authority to sustain or overrule the rejection. If the rejection is sustained the contractee's account will be debited a set penalty fee (20% of contract price) which is split and in turn credited to both the exchange and the contractor. If the rejection is overruled the contractor's account will be debited a fixed penalty (10% of contract price) to the exchange in addition to the original contract price. The contractee account will be credited the price of the contract.

Expiration of offers and timeliness of acknowledgements and responses will be determined exclusively by the Board of Directors. The Exchange reserves the right to review offers prior to posting to determine suitability, enforcibility, and legality. However, the Exchange makes no warranties or guarantees on any data or information contained in, or resulting from, offers or transactions.

Just as Adam Smith's Law of Population states that "Laborers are produced on demand," so we hypothesize, owners of information are produced upon demand. Certainly the demand for owners of information is growing explosively. The free-market system we are proposing uses well studied classical methods to promote a high volume intra-organizational data flow that is otherwise unachievable. "Laws of the market show us," writes Heilbroner, "how the drive of individual self-interest in an environment of similarly motivated individuals will result in competition; and they further demonstrate how competition will result in the provision of those goods that society wants, in the quantities that society desires, and at the prices society is prepared to pay."

3. An Information Exchange for Software Engineering

Software development and administration represents a large fraction of the worldwide investment in technology. Despite declining appropriations the U.S. Department of Defense (DoD) will increase spending on all aspects of software from $32 billion in 1990 to $46 billion in 2000. 67% of this Department of Defense software budget, over the remainder of this decade, will be spent on software re-engineering. That means that over two thirds of the software budget will be used to recode existing DoD-developed programs. Further, DoD is only a small, albeit influential and representative, part of the software marketplace. Computing constitutes a significant portion of the U.S. economy. In 1988 the U.S. computing industry accounted for 10% of GNP and almost 10% of all capital investment. The software investment in Japan and the Pacific Basin market economies exceeds $300 billion in 1990. By 1990, prognosticators expect, that there will be over 5 million programmers in Japan. The total worldwide investment in software will easily top $1 trillion this year.

One question confronting all government, industrial, and military agencies responsible for strategic planning and policy development in science and technology is, how to significantly improve software systems development and maintenance? Many software development projects are notorious for being way behind schedule and wildly over budget. Post-development systems costs frequently dominate the total cost of the system's development. The organizations that conduct these development and maintenance operations are processing massive amounts of volatile data and information. Moreover, no single employee or administrator can ever hope to master even a significant fraction of the information relevant to a modest size systems project.

A successful software engineering project relies on the cumulative expertise of an entire organization. All information relevant to the production of a final product must be accumulated and retained somewhere (or by someone) in the organization. In a hypothetical commercial software development project, for example, one team may be in charge of I/O architecture and writing device drivers while another team focusses on marketing strategies for the final product. Since there is no single individual that understands every aspect of a given project, it is frequently the case that important project information can be lost, or worse, never identified.

Federal agencies such as DARPA, NSF, and NASA are joining universities such as MIT and CMU and corporations such as DEC and Xerox to "develop a scientific theory that would explain how the activities of separate players, both individuals and machines, can be coordinated." The object of coordinating man and machine is to discover high leverage technologies that will promote more efficient software engineering. To illustrate the numerous technologies and variety of complex strategies employed in the pursuit of more efficient software engineering technology we quote a recent DARPA presentation on the Megaprogramming effort.

The following list identifies possible topics to be addressed in the development of a DoD Software Technology Plan.

(*a*) Software engineering environment frameworks. These are central to the integration of technology capabilities and the stimulation of commercial software tools responsive to DoD needs.

(*b*) Software engineering tools, including requirements, design, code, instrumentation and analysis, test, configuration management, and post-deployment support tools.

(*c*) Prototyping tools and their underlying prototyping support capabilities to support requirements engineering and design.

(*d*) Capabilities for classification, retrieval, and evaluation of reusable software assets, including code components, interface definitions, test cases, requirements specifications fragments, etc.

(*e*) Domain-specific software architectures, application generators, and domain-specific computational models. Opportunities exist in domains such as automatic target recognition, avionics, navigation, C^3I, and simulation and planning, as well as infrastructure areas such as real-time kernels, image processing, and signal processing.

(*f*) Software re-engineering. Apply or retrofit modern software technology (e.g., decompilers, code analyzers, testing aids, configuration management aids, Ada transition aids) to DoD's huge inventory of antiquated software.

(*g*) Management tools, including metrics and cost estimation. Candidates include group coordination and decision aids, knowledge-based software risk management aids, hypermedia and software visualization technology, gaming aids for training software managers, and automated support of modern software process models.

(*h*) Ultrareliable and secure software.

(*i*) Distributed and parallel software. Applications include large scientific and engineering modeling, embedded real-time applications, and AIS systems.

(*j*) Scalable Artificial Intelligence (AI) capabilities, interoperable knowledge base services, interoperability between AI services and conventional software services, and verification and validation for AI applications.

(*k*) Systems software, including support for security, ultrareliability, and real-time.

(*l*) Computer science base critical to addressing future DoD needs in software reliability, security, parallelism, and distributed real-time.

(*m*) Technology transition support including shadow projects, mid-life re-engineering and prototyping.

As the reader can verify by inspection of the DARPA list of promising software engineering technologies, there are a tremendous number of approaches to software productivity. Some promise very general long-term improvements while others promise quite specific short-term benefits. Any researcher would have difficulty selecting even one of these ideas as being clearly superior to the rest. Will "scalable artificial intelligence capabilities" be more helpful to computer programmers of the early 21st century than "domain specific software architectures?" Who knows?

The goals of software engineering research are pedestrian: we want to write less code, rapidly employ new software, eliminate duplication of programming effort, and efficiently produce working software. Yet the proposed techniques employed to reach these goals are all highly intellectual, requiring multidisciplinary expertise in areas ranging from Kalman filters and Church-Rosser reductions to Managerial Economics and Human Relations. As a group we promote various module composition algebras; establish highly descriptive process control structures called hierarchical, behavioral, semantic nets, threads, object-connection, value-set/entity, and petri-nets; and every decade or so we develop an entirely new generation of computer languages that more completely describe, at a very high level, how we intuitively want systems to behave.

The former U.S. Treasury Secretary, Bill Simon, used to shout at his young traders at Solomon Brothers, "If you guys weren't trading bonds, you'd be driving a truck. Don't try to get intellectual in the marketplace. Just trade." Establishing an Information Exchange allows industry and government to avoid making, almost certainly mistaken, intellectual judgements in software engineering without shrinking from their duty to obtain results. By facilitating a high volume information trading software development success is inevitable. It makes no difference if "hypermedia management tools" or "software engineering environment frameworks" are the key enabling technologies leading to increased productivity. In the final analysis, it is only important that software becomes cheaper to develop and maintain, easier to use, and more reliable.

The idea of utilizing a free-market approach to difficult systems problems is hundreds if not thousands of years old. Applying free-market principles to software engineering is new. Of course you can currently purchase both small and large software packages. You can send electronic mail questions to paid consultants who will return answers. And you can access massive databases, network bulletin boards, and software repositories containing a wealth of technical information. Each of these services are useful tools that can be successfully exploited; recall the recent unprecedented electronic mail collaboration on the IP = PSPACE Theorem reported in the scientific and popular press. But still, these are only isolated parts of a thriving social structure promoting high volume information sharing and man-machine collaboration.

The current software market is a good mechanism for moving general, multi-purpose, mature technology such as: compilers, spreadsheets, databases, operating systems, and communications applications into the hands of developers. All of these long-lived applications are examples of predictive information. In program development the need for a compiler, for example, is easily forecast. As a program developer I don't need a special compiler written to complete my job. Typically, many commercially available compilers are suitable for the programming task at hand. This is not the case for system specific cutting-edge technology required in software engineering development projects. There are no commercial software houses that will sell you reactive information such as: a unified flat file system for an optical jukebox simultaneously compatible with DOS 4.0, Mac OS, and Sun OS; or an example of a Mac II NuBus PCB PROM program for a small shared memory card that generates hardware interrupts. The current software market does not provide reactive information in the quantities desired by software developers.

Electronic mail and bulletin boards have opened up a new dimension in personal communication. Only in exceptional cases, however, are people highly motivated enough to ensure that the information that flows through the e-mail system is valuable to their organizations. A quick scan through any of the popular computer network services reveals that the majority of the digital information flowing through these networks are recipes, jokes, sports & news opinions, and random debates. Only a small portion of the information can be considered as having anything other than therapeutic entertainment value to workers in government and industry.

Automated repositories of information are only successful to the extent that they are capable of properly calculating a valuation for a particular piece of information. Since this problem is at least as hard as assigning a value to a stock or a bond it is not likely that total automated control of an information repository will be entirely satisfactory.

We envision using the Exchange for large systems development as well as yes/no questions. To the large systems developer the Exchange offers an immense international labor pool that would otherwise be unavailable. To the person with the short simple question like "Are the NuBus implementations on the NeXT and Apple Mac II machines compatible?" the Exchange offers a vast range of immediately accessible technical expertise. To illustrate the Exchange offer format we present the following three examples.

Example A: Information Exchange Offer A

(*1*) Offer price = $50 U.S.

(*2*) Offer duration = 5 days starting 12:00 am EST July 4, 1990.

(*3*) Terms - Nonexclusive

(*4*) Request - Provide a list of two (2) qualified senior hardware engineers available as consultants or permanent employment in Palo Alto, California. Suitable engineers must have significant experience producing prototype 32-bit bus (Apple Nu, IBM MCA, or VME bus) printed circuit boards using high speed optical components (150 MHz or faster TAXI chips from AMD, Gazelle, or Honeywell). Required Information includes:

- Two names, addresses, telephone numbers, and e-mail addresses

- C.V.'s of engineers on list
- Salary history of engineers on list

(*5*) No Restrictions

(*6*) Posting - Employment, Senior Engineer, Northern California, 32-bit bus hardware, Prototype printed circuit boards, Optics, TAXI, AMD, Gazelle, IBM, Apple, Honeywell, NuBus, MicroChannel Architecture, VMEBus.

Example A is a sample request for specific data, possibly available through other sources such as headhunters and consultants, that would be more efficiently serviced through the knowledge-base of a nationwide Information Exchange. Assuming that Exchange postings will be widely circulated, requests such as this gain direct access to a large pool of potential employees. This example also exploits the parallelism inherent in the execution of nonexclusive offers. Paying $50 for two leads is rather expensive, however, since this offer was made under nonexclusive terms the contractor may see tens or even hundreds of names simultaneously for the specified price. The contractor need only accept one response to fulfill his obligation under the Exchange contract. Potential contractees are exposed to the risk that their response will not be selected, however, significant financial compensation is being offered for bearing such risk.

Example B: Information Exchange Offer B

(*1*) Offer price = $30,000 U.S.

(*2*) Offer duration = 90 days starting 12:00 am GMT July 4, 1990.

(*3*) Terms - Exclusive

(*4*) System Specification - A 3278/79 emulation package written in AJPO verified ADA to be used on an IBM 3081-D mainframe running OS/MVS, TSO, and ISPF. Terminals will be attached through coaxial cable to a 3274/76 cluster controller communicating via a T1 carrier to a remote 3705 communications controller using IBM's SNA protocol. Required Services and Information include:

- 3,564 characters displayed - 27 lines x 132 chars.
- extended attribute bytes for 3278
- enhanced (16) color for 3279
- graphics adapter support 3279
- > 90% standard key and function duplication
- DOS/Xenix - terminal emulation switching
- File transfer - text and binary
- ASCII-EBCDIC translations
- Detailed Installation Guide
- User Level Documentation
- < other services >
- Source Level Documentation

- Ada Source.

(5) Restrictions

(a) Testing Data Set is as follows <.....test vectors...> Testing random input vectors in test configuration <.. test configuration specifications...> must find no more than 1 error in 10,000,000 execution trials.

(b) Contractor requires exclusive ownership of resulting systems - publicly available systems are prohibited.

(c) Contractee must be a West German Corporation.

(6) Posting - Ada, 3278/79 terminal emulation, IBM, OS/MVS, TSO, ISPF, SNA, 3705 communications controller, IBM 3081-D mainframe, DOS, Xenix.

Example B presents an outline of a standard software system specification similar to Broad Agency Announcements that appear in the Commerce Business Daily issued by the U.S. Department of Commerce and the U.S. Government Printing Office. In this case the contractor is simply specifying a software system that is functionally equivalent to commercially available terminal emulation packages. By specifying that the software must be written in Ada and that source must be supplied in the response it becomes clear that the commercially available systems do not fit this specification. Since the contractor specifies under item 5b) in Restrictions that he will retain exclusive ownership of the code we might assume that the resulting system will be part of some larger commercial system.

Example C: Information Exchange Offer C

(1) Offer price = $5,000 U.S.

(2) Offer duration = 50 days starting 12:00 am EST July 4, 1990.

(3) Terms - Exclusive

(4) Request - Voice recognition software module/hardware interface with Panasonic Composite Studio Digital Recorder AJ-D350 compatible with demonstration P&T multimedia system. The software module must be compatible with the following P&T source code running UNIX vXXX.

- <..... source code .....>
- <..... hardware specs...>
- <..... software environment specs ...>

(5) Restrictions

(a) Contractee must, in the past 3 years, have been an employee, consultant, or contractor of AT&T Bell Laboratories, Murray Hill, New Jersey - in Area 11.

(b) All testing of re-engineered code must be conducted at AT&T Bell Laboratories, Murray Hill, N.J. (guaranteed 8 hr./day single user access to equipment specified above).

(*6*) Posting - Software Re-engineering, Voice Recognition, AI, P&T System, Panasonic, UNIX vXXX, Multimedia, Area 11, AT&T Bell Laboratories, Northern New Jersey.

Example C illustrates the utility of the Exchange in exploiting domain specific knowledge and reuse of software modules. Corporate software expertise does not simply disappear even if expert employees change jobs. Even if the employee base becomes geographically distributed over time the accumulated expertise of a given set of employees is preserved. People with experience on a given piece of software may not currently be employed by the project, however, if expert employees are informed and offered proper compensation they may be willing to solve specific software development or administration problems.

4. Prognosis for The Information Exchange

An Information Exchange over an international computer network will be established in the near future because:

(*a*) Running a successful Information Exchange will be very lucrative and

(*b*) The economic and technological benefits of high volume information trading effect every corporation's profitability.

The success of such an Exchange hinges on the specification problem. That is, how well specifications and questions included in offers match the systems and answers included in the responses (did the customer get the product they ordered). All the other technical, legal, and administrative problems are of secondary importance. If the customer believes that he did not receive proper compensation for his data or that he was coerced into purchasing inappropriate or incorrect data then no amount of technical, legal, or administrative assistance will make the Exchange a success.

There are good reasons to believe that existing solutions to the specification problem are sufficient for the needs of the Exchange. Eminent mathematicians such as Paul Erdos routinely trade solutions to difficult problems for cash. Sending Erdos the first correct solution to a specific geometry problem might get you $500.00. The computer scientist Donald Knuth offers readers of his textbooks $2.00 for every new typographical error they find. Freelance software consultants, printed circuit board manufacturers, and system integration houses have built large systems to customer specifications for many decades. Thus it appears that there is significant experience within the technological community in writing, and responding to, complete systems specifications. Certainly, the Exchange will receive incomplete or ambiguous specifications like " I'll pay $5,000.00 for a fast computer program that will do my accounting today" or "I'll pay $300.00 to know how to program an Amiga," however, we expect that these requests will be easily identifiable and exceptional.

The technical problems that confront the establishment of a successful Information Exchange computer network are solvable with state-of-the-art technology. Issues of reliability, fault-tolerance, security, and interoperability are foremost among the technical concerns. The short list of critical nontechnical problems, both implementation and legal, include: establishing wide circulation for postings, copyright and other information ownership issues, contract enforcement, and adjudication. Many of these problems can be sidestepped initially if the Exchange operates entirely within a single large corporation (i.e., MEI, Mitsubishi, IBM). As the Exchange expands to a national or international scale issues of ownership and enforcement will become more challenging.

The proposed market uses automated hardware and software systems to enhance rather than replace human intelligence. We are proposing a cybernetic intelligence system as opposed to an artificial intelligence system. Perhaps it will be the case that researchers ultimately will completely automate the retrieval of data from the massive knowledge-base of corporate America. Such a general capability is not expected in the near future. For some specific domains, such as ground combat, aircraft, and submarines there exist realistic hopes of near term automated software repositories. Lotus' Notes and AT&T's Rhapsody represent a significant commitment on the part of industry to solve this problem. Building an automated information retrieval system around a motivated population of informed employees and expert technologists affords productivity gains and cost savings significantly larger than can be expected in the context of purely automated data retrieval.

The proposed market promotes explicit financial compensation for assuming risk. From the contractor's perspective the Exchange represents a large pool of laborers whom he can potentially direct to solve a specific problem. Theoretically, a commercial software contractor could replace his entire software development staff with Information Exchange contracts. In such a case the contractor's software development capability could be tied more closely to his fluctuating economic needs. The contractor can temporarily increase or decrease his software development capacity to a fine degree that is impossible today.

From the contractee's perspective a fluid market for his services enables him to consider assuming greater risk. The heightened risk might be in the form of responding to simple question posed in a nonexclusive contract or by responding to a large systems development offer, having only modest manpower leveraged by special productivity tools.

The proposed market would, according to the Fair Market Hypothesis, maintain fair prices for information and encourage competition in providing access to valuable information. The enhanced competition promoted by the Information Exchange has the obvious benefit of bringing the price down for both data and programs. Less obvious is the effect the Exchange will have on how the data and programs are produced. In the software engineering domain, for example, the software tools that dramatically increase productivity will become valuable assets. Contracts that employ these tools will be more likely to be successfully executed. Hence the contractees that learn how to use these tools will be financially rewarded. Ineffectual software engineering technologies, leading to failed contracts, will rapidly be identified as financial burdens.

Strong competition will not only sort out the profitable and unprofitable software engineering techniques, but it will spur the discovery of new more powerful techniques. Existing pieces of code may be submitted to the Exchange with the request that contractees try to find errors or demonstrate how the code breaks. Multiple exclusive contract offers from the same contractor may be simultaneously issued to the Exchange to ensure easier correctness testing (given the same input do the two programs do the same thing?). Assertions about a piece of code's behaviour may be tested by issuing a nonexclusive contract containing a true/false question (Is true that this program <... program specification ...> eventually terminates?). We expect that much of the software and hardware systems used to implement the Information Exchange will be developed through the Exchange itself.

The proposed market offers efficient exploitation of an existing underutilized corporate asset - the knowledge base of its employees. New opportunities for financial rewards will lead to a more highly motivated, responsive, and informed workforce. As a management tool an employee's track record on the Exchange will be a strong indicator of his reliability, responsiveness, and overall technical competence.