and association. Technologies for tracking people and conducting surveillance of public space risk chilling the freedom of association on which any possibility of democratic community is based. Yet the nature of that risk is still obscure. So long as privacy issues are located solely in the lives and relationships of individuals, it will be impossible to conceptualize either the potentials or the dangers of new media technologies for a democratic society.

Acknowledgements

I greatly appreciate the contributions of Marc Rotenberg, as well as the comments and suggestions of Colin Bennett, Trotter Hardy, and Robert Horwitz.

References

Agre, Philip E. 1994. Surveillance and capture: Two models of privacy. *The Information Society* 10, no. 2: 101–127.

Agre, Philip E. 1995. Reasoning about the future: The technology and institutions of Intelligent Transportation Systems. *Santa Clara Computer and High Technology Law Journal* 11, no. 1: 129–136.

Allen, Anita L. 1995. *Uneasy Access: Privacy for Women in a Free Society.* Rowman and Littlefield.

Bennett, Colin. 1992. *Regulating Privacy: Data Protection and Public Policy in Europe and the United States.* Cornell University Press.

Bertino, Elisa, and M. Tamer Ozsu, eds. 1994. *Distributed and Parallel Database Object Management*. Kluwer.

Bijker, Wiebe E. 1987. The social construction of Bakelite: Toward a theory of invention. In *The Social Construction of Technological Systems*, ed. W. Bijker et al. MIT Press.

Bowker, Geoffrey. 1994. Information mythology: The world of/as information. In *Information Acumen: The Understanding and Use of Knowledge in Modern Business*, ed. L. Bud-Frierman. Routledge.

Bowker, Geoffrey, and Susan Leigh Star. 1994. Knowledge and infrastructure in international information management: Problems of classification and coding. In *Information Acumen: The Understanding and Use of Knowledge in Modern Business*, ed. L. Bud-Frierman. Routledge.

Brackett, Michael H. 1994. *Data Sharing: Using a Common Data Architecture*. Wiley.

Branscomb, Lewis M., and James H. Keller, eds. 1996. *Converging Infrastructures: Intelligent Transportation and the National Information Infrastructure*. MIT Press.

Brenton, Myron. 1964. *The Privacy Invaders*. Coward-McCann.

Burnham, David. 1983. *The Rise of the Computer State*. Random House.

Casson, Mark. 1994. Economic perspectives on business information. In *Information Acumen: The Understanding and Use of Knowledge in Modern Business*, ed. L. Bud-Frierman. Routledge.

Cavoukian, Ann, and Don Tapscott. 1997. *Who Knows: Safeguarding Your Privacy in a Networked World*. McGraw-Hill.

Chaum, David. 1992. Achieving electronic privacy. *Scientific American* 267, no. 2: 96–101.

Clarke, Roger. 1994. The digital persona and its application to data surveillance. *Information Society* 10, no. 2: 77–92.

Coase, Ronald H. 1937. The nature of the firm. *Economica NS* 4: 385–405.

Cronon, William. 1991. *Nature's Metropolis: Chicago and the Great West*. Norton.

David, Paul A. 1985. Clio and the economics of QWERTY. *American Economic Review* 72, no. 2: 332–337.

Drake, William J., and Kalypso Nicolaidis. 1992. Ideas, interests, and institutionalization: 'Trade in services' and the Uruguay Round. *International Organization* 46, no. 1: 37–100.

Ernst, Morris L., and Alan U. Schwartz. 1962. *Privacy: The Right to Be Let Alone*. Macmillan.

Flaherty, David H. 1989. *Protecting Privacy in Surveillance Societies: The Federal Republic of Germany, Sweden, France, Canada, and the United States*. University of North Carolina Press.

Gandy, Oscar H. Jr. 1993. *The Panoptic Sort: A Political Economy of Personal Information*. Westview.

Goffman, Erving. 1957. *The Presentation of Self in Everyday Life*. Anchor.

Gostin, Lawrence O. 1995. Genetic privacy. *Journal of Law, Medicine, and Ethics* 23, no. 4: 320–330.

Information and Privacy Commissioner (Ontario) and Registratiekamer (The Netherlands). 1995. Privacy-Enhancing Technologies: The Path to Anonymity (two volumes). Information and Privacy Commissioner (Toronto) and Registratiekamer (Rijswijk).

Inness, Julie C. 1992. *Privacy, Intimacy, and Isolation*. Oxford University Press.

Latour, Bruno. 1987. *Science in Action: How to Follow Scientists and Engineers Through Society*. Harvard University Press.

Lyon, David, and Elia Zureik, eds. 1996. *Computers, Surveillance, and Privacy*. University of Minnesota Press.

Nugter, A. C. M. 1990. *Transborder Flow of Personal Data within the EC: A Comparative Analysis of the Privacy Statutes of the Federal Republic of Germany, France, the United Kingdom, and the Netherlands*. Kluwer.

Packard, Vance. 1964. *The Naked Society.* McKay.

Porter, Theodore M. 1994. Information, power and the view from nowhere. In *Information Acumen: The Understanding and Use of Knowledge in Modern Business*, ed. L. Bud-Frierman. Routledge.

Posner, Richard A. 1981. *The Economics of Justice*. Harvard University Press.

Regan, Priscilla M. 1995. *Legislating Privacy: Technology, Social Values, and Public Policy.* University of North Carolina Press.

Registratiekamer. See Information and Privacy Commissioner.

Rotenberg, Marc. 1996. Personal communication.

Schoeman, Ferdinand David, ed. 1984. *Philosophical Dimensions of Privacy: An Anthology*. Cambridge University Press.

Sciore, Edward, Michael Siegel, and Arnon Rosenthal. 1994. Using semantic values to facilitate interoperability among heterogeneous information systems. *ACM Transactions on Database Systems* 19, no. 2: 254–290.

Smith, H. Jeff. 1994. *Managing Privacy: Information Technology and Corporate America*. University of North Carolina Press.

Smith, Robert Ellis. 1979. *Privacy: How to Protect What's Left of It.* Anchor Press.

Turkington, Richard C. 1990. Legacy of the Warren and Brandeis article: The emerging unencumbered Constitutional right to informational privacy. *Northern Illinois University Law Review* 10: 479–520.

Westin, Alan F. 1967. *Privacy and Freedom*. Atheneum.

Williamson, Oliver E. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press.

1

Beyond the Mirror World: Privacy and the Representational Practices of Computing

Philip E. Agre

In January 1996, the California Air Resources Board (ARB) issued a Request for Proposals (RFP 95-7) entitled Incorporation of Radio Transponders into Vehicular On-Board Diagnostic Systems. The ARB observed that, starting in 1996, new cars and light trucks in California will be equipped with an on-board diagnostic (OBD) system that illuminates a dashboard indicator when it detects a malfunction in the vehicle's emissions system. It also observed that drivers may not actually get their malfunctioning cars fixed until their next scheduled inspection check. With an eye to enforcing compliance with emissions laws, therefore, the ARB proposed to build a fleet of ten test cars, each equipped with a transponder capable of transmitting the OBD system's error codes to roadside or satellite-based receivers. Specifically, in response to a query from the receiver, the RFP specifies that the transmitter be able to supply the following information:

- the full 17-digit vehicle identification number
- any fault codes reported by the OBD system
- the location of the vehicle at the time of query
- a status code.

The receiver is to be able to store the dates and times of all queries to passing cars, along with the information that the cars return, including "vehicle location (to the zip code level, and city)." Although this RFP only envisions the construction of test vehicles, the successful bidder is asked to analyze the system's cost effectiveness "assuming that 1,000,000 vehicles equipped with the [transponder-equipped diagnostic system] are added to the fleet beginning in the year 2000."

The ARB system is presumably intended to be used only for its stated purpose: ensuring compliance with emissions regulations. The system as specified, however, could easily be used to track the location of every car that is equipped with it, regardless of whether it has any fault codes to report. The potential for abuse ought to figure significantly in any weighing of risks and benefits from the system, especially given that the newer cars on which the system would first be installed are the cars least likely to experience emissions-system malfunctions. What is most striking about the RFP, though, is that it takes for granted a whole vocabulary of technical methods that has become familiar and widespread. The winning bidder, for example, is instructed to "investigate the possibility of coordinating this effort with other agencies or entities," so as to "suggest the most effective and efficient infrastructure for statewide electronic fleet monitoring." The "other agencies" include the California Department of Transportation, which has been working on transponder-based toll-collection projects for several years, using a transponder architecture that is expressly designed to be extensible to other applications.

The ARB is not unaware of the privacy concerns that its plans may raise. These concerns are described at length in a report on legal issues (Di Genova and Macomber 1994) by Sierra Research, the previous contractor on the monitoring project. They have also been reported in the press (Rogers 1996; Halpert 1995). Although the project may be derailed by protests from affected industries, on a technical level the system retains a sense of inevitability; it seems clearly the logical culmination of a path of technical development. Although wireless capture of onboard fault codes has been rendered practical only recently, through decreasing costs of microprocessor power and digital wireless communications, little about the RFP is technically new or difficult; any competent computer systems designer can imagine how it might be embodied in a working system. Such a system, moreover, would be similar in its workings to hundreds of other computer-based technologies for tracking human activities that have arisen in recent years, or that are currently being developed. Together these technologies constitute an emerging infrastructure—decentralized to be sure but pervasive nonetheless—for the systematic surveillance of numerous spheres of human life.

Ever since the rise of organizational computing in the 1950s, social imagination has always associated computing with surveillance. The reasons for this association are clear enough: computer systems analysis and design promptly took up and generalized the methods of rational administration that organizations had developed throughout the modern era (Clement 1988). The technical concept of "algorithm" was assimilated to the bureaucratic concept of "procedure," and the technical concept of "data" was assimilated to the bureaucratic concept of "files." So long as computers have been understood as inherently suited to social relationships of rationalized control, debate about privacy has been structured as a series of tradeoffs and by the assertion of abstract rights (dignity, autonomy, association, self-determination) against specific encroachments (each accompanied by a compelling justification). Lyon (1994), for example, accurately notes that the same technologies that have raised concerns about a "surveillance society" have historically made possible many benefits that most citizens would prefer not to surrender. Society, prodded in part by the resistance and agitation of concerned groups, may respond to such dilemmas by shaping data-protection regulation, but the resulting regulatory process must always contend with the putative benefits of particular surveillance technologies.

The technical methods of computing have always defined the landscape across which these controversies are conducted (Pool 1983, p. 251). Until very recently, though, these technical methods were not themselves matters of dispute. Computers as cultural objects have, in effect, been essentialized—treated as a stable category with fixed relationships to larger social practices. This ahistorical understanding of computers has eroded somewhat in recent years as citizens have experienced the endless turmoil of the personal computer industry firsthand. The emergence of computer networking technologies such as the Internet, moreover, has popularized an alternative, liberatory construction of computers that focuses on their communications function (Feenberg 1995). Significant changes in computer architecture do appear occasionally, such as the emergence of RISC architectures in the early 1980s (Molina 1993; Radin 1983). Nonetheless, certain fundamental aspects of computer architecture have remained stable. The vast majority of computer processors are still serial, and virtually all of them employ the register-transfer model that first arose in the

1940s. Calculations are still overwhelmingly digital, and computer memory still consists of a hierarchy of media, each indexed by a linear sequence of addresses. On the whole, computing has progressed by augmenting and accelerating basic architectural choices, not by replacing them. Although technical research projects have proceeded in many directions, the innovations that have become widespread in actual practice have generally been built on top of existing methods, or they provide generalizations or rational reconstructions of existing practice.

The architecture of a computer does not completely determine its use in institutional contexts. Even so, the design of computer systems and the practices of rationalization have evolved together in a slow, continuous motion for 50 years. The purpose of this chapter is to explore this evolution. Although some valuable histories of system-design practice have been produced (Edwards 1996; Friedman 1989), my particular goal here is to gather some fragments toward a history of data. Computers raise privacy issues when organizations use them to represent human beings and their lives, and by analyzing the history of these representational practices I hope to make the architecture and the social practices of computing seem less immutable. This history, as subsequent sections will make clear, is not a matter of linear progress. The concept of data changes over time, impelled both by the metaphors of technical language and by technical practitioners' experiences in trying to get their systems to work. Toward the end of the story, a mainstream, data-centered model of computing emerges alongside a model based on public-key cryptography. I will argue that the various models of data, despite their seeming differences, are coherently related to one another and to the institutions of computing.

Although I will sketch these models of data, my main purpose is not technical exposition. Instead, I wish to explore the "practical logic" of computing work (Agre 1997). What happens when a technical community with a definite worldview, perhaps influenced by certain metaphors, puts that worldview into practice? Any given metaphor will necessarily consign certain aspects of reality to the practitioners' peripheral vision. A metaphor may even misrepresent the very methods that supposedly embody it. To the extent that practitioners apply their methods in concrete situations, though, the marginalized features of reality will reassert themselves, most likely in the form of recurring patterns of difficulty whose nature may not be clear at first. The progress of technical work is, in this sense, largely reactive: the technical community applies certain metaphors in organizing its work, a certain pattern of difficulty emerges, and the existing methods are modified or extended according to some understanding of the difficulties. Only rarely does an entirely new metaphor arise, and then the old and new metaphors are likely to coexist, whether as competitors or complements. New methods arise when a new metaphor makes sense of large parts of practitioners' experiences, but they are adopted only when they can be fitted into the sprawling network of methods by which the technical community as a whole designs and builds complex systems. In this way, computer work is influenced by its language, by the mutual constraints of its methods, and by the recalcitrance of the world that lies beyond both. Reconstructed historically, this development makes a certain kind of sense, and this sense is the practical logic of the work.

Representing Activity

The first methods for representing human activities on computers were derived from the work-rationalization methods that industrial engineers had been developing since the 1910s. Hence the phrase "information processing": the idea was that computers automate a kind of factory work whose raw material happens to be information rather than anything physical or tangible.

Industrial engineers, starting with Gilbreth (1921) and Taylor (1923), developed two broad categories of representations: process charts for representing the flow of materials through a manufacturing process, and time-and-motion diagrams of particular jobs. Much has been written about the sociology of early work rationalization (Braverman 1974; Guillén 1994; Montgomery 1987; Nelson 1980; Thompson 1989), but I am specifically concerned with work rationalization as a representational practice. How, and in what sense, do computers come to embody representations of human activities?

Figure 1 presents an example of a process flow chart for the baking of soda crackers. Note, first of all, that the representational scheme consists

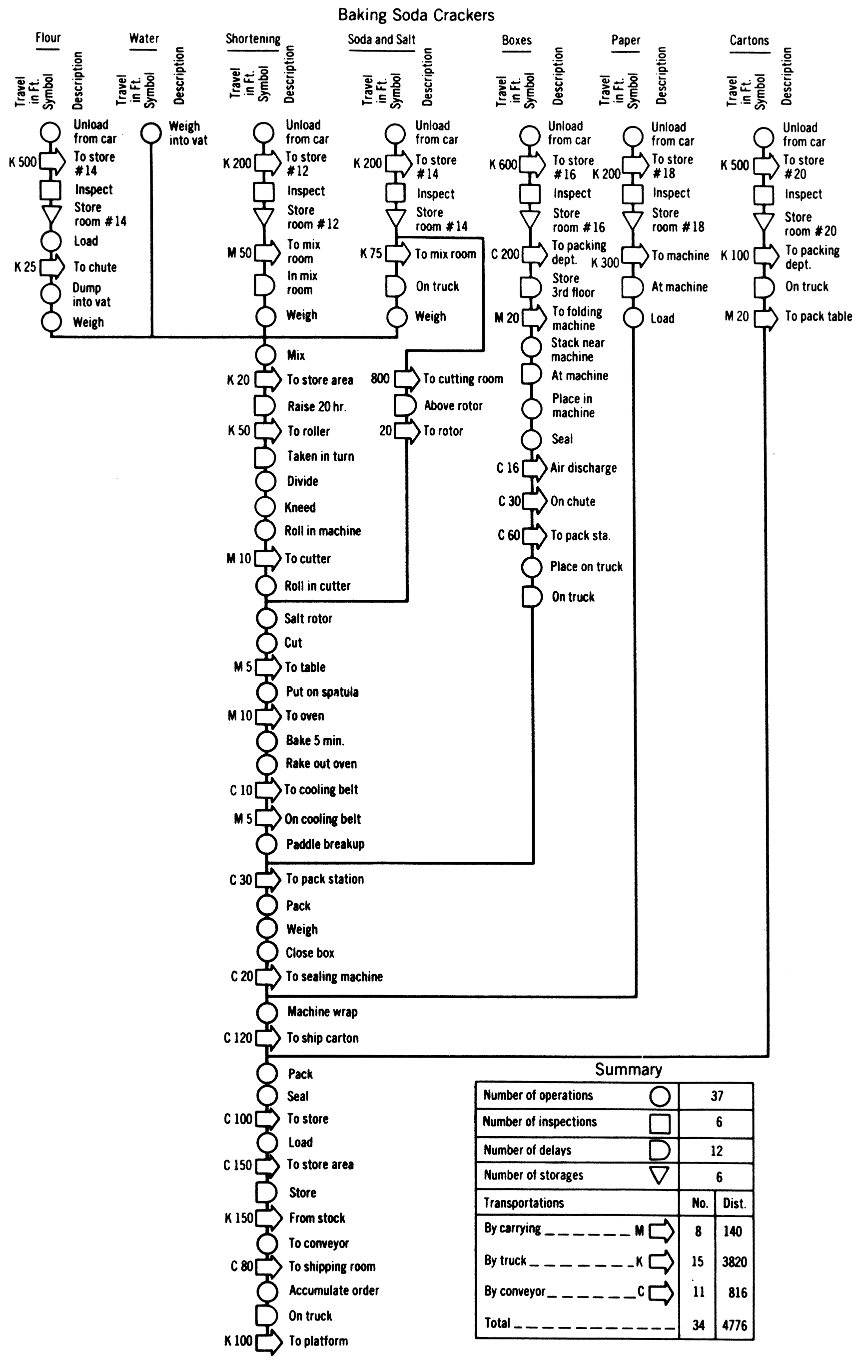

Figure 1
A process flow chart. Source: Ralph M. Barnes, *Motion and Time Study*, seventh edition (Wiley, 1940), p. 76.

of a small vocabulary of discrete elements. These elements can be used both to represent an existing form of activity and to prescribe a new one. In each case, the scheme presupposes that activities can be broken down into discrete units. These units need not be sequential (i.e., they might overlap), but they identify discrete categories of action into which particular empirical streams of action must be parsed.

Figure 2 presents Gilbreth's elements of motion, "therbligs." Once again, the idea is that work activities can be decomposed into a fixed set of categories (search, find, select, etc.). Figure 3 provides an example of the therblig notation scheme in action; in this case, each action is decomposed into to the respective roles of the worker's left and right hands. The therblig system includes units for representing rest and delay, and it includes a unit to represent moments when the worker must stop to construct a plan. This latter unit is distinctive. In Derrida's (1976 [1967]) terms, the "plan" therblig is the supplement of the system: an extra term that closes up a gap in the system while simultaneously embracing it as a whole. The whole idea of motion study, after all, is to minimize or eliminate this extra term by displacing the planning effort into the work of engineers. By making this supplementary term an explicit part of his representational scheme, Gilbreth acknowledges, if only indirectly, that reality always surpasses any particular representation of it. Subsequent schemes in the same tradition often did not make room for explicit planning, but they did build upon a vast body of practical experience in getting such schemes to work.

Couger's (1973) survey of the emergence of systems-analysis methods in computing makes explicit the connection between Gilbreth's use of process charts and later computing practice. The initial concern was the flow of forms in an organization. The rationalized organizational use of paperwork was highly developed before the invention of the computer (Yates 1989), and computing practice drew on an established tradition of automating paper-based work. For example, figure 4, a process chart, represents the flow of Hollerith cards through various tabulating machines. Each card corresponded to a particular record, in this case records specifying orders and production levels for various models of a company's products. As this figure makes clear, techniques used to map the flow of parts in manufacturing were carried over to the flow of documents in the

Name of Symbol	Therblig Symbol		Explanation-suggested by	Color	Color Symbol	Dixon Pencil Number	Eagle Pencil Number
Search	Sh		Eye turned as if searching	Black		331	747
Select	St		Reaching for object	Gray, light		399	734½
Grasp	G		Hand open for grasping object	Lake red		369	744
Transport empty	TE		Empty hand	Olive green		391	739½
Transport loaded	TL		A hand with something in it	Green		375	738
Hold	H		Magnet holding iron bar	Gold ochre		388	736½
Release load	RL		Dropping content out of hand	Carmine red		370	745
Position	P		Object being placed by hand	Blue		376	741
Pre-position	PP		A nine-pin which is set up in a bowling alley	Sky-blue		394	740½
Inspect	I		Magnifying lens	Burnt ochre		398	745½
Assemble	A		Several things put together	Violet, heavy		377	742
Disassemble	DA		One part of an assembly removed	Violet, light		377	742
Use	U		Word "Use"	Purple		396	742½
Unavoidable delay	UD		Man bumping his nose, unintentionally	Yellow ochre		373	736
Avoidable delay	AD		Man lying down on job voluntarily	Lemon yellow		374	735
Plan	Pn		Man with his fingers at his brow thinking	Brown		378	746
Rest for overcoming fatigue	R		Man seated as if resting	Orange		372	737

Figure 2
Therbligs (Frank Gilbreth's units of motion). Source: Ralph M. Barnes, *Motion and Time Study*, seventh edition (Wiley, 1980), p. 117.

course of "manufacturing" other documents. The focus was on the documents as physical artifacts, to be processed like any others, and not on the meaning of the various representational elements in the documents. The representational nature of the documents, in other words, was secondary. The systems analyst's job was to represent the forms, not the things that the forms represented. The transition from paper forms to Hollerith cards

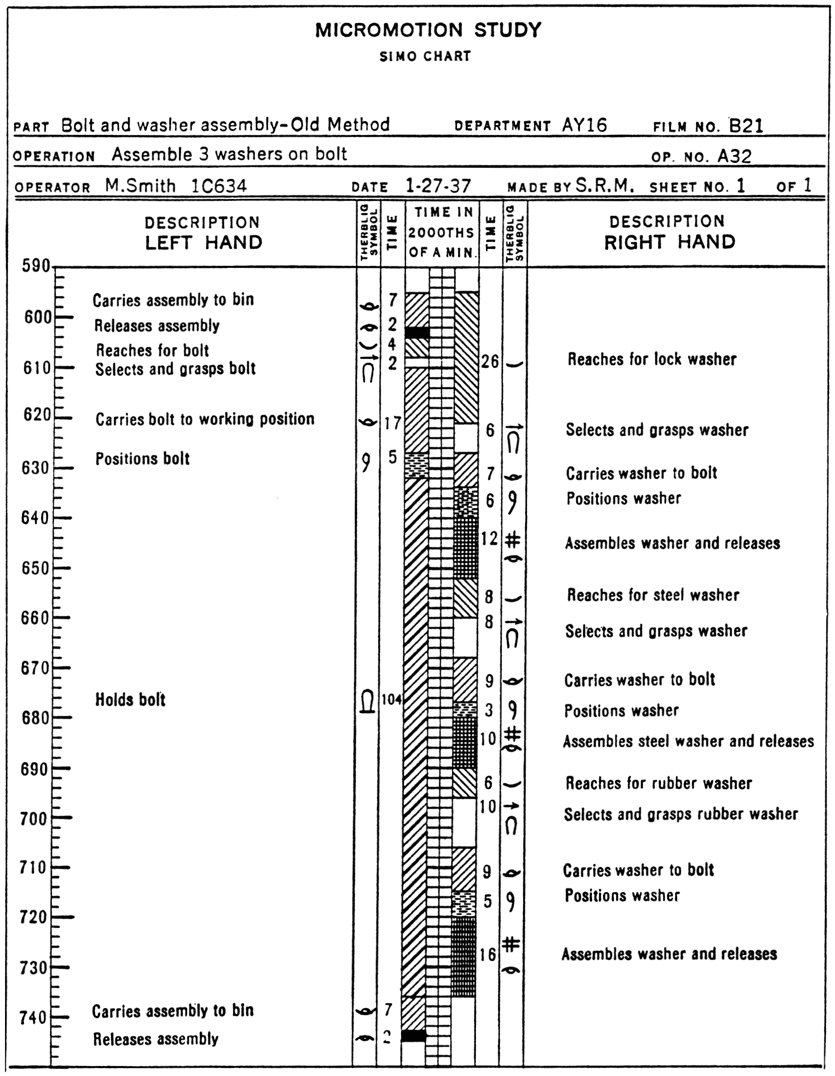

Figure 3
Simo-chart analysis. Source: Ralph M. Barnes, *Motion and Time Study*, second edition (Wiley, 1940), p. 93.

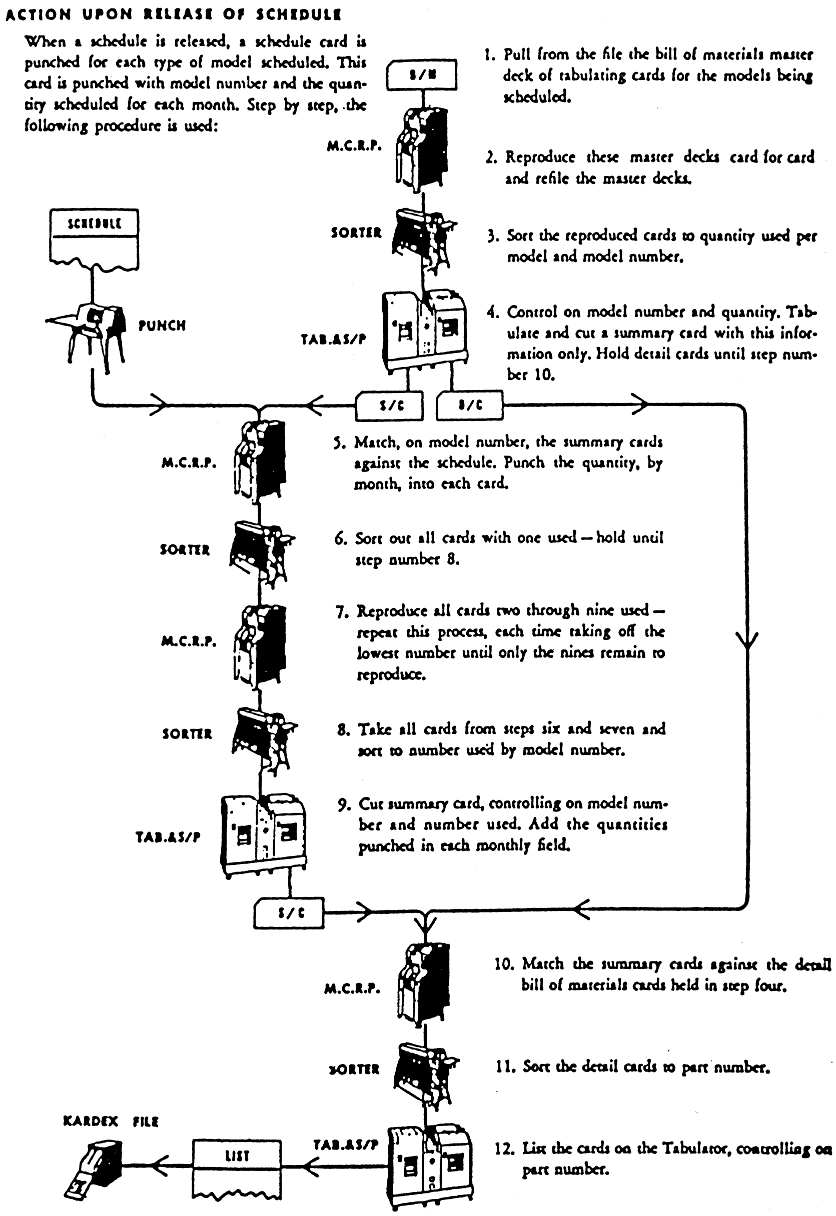

Figure 4
A tabulating-machine process chart. Source: J. Daniel Couger, Evolution of business system development techniques, *Computing Surveys* 5, no. 3 (1973): 167–198. Courtesy of Unisys.

to database entries was gradual and uneven, and the concept of information served to abstract the common functionality of these media without drawing any particular attention to the representational nature of the media, or to the facts about the world that these media conveyed. This accounts for a curious feature of the metaphor of information processing: information is treated as a tangible, measurable substance, but only in a minimal sense does the information seem to be *about* anything (Agre 1995a). Couger (1973, p. 196) concludes his survey by pronouncing it "amazing" that systems analysts had made little use of computers to construct the representations that they employed in their work, even though their work consisted in the design of computerized representations. But in retrospect the discrepancy makes sense: the early systems analysts did not recognize much commonality between the representations (process charts) that they manipulated in their work and the representations (office documents) that were manipulated by the people they studied.

The Mirror World

The metaphor of information processing proved misleading in important ways. Having been motivated by the automation of existing work methods based on paper forms, it was not particularly helpful to database designers, whose task was to devise record structures for new categories of digital information. The original databases were inevitably modeled on the tabular information that had long been stored on paper forms, with rows that corresponded to particular things in the world and columns that corresponded to particular attributes of those things.

This view of data was formalized by Chen (1976) in the entity-relationship model of data. Although many databases had been designed in the two decades before Chen's influential paper, the entity-relationship model brought database design into focus as a body of practitioner experience. Perhaps its greatest conceptual innovation was a clear distinction between data records and the entities that they represent. Though seemingly abstract, the point was not philosophical but organizational; by starting the design process with an enumeration of entities, the database designer is supposed to ensure that the resulting design reflects actual business needs. Database designers, after all, were generally trained as general-purpose system analysts rather than in the particular industry whose substance the data was to represent, and their attention was naturally focused more on databases as such than on the categories that were implicated in particular applications domains. Nor was the point at all obvious; even a recent manual of data modeling (Reingruber and Gregory 1994), in trying to define the concept of an entity, repeatedly blurs the distinction between records and things. Having surveyed various definitions of "entity" in the literature, Reingruber and Gregory (p. 64) extract the following themes:

An entity is a "thing" or "concept" or "object." . . . Or it may represent the fact that we need to capture information specific to only a subset of the instances of an entity. . . .

An entity is not a single "thing," but rather a representation of like or similar things that share characteristics or properties. For example, *King Lear* and *Hamlet* are both plays, and have properties such as name, author, cast of characters, and lines of verse. The entity describing these might be PLAY, with *King Lear* and *Hamlet* as examples or instances or occurrences of PLAY. . . .

Entities involve information. The "things" that entities represent are things about which the enterprise wants or needs to retain information.

These sentences shift repeatedly between treating entities as things in the world and treating entities as *representations* of things in the world. Reingruber and Gregory's choice of example facilitates the confusion: the word "play" refers to both the text and the performance—both the representation and the thing represented. Reflecting on the above points, they "refine our view of entities" as follows:

Entities are collections of attributes that satisfy a particular set of rules established for relational modeling. (ibid., p. 64)

This definition is perfectly ambiguous. Real things can, with some effort, be viewed figuratively as collections of attributes, but entity-relationship data records are precisely collections of data elements called attributes. Reingruber and Gregory may have been misled by the practice, common but usually well defined in the literature, of using the word "entity" to refer both to categories of things (cars, people, companies) and to instances of those categories (my old 240Z, Jacques Martin, IBM). Be this as it may, the conflation of representation and reality is a common and significant feature of the literature of computer science. (In artificial intelligence, for

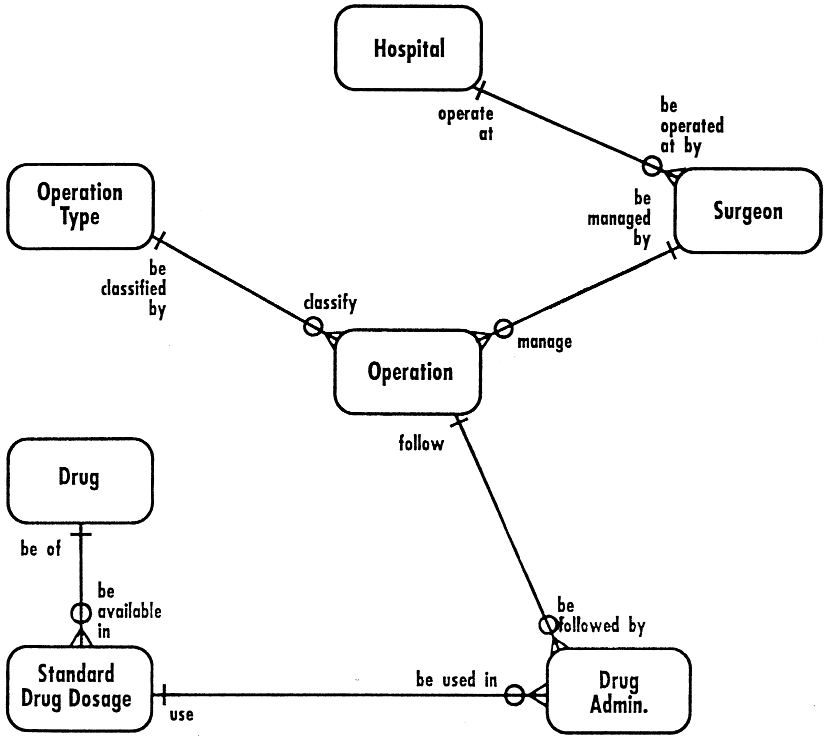

Figure 5
A data model. Source: Graeme Simsion, *Data Modeling Essentials* (Van Nostrand Reinhold, 1984), p. 59.

example, formally specified activities such a chess and logical theorem proving serve the purpose that plays serve for Reingruber and Gregory.)

The entity-relationship approach employed a graphical notation, a recent and relatively streamlined version of which is shown in figure 5. Each box in this figure stands for an entity and each link stands for a relationship; the symbols on the links indicate whether the relationship is one-to-one (for example, cars and engines) or one-to-many (for example, cars and wheels), and whether the relationship exists necessarily (every car has at least one seat) or optionally (a car may have a roof but need not). This notation scheme can express distinctions of some ontological delicacy, but only up to a point. During the same period, AI research in knowledge representation developed a complex multilevel analysis of ontological and epistemological categories for the purpose of representing structural relationships that are much more complex than have usually arisen in business data (Brachman 1979; Borgida 1991). Within database research, similar but less ambitious work on semantic and conceptual data modeling has had little influence on mainstream practice. Instead, the mainstream focus has been on the business significance of large databases of relatively simple records that can be searched quickly.

During the 1980s, then, data became increasingly central to the practice of computing. Even though the design of a database accounts for a minor portion of the overall cost of designing a system, the structure of data has profound consequences for the design of algorithms. Once created, data can be used for a variety of purposes, including purposes that were not envisioned when the database was originally designed. Database designers therefore increasingly viewed their work as independent of program design, and they evolved canons of good practice that responded to cases where database designs failed to generalize to new applications. Amidst these imperatives, there arose a new model of computing. Whereas information processing was primarily concerned with the structure of processing, the newer model was primarily concerned with the structure of data. And the purpose of data was to reflect the repertoire of entities in which the business dealt.

The apotheosis of the emerging data-centered view of computing is a popular book entitled *Mirror Worlds*, written by the prominent computer scientist David Gelernter (1991). *Mirror Worlds* repays close reading. As the title makes clear, the book's central metaphor is the computer as a mirror of the world. Gelernter suggests that the progress of computing will inevitably produce a single vast distributed computer system that contains a complete mirror image of the whole of reality. Although the computer as mirror is at best a tacit theme in the database literature, Gelernter makes the mirror metaphor thoroughly explicit and offers it as a way of extrapolating current directions in technical development:

> You set up a software mirror wherever you like, then allow some complex real-world system to unfold before it. The software faithfully reflects whatever is going on out front. (p. 6)

The Mirror World, Gelernter suggests, will permit individuals to investigate reality without leaving home, simply by "traveling" in the digital mirror:

> Capturing the structure and present status of an entire company, university, hospital, city or whatever in a single (obviously elliptical, high-level) sketch is a hard but solvable research problem. The picture changes subtly as you watch, mirroring changes in the world outside. But for most purposes, you don't merely sit and stare. You zoom in and poke around, like an explorer in a miniature sub. (p. 15)

Video cameras figure prominently in Gelernter's story, as befits the optical metaphor of the mirror:

> Eavesdrop on decision making in progress. Among other things, you will discover video feeds down here. When you dive into City Hall, one part of the display on your screen might be a (little) TV picture. You can mouse over there and enlarge the thing, if you want to hear the mayor's press conference or the planning board meeting. (p. 17)

Gelernter is well aware of the privacy concerns that computers have raised, and much of his book is concerned with the nature of public space in the era of the Mirror World:

> The Mirror World isn't snoopware. Its goal is merely to convert the *theoretically* public into the *actually* public. What was always available in principle merely becomes available *in fact*. Organized, archived, spiffily presented, up to the minute and *integrated* into a whole. (p. 19)

As the last phrase suggests, one purpose of the Mirror World is to repair social fragmentation by making it possible to grasp the world as a totality through its computer representation. Gelernter refers to this form of understanding as "topsight":

> If *insight* is the illumination to be achieved by penetrating inner depths, *topsight* is what comes from a far-overhead vantagepoint, from a bird's eye view that reveals *the whole*—the big picture; how the parts fit together. (p. 52)
>
> It used to be generally conceded that whole-sightedness—a due respect for what Madison calls "the permanent and aggregate interests of the community"—was a good thing. Today, all sorts of angry factions are proudly dedicated to the methodical tearing-to-shreds of public life. Rapacious PAC lobbyists in Washington and multiculturalists at Stanford are quite agreed that a little E Pluribus *Unum* goes a long way. (p. 31)

Gelernter views Mirror Worlds (which are sometimes singular, sometimes plural) as extensions of public spaces that can facilitate democratic decision making:

. . . Mirror Worlds function *in part* as fire walls opposing the onslaught of chaos. But they aren't *mere* fire breaks. They are beer halls and grand piazzas, natural gathering places. . . . (p. 6)

Every Mirror World neighborhood is equipped with a public message board. Candidates can post statements. Towards election time, we can set up a special political playground, where they can hold forth at greater length. (p. 23)

In response to fears that computers are necessarily instruments of social control through surveillance, Gelernter argues that the Mirror World will make it possible for ordinary citizens to turn that same power of surveillance back against the state:

A Mirror World is some huge institution's moving, true-to-life mirror image trapped inside a computer—where you can see and grasp it whole. The thick, dense, busy sub-world that encompasses you is *also*, now, an object in your hands. A brand new equilibrium is born. (p. 3)

[Mirror Worlds] scotch that great primal modern fear, to be entangled by the sucker-arms of the modern institutional state, and all those private mini-states within which . . . we hang embedded. They offer penetrating vision; they repair the shattered whole. (p. 184)

In this latter passage, Gelernter suggests that Mirror Worlds will offer defenses against both the state and the "private mini-states," but he does not develop this metaphor of private organizations as states or address the politics of making private organizational information public.

In response to critiques of formal representation that associate it with emotional drives for domination and control (e.g., Walkerdine 1988), Gelernter embraces those drives and promises to democratize them:

Don't mistake this topsight search for a mere *intellectual* drive. It's an emotional quest too. . . . our grandest accomplishments reflect the innate urge to dominate. . . . The pursuit of topsight is intellectually compelling because it is *emotionally* compelling. . . . what lends [Mirror Worlds] a uniquely *potent* potential is the submerged iceberg mass of their emotional appeal. . . . topsight to the millions. . . . (pp. 183–184)

Gelernter is aware that security concerns will arise:

Many Mirror Worlds will contain a good deal of confidential information. Professional thieves will certainly be attracted. (p. 197)

Moreover, despite his major emphasis on those parts of the world that are normatively open to public view, such as public spaces and the state, Gelernter also offers a model of selective access to the sensitive information that these vast databases will surely contain:

> Granted, lots of new information-gathering devices have been installed. But the information they are gathering and feeding into the Mirror World is strictly *public* information—or information to which this particular Mirror World's clientele is entitled. And the Mirror World will discriminate judiciously among visitors. The public at large, for example, is entitled to enter the City Hospital Mirror World, and to learn a good deal about what's going on. Furthermore, anyone is entitled to see his own medical records. But very few people have access to anyone else's, although they are all *stored* down here. Access to *private* information is closely controlled. (pp. 18–19)

To implement this version of access control, Gelernter complicates his metaphor of travel in a Mirror World landscape:

> [The] Chief Cameraman is a crucial figure in the Mirror World landscape. . . . he knows who you are. He knows precisely which televiewers belong to you; which are public; which are private, but you're allowed to see; and which you are *not* allowed to have anything to do with. (p. 196)

Gelernter, clearly, is a technological utopian who seeks to refute the technological pessimism of the social theorists. He regards the Mirror World and its political benefits as inevitable—immanent in the technology as it is already developing. Despite his contempt for interest-group activism and for the entangling tendencies of the modern state, he is not a libertarian. To the contrary, he seeks a democracy of liberal individuals, and he thinks of the Mirror World as a collective undertaking through which the conditions of liberal democracy might be established:

> The real software revolution . . . will center . . . on software that steps over the crucial boundary between private and public. It will have to do with "public software works," with civil software-engineering, with Mirror Worlds. The software revolution hasn't begun yet; but it will soon. (pp. 7–8)

Despite Gelernter's focus on the public sphere, the traditional rhetoric of organizational control (Beniger 1986) resurfaces whenever his standpoint shifts from that of the ordinary citizen to that of the organizational leader:

> A parable: consider the modern, state-of-the-art fighter aircraft. It's so fantastically advanced that you can't fly it. It is aerodynamically unstable

> . . . . Modern organizations are in many cases close to the same level of attainment—except that, when they're out of control, they don't crash in flames; they shamble on blindly forever. (p. 19)

> It isn't easy for a hospital administration with the best of intentions to make sure that bad or inappropriate or newly outmoded or wrongly sequenced or mutually inconsistent or maximally expensive or unnecessarily dangerous procedures are never scheduled. A battery of software agents planted in a Mirror World can watch for bad practices full-time. (p. 20)

This is a point of instability in Gelernter's argument. At some moments he presents the Mirror World as a mirror of the *public* world, but at other moments he presents it as a mirror of the *entire* world. The Chief Cameraman appears on moments when it is necessary to negotiate this difference.

Gelernter's conception of data differs greatly from the metaphor of information processing. That metaphor has not vanished, but it has gone backstage. The computational mechanisms concerned with constructing, maintaining, and using the Mirror World are still called "processes," and Gelernter uses a hydraulic metaphor to describe them:

> Oceans of information pour endlessly into the model (through a vast maze of software pipes and hoses): so much information that *the model* can mimic *the reality's* every move, moment-by-moment. (p. 3)

> A Mirror World is an ocean of information, fed by many data streams. . . . Data values are drawn in at the bottom and passed upward through a series of data-refineries, which attempt to convert them into increasingly general and comprehensive chunks of information. As low-level data flows in at the bottom, the big picture comes into focus on top. (p. 27)

Despite the change of metaphor, then, the technical methods that Gelernter describes are not radically different from those of classical systems design. The mirror metaphor and the information processing metaphor are more like two sides of the same coin, logically independent perhaps but each connected to the other within the discourse of computing as it has evolved historically.

Representation and Professional Distance

These metaphors are constituents of a discourse; what, though, about the practices? Gelernter intends to project the future development of computing technology, and in one sense he is correct: the mirror metaphor provides one way of expressing a central theme of recent technologists' stories about computing. Research on virtual reality, for example, is predicated on the idea that computerized gear attached to a user's body can create immersive "worlds" that deserve to be called a substitute for reality (Benedikt 1991). Likewise, a large community within the world of computer research is concerned with simulating reality for a wide variety of purposes. The US military has heavily funded such research in support of its ambitious attempt to blur the distinction between real and simulated battlefields. One battle during the Persian Gulf War, for example, was thoroughly instrumented and recorded so that later training exercises could recreate it in fine detail and explore different directions that it might have taken (Sterling 1993). This project is continuous with the larger phenomenon that Edwards (1996) calls "cyborg discourse," which treats human beings, social institutions, and the natural world as components of one overarching "system." Even when extravagant ideologies of this kind are missing, the mirror metaphor aptly expresses the most basic procedure of traditional computer system design: prepare a roster of the entities that are present in some part of the world, determine which relationships among those entities are significant for the tasks at hand, and create data records whose structures and interconnections map those entities and relationships one for one. Simply put, the computer profession is trying to reproduce the whole world in the workings of machinery.

Gelernter's story, then, is congruent with much of the overtly articulated motivation for pervasive computerization in contemporary industrial society. His optimism about this story, though, is hard to evaluate. One approach to a critique would be to accept the metaphor and dispute his sense of proportion: Mirror Worlds, while increasing the capacity of citizens to conduct computer-mediated surveillance against the powerful, might be held to increase the already vastly greater capacity of the powerful to conduct surveillance of their own. But this argument, while perhaps accurate, is not particularly compelling.

A more sophisticated analysis must begin with the mirror metaphor itself. In fact this metaphor permits Gelernter, presumably without intending it, to gloss over numerous features of the actual practice of computing, namely the diverse forms of mediation between representations and reality (Agre 1995b; Bowers 1992; Star 1995). Let us consider a few of these.

ontology A computer can only capture and calculate with what it can represent. A Mirror World can feed video streams into citizens' living rooms and store them in databases, but it will not be able to interpret them in any significant way. The Mirror World's ability to summarize, abstract, or evaluate the world will depend on the ontology embodied in its representations.

standards To have any hope of representing the world in a compact and useful way, a Mirror World will require that the same ontology, or a manageably small set of ontologies, be used to represent human affairs the world over. But this kind of ontological uniformity will require the uniform imposition of standardized categorizations and measurement schemes (Bud-Frierman 1994).

instrumentation A mirror can be installed next to an activity without having much effect on it. But if an activity is to be reflected in a computer database, its participants must somehow be provided with equipment that can capture the necessary data. This instrumentation will normally require that the activity be reorganized in some way, for example by requiring employees to enter certain data after each job. The introduction of new formal methods can lead to subtle and pervasive reorganizations in the equilibrium of work at a given site (Berg 1997).

authentication A Mirror World that claims to report the activities of particular people and things must presuppose some material process to verify their identity. The mirror metaphor postpones this requirement, displacing the need for identification to the end user of the system who observes the activities at a digitally mediated distance.

interpretation Gelernter predicts that Mirror Worlds will eliminate the need for political interest groups by providing everyone with equal access to an overall view of society. This assumes that the necessary syntheses of raw data can be performed automatically. In reality, one major function of an interest group is its interpretation and synthesis of information. The work of a legislative analyst, for example, is nowhere near being automated, although perhaps it can be displaced geographically from the legislature as drafts of legislation become available on line.

selection If a Mirror World is easy to set up, then presumably anybody can set one up anywhere. But as the complexities of mirroring become more evident, the question arises of which data are created in the first place and which data are stored. Gelernter (1991, p. 27) acknowledges that data streams generated by environmental sensors can easily create more data than could possibly be stored, but he assumes that choices about filtering and storage of the data can be made automatically while preserving the utility of the system for users.

bias Citizens using a Mirror World to gain an overview of the real world will depend on a great deal of software. But software can embody biases that are not obvious to its users or even to its authors (Bowers 1988; Friedman and Nissenbaum 1994). Seemingly objective calculations are routinely deployed for strategic purposes (March and Olsen 1989). The mirror metaphor promises an undistorted image of the world, or at worst an anamorphic image if the mirror is warped, but other types of bias are harder to conceptualize within the framework that the metaphor provides.

performance Once Mirror World surveillance becomes pervasive, people will design their activities with an eye to the consequences of surveillance (Dourish 1993; Suchman 1992). The activities will become, in effect, staged performances that project, to the greatest possible extent, the participants' preferred construction of reality.

The mirror metaphor, then, minimizes our sense of representations as material things. Computer representations differ from the reality of human activities in many ways, and the ability of computers to create and maintain representations of human activities presupposes a great deal of prior work (Agre 1994). The language and the practices of database design do not reflect this work, for the simple reason that it is someone else's job. Databases receive data, and it is parsimonious to describe the various data elements in a straightforward fashion as proxies that map a preexisting reality. A database designer might be told that actions or statuses fall into certain discrete categories, for example, without needing to know how that categorization is performed in real life.

The space between representation and reality does sometimes become visible in database designers' disciplinary texts. For example, Simsion's outstanding text on data modeling points out that the designer must often choose whether a given concept is to be understood as an entity or as a relationship, and Simsion warns against imagining that the answer can be read straight from the underlying reality:

> If we think of the real world as being naturally preclassified into entities and relationships, and our job as one of analysis and documentation, then we are in trouble. On the other hand, if we see ourselves as designers who can choose the most useful representation, then this classification into entities and relationships is a legitimate part of our task.
>
> Some in the academic world choose the first option: that there are entities and relationships "out there" that we need to document consistently. When reading the literature, watch out for this assumption—it isn't always explicitly stated. (Simsion 1994, p. 93)

Simsion is referring here to the entity-relationship approach's distinction between entities and records, and to its insistence on giving priority to entities as pre-given features of the business. Nonetheless, he regards the designer's job as a choice among equally objective representation schemes:

> It is helpful to think of E-R modeling as "putting a grid on the world": we are trying to come up with a set of non-overlapping categories so that each fact in our world fits into one category only. Different modelers will choose differently shaped grids to achieve the same purpose. Current business terminology is invariably a powerful influence, but we still have room to select, clarify, and depart from this. (p. 80)

Simsion's text provides a valuable record of database designers' collective experience. As designers run up against a certain pattern of difficulty, it is possible for an author such as Simsion to identify the pattern and to warn future practitioners against the errors that caused it. He warns, for example, that many commonly used business terms are poor names for entities in a data model because these terms have accumulated a variety of context-dependent meanings that ought to be mapped onto several distinct categories of data (p. 64). These warnings outline the contours of the database designer's epistemological standpoint; they reflect the manner in which problems become visible from that particular location in the organizational and disciplinary division of labor.

This effect is most instructive in the case of the designer's choice of primary keys for a database. The primary keys are those columns in a given table (i.e., those attributes of a represented entity) that serve to identify uniquely the entity instance represented by a given row in the table. Historically, the notion that these keys are plural derives from practices such as identifying individual human beings in terms of their names and their dates of birth. Paper forms, for example, commonly ask for an individual's birth date in order to use it as a "tie breaker" in cases when two people have the same name. This practice is deeply entrenched in the profession—not least because, in the normal course of business, queries to databases tend to employ these "natural" keys. (The effect reinforces itself: use of the birth date as an identifier in existing databases requires it to be collected, and then its availability facilitates its use as an identifier in future databases.) Nonetheless, Simsion argues against the practice. Some of his arguments are more aesthetic than practical: it is hardly good engineering, he suggests on pp. 202 and 205, to rely on a probabilistic assurance of having gotten the right record. A more clear-cut problem can arise when the scope of a system is broadened after the database has been designed; a choice of keys for the narrower range of things might not suffice for the broader range (p. 204). Simsion's strongest argument, though, is that foreign keys—that is, keys whose values are assigned by someone other than the database designer—can change. A person might change his or her name, or a birth date might be discovered to be in error. Changing existing online records to reflect the new values will be hard enough, since the primary key may be used to identify the individual in numerous tables. But it may be impossible to update records that have been archived on fixed media. "[The] foreign key maintenance problem is usually the most effective method of convincing programmers and physical database designers of the need for stable primary keys." (p. 206)

The point has significant consequences for privacy. The mirror metaphor, as well as the conceptual and linguistic tendency to conflate records with the entities they represent, makes it seem reasonable to employ a "natural" identifier, such as name and birth date, as the primary key. But, Simsion suggests on p. 202, a cold look at the technical requirements for a primary key provides a compelling case for the use of "surrogate keys"—arbitrary numbers used solely to signify an entity instance's existence and not its other attributes. So long as representations and reality are conflated, it is difficult to comprehend privacy in any terms other than control over access to individually identified data records. Practitioners' experiences, though, have led to a tacit and indirect awareness of the need to distinguish cleanly between representation and reality. Principled application of data modeling methods and surrogate keys, it must be said, is hardly the norm in current industrial practice. Nonetheless, consistent use of surrogate keys to identify data subjects affords other technical options for protecting privacy.

New Privacy Strategies

At the conclusion of their 1996 meeting in Manchester, England, the European Union Data Commisioners released a statement reporting their decision to "work over the next year on Privacy Enhancing Technologies to promote these techniques for reducing the extent to which personally identifiable information is used in information systems, with a view to developing pilot projects which can be examined at their next annual meeting and used as models to promote the further adoption of these techniques in industry, commerce and public administration."

The analyses above help explain why the data commissioners' initiative is possible in theory and why it is exceedingly difficult in practice. Ever since Sweden passed the first data-protection law in 1973, privacy policies have followed a common model grounded in the Fair Information Practices: presupposing the existence of computerized data records that identify individuals, governments have sought to regulate the conditions under which such records are created, used, and propagated (Flaherty 1989; Bennett 1992).

The data commissioners' analysis of privacy-enhancing technologies (PETs) relied upon a study by the data commissioners of Ontario and the Netherlands on technologies of anonymity (Information and Privacy Commissioner and Registratiekamer 1995; see also Burkert's chapter in this volume). This study suggested the use of a system element called an "identity protector" to shield individual identities from organizations holding data records about them. The underlying technology for identity protection was developed in large part by David Chaum (1985, 1986, 1990, 1992; see also Chaum, Fiat, and Naor 1988). Chaum observes that data records cannot cause privacy problems unless they can be traced back to the individuals whose lives they represent. So long as databases identify individuals by a universal identifier, such as a name or a government-issued identity number, records can easily be propagated and merged, and thus they can be employed for secondary purposes to the individual's detriment. Chaum's alternative approach is to employ digital pseudonyms, also called pseudoidentifiers. Individuals wishing to transact business with an organization would not have to identify themselves. Instead, one would present proof that one owned a particular pseudonym, and would then use a cryptographic protocol to present "credentials" (Chaum 1985) that would warrant that some relevant predicate holds (e.g., being old enough to purchase alcoholic drinks) without revealing the information upon which the credential was based (e.g., the individual's precise age). The purpose of the identity protector is to manage and authenticate these pseudonyms, which are then used as surrogate keys in the relevant organizational databases. Figure 6 illustrates one possible configuration of an identity-protecting architecture. In this case, the identity protector is interposed between a system user and a variety of "pseudo-domains," each of which employs a different pseudonym to identify the user. By providing a different pseudonym to each organization, the user can control which personal information moves across organizational boundaries (Chaum 1990). Other configurations are possible, depending on where the sensitive identification is stored, which parts of the system actually need to know the user's identity, and what kind of auditability might be required. For example, the California Air Resources Board's scheme to detect drivers who fail to repair their emissions systems, discussed at the outset, could employ a kind of "identity escrow" that discloses the identity of a driver (or the car) only when fault codes consistently show up over a certain period.

Despite their complexity, schemes based on digital pseudonyms offer certain advantages beyond the protection of individual privacy. Because databases indexed by pseudonyms no longer contain individually identifiable information, they need not be secured as tightly. Information can also be more readily transferred across organizational boundaries for purposes such as statistical research. Authentication poses a problem for these schemes, however, and the most appropriate method of authentication will

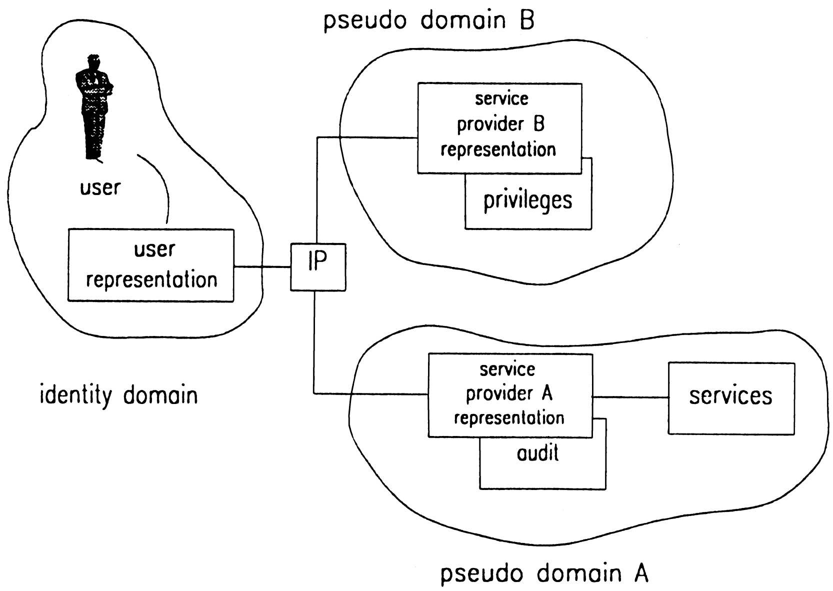

Figure 6
A proposed architecture for identity protection. Source: Privacy-Enhancing Technologies: The Path to Anonymity (Information and Privacy Commissioner and Registratiekamer, 1995), volume 2, p. 19.

depend on the particulars of each application. Photographic or biometric identification of individuals (Clarke 1994) would seem to defeat the whole purpose of an anonymous system. Smart cards carry the risk of loss or theft. Users probably cannot be expected to remember large numbers of distinct passwords. Technologies based on biometric encryption, though, may be useful for pseudonymous authentication in certain applications. In the biometric encryption system marketed by Mytec Technologies of Toronto, an individual initiates a relationship with a given organization by permitting a Mytec device to encode a known data string using an optically transformed fingerprint image. On subsequent occasions, the individual is able to initiate a transaction by placing a finger on a scanning device to decrypt the encoded information. The organization, meanwhile, does not retain an image of the individual's fingerprint or any other individually identifying information.

Encouraging the adoption of such technologies, unfortunately, requires more than technical existence proofs. One significant issue is trust in the system. After all, the system could also be used as a traditional biometric authentication system through the encoding of a universal identifier; an organization that simply claims to be using biometric encryption is therefore not necessarily protecting anybody's privacy. The issue has already arisen in Ontario, where a social activist group, the Guelph Coalition Against the Cuts (1996), has begun a campaign against the provincial government's plan to employ biometric encryption in distributing welfare payments. The Coalition suggests that this "technical proposal for treating the poor like criminals" is a suitable symbol of the conservative government's cuts in social-welfare programs. Fingerscanning systems for welfare recipients have, after all, been implemented in the United States; the issue here is whether the Ontario system based on biometric encryption deserves to be assimilated to the police-state metaphors that have long shaped the public's understanding of universal identification. It is conceivable that any misunderstandings could be repaired through a suitable public communication campaign. One potential model might be the successful campaign that Pacific Bell conducted in the spring of 1996 as a condition of its introduction of Caller ID service to California. Pseudoidentity, unfortunately, is a more technical concept than Caller ID, and the technology itself is counterintuitive even to computer professionals.

Indeed, the force of habit among system designers is another significant obstacle to the adoption of PETs. So long as the language of computer system design tends to blur the difference between representations and reality, treating data records as mirror images of the world, the protection that PETs provide for individual identity will remain nearly incomprehensible—designers will often simply assume that records pertaining to individuals can be traced back to them. The problem does not lie in the practices of database design, which may evolve toward the use of identifers that can equally easily be universal or pseudonymous—in that limited sense, the methods of computer system design do not, as social theorists have often assumed, inherently lead to the invasion of privacy. Instead, privacy invasion results from the way in which the technical methods are customarily applied. The necessary change in customs can be encouraged through model programs and exhortation from policymakers.

Ultimately, though, it will be necessary to revise the training of system designers to integrate the new range of technical options. Lessons will be needed in cryptography and protocols for the management of pseudonyms, of course, but other necessary revisions are less obvious. The examples provided in database texts, for example, almost invariably identify human beings by their names, and they rarely provide any sense of the moral issues that are at stake in the selection of primary keys.

Much experience will have to accumulate before PETs can be integrated into the wide variety of institutions and concrete life situations to which they are applicable in theory (see Bijker and Law 1992). Troubles arise immediately when organizational relationships are not confined to digital media. Customer-service telephone numbers, for example, have typically required individuals to identify themselves, and telephone interactions based on pseudonyms will inevitably be clumsy. Customers who enjoy the record-keeping benefits of periodic statements (from credit cards, banks, the phone company, and now automatic toll-collection systems) will require some way of obtaining such statements without disclosing their names and mailing addresses.

Perhaps the most significant accomplishment of PETs to date is to have initiated a shift of imagination. When records of personal information can always be traced back to the individuals whose lives they represent, privacy interests trade off against system functionality. The result, as the historical record plainly shows, is that privacy interests almost invariably lose out in the end. Data-protection regulation can contain the abuse of personal information within this framework, but it cannot contain the proliferation of technologies that create and use personal information for ever-broader uses. Markets in privacy can operate at the boundaries of the system (for example, to express customer choice about secondary uses of transaction-generated information), but such markets require that customers be able to estimate the disutility of secondary use despite the enormous credit-assignment problem. (The issue is one of transparency: "How did they get my name?") PETs, however, promise to provide data subjects with much more detailed control over the use of their information, and to greatly lower the price of refusing to disclose it. Whether that promise is fulfilled, however, depends on a great diversity of factors.

Conclusion

I hope to have suggested the utility of historical investigation of computing as a representational practice. A brief chapter cannot survey the whole field of computer science or even the whole subdiscipline of systems analysis; numerous episodes in the complex history of data have been glossed over, and relevant areas such as data communications and security have been left out entirely. Nonetheless, some clear trends can be discerned. At stake is the sense in which a technical field has a history: what it inherits from past practice, how this inheritance is handed down, the collective thinking through which the field evolves, and how all this is shaped by the institutional contexts within which the work occurs.

To answer these questions, one might suggest that privacy problems have arisen through the application of industrial methods to nonindustrial spheres of life, where normative relations of representation and control are different. But that statement alone is too simple. Technical work is not indifferent to the contexts in which it is applied; quite the contrary, it is continually confronted by the practical problems that arise in pursuing instrumental goals in particular concrete settings. Another part of the problem, then, is the manner in which "problems" arise in the course of practitioners' work, and how these problems are understood. Database designers, for example, have been forced to clarify their methods on numerous occasions when existing databases have been used for new and unforeseen purposes.

Yet these mechanisms of practical feedback were evidently not adequate to stimulate the invention or the ready adoption of methods to decouple data records from human identity. The mathematicians who did invent such methods in the 1980s were explicitly motivated by a desire to protect privacy, and they have faced an uphill fight in getting their ideas adopted. Much of this fight, of course, has been overtly political (see Phillips's chapter in this book). Organizations with pecuniary interests in the secondary use of personal information will presumably not be enthusiastic about the new technologies. On another level, however, the fight concerns basic understandings of representation and its place in human life. Information is not an industrial material or a mirror of a pre-given reality. It is, quite the contrary, something deeply bound up with the material practices by which people organize their lives. Computer systems design will not escape its association with social control until it cultivates an awareness of these material practices and the values they embody.

Acknowledgements

One version of this paper was presented at the University of Toronto, where I benefited from the comments of Andrew Clement. Another version was presented at AT&T Labs, where Paul Resnick, Joel Reidenberg, and several others offered helpful comments. Geoff Bowker, Michael Chui, Roger Clarke, and Leigh Star kindly offered suggestions on a draft.

References

Agre, Philip E. 1994. Surveillance and capture: Two models of privacy. *Information Society* 10, no. 2: 101–127.

Agre, Philip E. 1995a. Institutional circuitry: Thinking about the forms and uses of information. *Information Technology and Libraries* 14, no. 4: 225–230.

Agre, Philip E. 1995b. From high tech to human tech: Empowerment, measurement and social studies of computing. *Computer Supported Cooperative Work* 3, no. 2: 167–195.

Agre, Philip E. 1997. *Computation and Human Experience*. Cambridge University Press.

Benedikt, Michael, ed. 1991. *Cyberspace: First Steps*. MIT Press.

Beniger, James R. 1986. *The Control Revolution: Technological and Economic Origins of the Information Society*. Harvard University Press.

Bennett, Colin J. 1992. *Regulating Privacy: Data Protection and Public Policy in Europe and the United States*. Cornell University Press.

Berg, Marc. 1997. *Rationalizing Medical Work: Decision-Support Techniques and Medical Practices*. MIT Press.

Bijker, Wiebe E., and John Law, eds. 1992. *Shaping Technology / Building Society: Studies in Sociotechnical Change*. MIT Press.

Borgida, Alexander. 1991. Knowledge representation, semantic modeling: Similarities and differences. In *Entity-Relationship Approach: The Core of Conceptual Modeling*, ed. H. Kangassalo. North-Holland.

Bowers, C. A. 1988. *The Cultural Dimensions of Educational Computing: Understanding the Non-Neutrality of Technology*. Teachers College Press.

Bowers, John. 1992. The politics of formalism. In *Contexts of Computer-Mediated Communication*, ed. M. Lea. Harvester Wheatsheaf.