

Privacy in Sensor Webs and Distributed Information Systems

J.D. Tygar

EECS and SIMS
University of California, Berkeley
USA
tygar@cs.berkeley.edu

Abstract. We are seeing rapid development of sensor webs that collect information and distributed information aggregation systems – and these trends are present both in industry and government. This paper outlines a research agenda for address privacy questions raised by sensors webs and distributed information systems.

1 The Role of Privacy

Privacy has taken on new importance after the tragic events of September 11, 2001. In response to terrorist attacks, governments are preparing systems that

- Anticipate potential threats; and
- Respond to actual threats.

The first item often includes improved intelligence systems; the second item includes systems that accurately report on the results of terrorist attacks. For example, in the US, the Defense Advanced Research Projects Agency has created two new offices: the Information Awareness Office and the Information Exploitation Office [9], to deal with these two phases respectively. The Information Awareness Office is developing a major new system called Total Information Awareness [5,8] to address issues of anticipation. Work going on in several places on wireless sensor webs [4] is partially motivated by issues of response. These systems have much potential in providing increased intelligence and emergency response capabilities, but at the same time, provide the capability for sophisticated and universal surveillance, to a degree previously not seen.

But the contemporary computer security community is ill-prepared to address questions of privacy. Indeed, privacy is usually treated as the poor step-child of computer security. Although a leading IEEE conference in our field is called *The Symposium on Security and Privacy*, even a casual glance through any recent proceedings will reveal that privacy usually receives little attention – and when it does, it is ancillary to a much larger and more important point about computer security. Other leading

conferences, run by the ACM, USENIX, or the Internet Society don't mention privacy at all in their titles, and privacy papers in those conferences are no more likely than in our IEEE conference.

It is more than a little surprising that privacy receives so little attention¹. Such systems are ubiquitous in the private sphere (consider, for example, the stunning amount of personal data collected by the private firm Acxiom.) But although the most serious privacy questions exist in private systems, concerns are also raised in government systems – both existing and proposed. In the United States, for example, a large number of government agencies, ranging from the Internal Revenue Service to the Federal Bureau of Investigation to the Social Security Administration to the Immigration and Nationalization Service today collect large amounts of personal information.

A number of continuing trends (such as the spread of the Internet, the increase in electronic payment methods, near-universal use of cellular phone communications, ubiquitous computation, sensor webs, etc.) mean that private organizations will have increased ability to build detailed profiles of individual Americans. And on the government side, there have been widespread calls in the media in the wake of the September 11, 2001 terrorist attacks for government organizations to share information and “connect the dots.”

Privacy poses significant, challenging problems. These are not merely hypothetical problems: commercial use of data means that these problems are present, important, and pressing today.

This paper presents some preliminary results on the problem of privacy, but its primary focus is to set forth proposals for a research agenda in privacy. Now the problems associated with privacy are well-understood to be difficult. But this paper presents concerns in a true spirit of research – merely stating that a problem is difficult is no reason to avoid studying it. Indeed, in the face of nearly certain increased private commercial and governmental surveillance, the problems of privacy demand attention, regardless of their inherent difficulty.

Privacy is an issue that includes both technical and policy elements, and a full consideration of privacy must necessarily touch on policy issues. However, for this paper, my focus is entirely on the technical aspects of privacy.

If we can make progress on the privacy problem, we will benefit from many powerful spin-off technologies. Better privacy means better handling of sensitive information, which can directly lead to better computer security. If we can deny attackers access to sensitive information, we raise the bar on the ability of attackers to successfully attack systems. This point lies behind much of the concerns raised by the President's Critical Infrastructure Protection Board's September 18, 2002 *Draft National Strategy to Secure Cyberspace* [6].

In addition, we can also hope that this technology will see use in the commercial sphere, as mentioned above.

Finally, in the wake of the terrorist attacks of September 11, 2001, intelligence handling by several United States government agencies, including the Federal Bureau

¹ The excellent paper on *Negotiated Privacy* by Jarecki, Lincoln, and Shmatikov in this anthology is a happy exception to this general rule.

of Investigation, the Immigration and Naturalization Service, and the National Security Agency were widely criticized in the media. Part of this poor handling came from awkward handling of “private” information (that is, information that was about private individuals – if we had better privacy controls, we could hope that critical information could reach the right people while protecting irrelevant personal information from being disclosed, yielding improved intelligence.

2 Privacy Challenge: Distributed Information Systems

In November 2002, numerous newspaper articles in the United States pointed to one example of a proposed distributed information system, the Total Information Awareness (TIA) system proposed by the United States Defense Advanced Research Projects Agency. In many ways, the focus on TIA is unfair: TIA is actually behind what is going on in the commercial sphere. Neither is the debate on the role of government databases unique to the United States. In Canada, Europe, and Japan, similar concerns have arisen – for example, just consider the concern of Japan’s computerized citizen information system *Zyuumin Kihon Daityou*.

But perhaps TIA stands out since privacy has been a central goal of the program since the beginning, and since it has the support of the United States Defense Advanced Research Projects Agency (DARPA), an organization famous for solving apparently unsolvable problems. So let us begin by considering TIA. (To be clear, this paper is not a review of TIA, and does not mean to endorse or criticize TIA. However, the reality is that commercial firms such as Acxiom have substantial data aggregation and surveillance capabilities today. A government system such as TIA is likely – perhaps less ambitious, perhaps more ambitious, perhaps with different sorts of data. But it simply defies political imagination to assume that government will not play a role in using data aggregation techniques that are now ubiquitous in industry. The question raised by this paper is: *Given the contemporary, real development of sensor webs and distributed information systems, what can researchers do to help find ways to protect privacy?*) Doug Dyer of DARPA describes the rationale behind Genisys, a significant component of TIA, as follows [1]:

“In 1995, Aum Shinri Kyo attacked the Tokyo subway system using sarin gas, killing 11 and sending hundreds to the hospital. This attack is a prototypical, perhaps extreme example of predictive precursor events. Prior to the 1995 attack, Aum Shinri Kyo cultists led by Shoko Asahara had tried to buy US and Russian chemical munitions. When that failed, they engaged in a substantial weapons development program aimed at mass effects. They created elaborate development facilities for producing sarin, purchased expensive equipment, and accumulated huge chemical stockpiles. Following several malicious attempts with ineffective agents, they created a test facility in Australia and tested sarin by killing sheep whose remains were later discov-

ered. Noxious chemical leaks were reported by neighbors near their development facility in Japan, but police still made no arrests. Asahara broadcast his intentions clearly via radio. And months before the subway attack, cultists used sarin in an overt attempt to kill three judges in the city of Matsumoto. In this example, just as in the case of 9/11, fragments of information, known by different parties in different organizations, contained a predictive pattern that could have been identified had the information been shared and properly analyzed.”

Among the types of information anticipated to be collected in TIA is transactional information including communication records, financial records, education records, travel records, records of country entry, place or event entry, medical records, veterinary records, transportation records, housing records, transportation records, housing records, critical resource records, and government records. These records may be gathered from private sources or governmental sources. [5]

What techniques can we use to address privacy concerns in distributed information systems? We can clearly leverage from recent technical advances including:

- Powerful strong audit systems, that maintain tamper-resistance audits (for example, see [2]).
- The ability to search on encrypted data without revealing information about the query itself or its response. For example Dawn Song, Adrian Perrig, and David Wagner present powerful techniques that allow information to be searched in data repositories, without revealing the nature of the of the queries or the results either to eavesdroppers or the data repositories itself (see Figure 1). [7] This work, which has since been built on by a number of researchers, offers a surprising result – we can build databases that are secure, both in a conventional sense but in a new stronger distributed sense. This suggests that a number of extensions may be possible: (a) government agencies may be able to use data from private organizations such as Acxiom without revealing the nature of inquires or searches to Acxiom. Since private commercial organizations often have appalling security (example: the widely reported recent “identity theft” of highly sensitive private information using poor security at private credit agencies and their clients), protecting the nature of queries is a central concern for effective government use of private information. (b) To the extent that privacy handling rules can be expressed in computer readable format, it may be possible to enforce privacy restrictions within the framework of an automated processing system. Now, I do not mean to imply that this technology is ready to use “off the shelf” – it does not fully support the functionality listed above and has efficiency issues. But the theoretical success of the Song/Perrig/Wagner approach suggests that we will be able to make real progress on an ideal system that will be efficient and support the above goals.

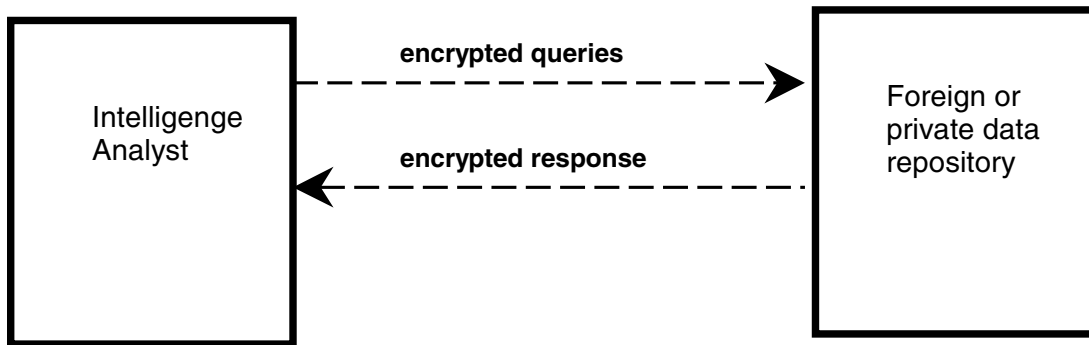


Fig. 1. A schematic representation of the Song/Perrig/Wagner secure query system

- Improved techniques in program analysis, including advances such as Peter Lee and George Necula’s Proof Carrying Code to prove that mobile code has certain properties, static analysis of programs, and projects such as UC Berkeley’s Athena system for automatically synthesizing cryptographic protocols with given properties.
- Stunning advances in dynamic coalitions that can allow diverse, constantly changing sets of parties to cooperate and work together.

2.1 An Outline of Privacy Research Goals for Distributed Information Systems

Here are three techniques that show strong promise for protecting distributed information systems:

- **Selective revelation.** Rather than revealing all data to queries, data can be distributed in discrete chunks that require positive approval to advance. The goal is to minimize revelation of personal data while supporting analysis. The approach revolves around partial, incremental revelation of personal data. Initial revelation is handled by statistics and categories; subsequent revelations occur as justified by earlier results. Such a system is powerful – it can support both *standing queries* (“tell me all of the foreigners from watchlist countries who are enrolled in flight school”) as well as specific real-time queries (“tell me the where the money involved in Transaction X may have come from”). Selective revelation is a method for minimizing exposure of individual information while supporting continuous analysis of all data. The challenge in doing this is that much of the data is private information about people which cannot necessarily be revealed to a person. Even data that is derived via an analysis algorithm may be private, depending on the status of the data from which it was derived, and on the nature of the algorithm.

The idea of selective revelation is that initially we reveal information to the analyst only in sanitized form, that is, in terms of statistics and categories that do not reveal (directly or indirectly) anyone’s private information. If the analyst sees reason for concern he or she can follow up by seeking permission to get more precise information. This permission would be granted if the initial information provides sufficient cause to allow the revelation of more information, under appropriate legal and policy guidelines.

For example, an analyst might issue a query asking whether there is any individual who has recently bought unusual quantities of certain chemicals, and has rented a large truck. The algorithm could respond by saying yes or no, rather than revealing the identity of an individual. The analyst might then take that information to a judge or other appropriate body, seeking permission to learn the individual's name, or other information about the individual. By revealing information iteratively, we prevent the disclosure of private information except when a sufficient showing has been made to justify that revelation.

Selective revelation works by putting a security barrier between the private data and the analyst, and controlling what information can flow across that barrier to the analyst. The analyst injects a query that uses the private data to determine a result, which is a high-level sanitized description of the query result. That result must not leak any private information to the analyst.

Selective revelation must accommodate multiple data sources, all of which lie behind the (conceptual) security barrier. Private information is not made available directly to the analyst, but only through the security barrier.

A key technology challenge is making data relationships successively refinable, and thus selectively revealable. One way to address this challenge is to develop data ontologies that provide structured and logical ways for the analyst to make appropriate queries.

But of course, as the careful reader will have already noted, this creates an architectural problem. It is easiest to think of a protective privacy/security barrier as existing outside a single monolithic repository. However, for the applications listed above, the single monolithic repository will not exist. In the sort of systems necessary to support detailed analysis, information must be collected and aggregated from multiple repositories. In this sort of system, there can be no central privacy/security barrier – each repository must have its own barriers, and those barriers must be coordinated to support privacy restrictions of the system as a whole and of the individual repositories.

- **Strong audit.** Perhaps the strongest protection against abuse of information systems is strong audit mechanisms. We need to “watch the watchers.” These audit systems must be tamper-evident or tamper-resistant, and since repositories span different organizations, must themselves span different organizations. If such audit mechanisms exist, we will realize substantial advantages. (For example, a strong audit mechanism would have been likely to identify a spy such as Aldrich Ames or Jonathan Pollard very early on.) However, these audit systems themselves pose a substantial challenge. Audit data will be voluminous and highly sensitive (certainly, foreign intelligence agents would be very interested in finding out what sorts of queries are run through a government's commercial or governmental information systems.) How can we find instances of inappropriate queries? In many ways, this is a recursive instance of the general intelligence data mining problem, and should probably be considered in conjunction with that problem. This hall of mirrors presents a number of technical challenges, and would benefit from the research community's attention.

- **Better rule proceeding technologies.** Distributed information systems combine data from diverse sources. Their privacy systems must support privacy constraints: both systemic privacy constraints and privacy constraints specific to a particular set of information repositories. (For example, information derived from a foreign source, such as country X's information repositories, may come with specific privacy concerns attached.) Since computers in general can not understand the underlying representation of private information, it is necessary to label data with information that will allow it to be properly processed, both with respect to privacy constraints but also with respect to general constraints. Information varies tremendously in quality as well. For example, substantial anecdotal evidence supports the claim that a significant data appearing in commercial credit bureau sources is not always accurate. Information from foreign sources may be tampered with. Government agencies vary in the degree of scrutiny they apply to keep data accurate. All of this poses issues for accurate labeling. Further concerns arise because even a new information system will likely build on substantial amounts of (unlabeled or inaccurately labeled) previously existing *legacy data*. And problems continue to increase in complexity: what happens when data is combined to produce *derived data*. How should the derived data be labeled? A conservative approach would suggest labeling the derived data with the most conservative labels possible. However, in many cases, this will be inappropriate – derived data is often sanitized and poses less privacy restrictions than the original source data used. On the other hand, in some cases derived data may actually be more sensitive than the original source data.

Data labeling is actually an old idea – it dates back to some of the earliest discussions in the security community. And yet, labeling enjoys at best a checkered reputation: for example, as this author has elsewhere argued, data labeling was a spectacular failure in the attempts to define standards such as the US Department of Defense's Trusted Computer Systems Evaluation Criteria (the "Orange Book.") But there is reason for optimism: when labeling is used to support advisory functions or review functions, one would expect considerably better performance than in the classic mandatory security systems that formed the heart of higher levels of the Orange Book. Indeed, labeling is going through something of a renaissance at present: consider its use in a variety of digital rights management systems (DRMs) (although, of course, these DRM systems are now widely recognized to have only partial effectiveness.)

If labeling can take place, substantial advantages can be realized. For example, consider the complexity of privacy law in the United States. Different laws apply to protecting video rentals and airline tickets, audio recordings and video recordings, information from foreign sources and domestic sources, information concerning US persons and foreigners, etc.

This poses risks in multiple directions. On the one hand, there is the risk that current complexity of US privacy laws and rules may result in inappropriate disclosure of information. But perhaps, as some former US Justice Department members have argued, in many cases, intelligence analysts and law enforcement personnel miss the opportunity to use essential intelligence information: in their

desire to comply with privacy rules, the government officials fail to use material that they are legally entitled to use. In other words, the haze of privacy law makes officials go beyond legislative and regulatory privacy restrictions and means that the government misses the chance to “connect the dots.”

Clearly, we have a significant challenge in allowing users of databases (whether employees of companies such as Acxiom or law enforcement officials or intelligence analysts) reasonably understand what the real privacy restrictions are. Here is a place where technology can help – if we can develop a “privacy toolbar” that helps inform users of privacy restrictions, we can help eliminate mistakes caused by human misunderstanding of the United State’s currently complex privacy law. This especially applies when rules interact. If a privacy restriction is reached, such a system can help explain the procedures necessary to legally access data.

Moving towards more advanced and speculative research, we envision a system which can simulate different information handling policies. Such a system would use synthetic data and run queries against them. By comparing different privacy policies, we aspire to find examples that will help illustrate respective advantages and disadvantages of a variety of privacy policies. This could help inform debate by policy makers as they consider different privacy policies. (This stands in marked contrast to contemporary approaches to privacy policy making, which is often marked by political rhetoric and vague sociological characterizations.) However such a simulator faces substantial challenges: we need to design the simulator itself, we must find a way to generate meaningful synthetic data, and we must find ways to verify or validate the accuracy of reports from the simulator. These are all hard problems, and their solution is far from obvious. This is an example of “high-risk” (the challenges are real), “high-payoff” (even partial solutions can help shed considerable light on policy making) research.

But within the framework outlined here, considerable questions remain. Consider the problem of adaptation. As people realize that certain data is subject to surveillance or response, they change their behavior. Here is an example familiar to any frequent flier in the United States: prior to the terrorist attacks of September 11, 2001, many experienced airline passengers angled to be among the very first in line to board commercial airplanes – in that way, the passengers could place their carry-on luggage and settle in before the rest of passengers boarded. In the wake of the September 11, 2001 terrorist attacks, authorities instituted thorough searches of some passengers (and their carry-on luggage) boarding flights. While these checks are ostensibly random, they in fact tend to select more highly from the first people to board a flight. Now, experienced travelers angle to be the tenth in line instead of the first in line.

In the same way, information systems designed to identify certain groups of people are likely to result in different behavior from both the people they are intended to track as well as innocent people who for personal reasons desire to evade surveillance (for example, observe the “arms race” in the United States between telephone caller-ID systems, those who desire to make anonymous calls, those who desire to reject anonymous calls, etc.) Failure to correctly anticipate

this sort of adaptation is likely to lead to unexpected and (often undesirable) results. In the worst case, this pattern of adaptation could lead to widespread evasion of surveillance, followed by a counter-move of analysts who need to dig deeper into private data, leading to a spiral resulting in markedly decreased privacy in practice. To some degree, simulation as mentioned above may help prevent this. But beyond that, it makes sense to attempt to user large scale real-user experiments. One way to accomplish this may be to adapt commercial multiplayer games (along the lines of *Everquest* or *Baldur's Gate*) to see how real users (albeit, a highly self-selected group) adapt to specific policies.

3 Privacy Challenge: Sensor Webs

Sensor webs use small, independent sensors that communicate wirelessly over micro-cells to collect and share a wide variety of information about their environments, including people in the environment. As part of the CITRIS project at Berkeley, the author and his colleagues at Berkeley have been developing and deploying such sensors. Among the explicit purposes of such devices are to report on people in the building and the status of the building in case of a disaster such as an earthquake, fire, or terrorist attack on a building. These reports are explicitly designed to report on the position and status of people in a building. This raises obvious privacy concerns.

Today, the sensors used are obvious – while small, they are still visible to the naked eye. However, with successful further development, some aspire to develop truly dust-sized sensors that can even be mixed in a paint can painted onto the wall. Using ambient power sources (such as temperature differentials or small movements), these dust-sized sensors could function indefinitely. Forming an ad hoc network on the fly, these sensors could pass along a variety of information, eventually linking that information up with a base station.

In [4], the author of this paper and Adrian Perrig describe how to arrange sensors into a secure network. We describe the configuration that we used for our actual system:

“By design, these sensors are inexpensive, low-power devices. As a result, they have limited computational and communication resources. The sensors form a self-organizing wireless network and form a multihop routing topology. Typical applications may periodically transmit sensor readings for processing. Our current prototype consists of nodes, small battery powered devices, that communicate with a more powerful base station, which in turn is connected to an outside network. . . . At 4MHz, they are slow and underpowered (the CPU has good support for bit and byte level I/O operations, but lacks support for many arithmetic and some logic operations). They are only 8-bit processors (note that . . . 80% of all microprocessors shipped in 2000 were 4 bit or 8 bit devices. Communication is slow at 10 kilobits per second. The operating system is particularly interesting for these devices. . . . [TinyOS, a] small, event-driven operating system consumes almost half of the 8K bytes of instruction flash memory, leaving just 4500 bytes for security and

the application. Significantly more powerful devices would also consume significantly more power. The energy source on our devices is a small battery, so we are stuck with relatively limited computational devices. Wireless communication is the most energy-consuming function performed by these devices, so we need to minimize communications overhead. The limited energy supplies create tensions for security: on the one hand, security needs to limit its consumption of processor power; on the other hand, limited power supply limits the lifetime of keys (battery replacement is designed to reinitialize devices and zero out keys.)”

Given the severe limitations of such an environment, the technical reader is likely to doubt that security is even possible. However, in the book, we go on to describe in great detail a set of algorithms which do secure these systems, and we go on to describe an actual implementation of those algorithms on these very sensors.

3.1 An Outline of Privacy Architecture for Sensor Webs

Privacy on sensor webs can be considered in three levels:

- **Fundamental security primitives:** Examples include low level encryption, authentication, secure broadcast, and synchronization protocols. These are challenging because of the limited computational power of sensor webs, but as outlined in [4], substantial progress has been made on these topics, and this particular level seems particularly tractable.
- **Privacy policies²:** We need a way to specify what information is recorded, made available and distributed: to whom and under what conditions. A large variety of privacy policies can be considered – here are some examples: “location information of individuals is not recorded or made available except in case of an emergency; and such an emergency must be accompanied by a loud alarm”, “users can specify whether they want their location information revealed or not”, “location information is made available, but the identities of the persons are anonymized.” Not only must these policies be specified, but we must provide a formal, machine understandable language to interpret these policies.
- **Human interfaces:** Humans interact with sensor webs in at least two ways: as subjects who are potentially under surveillance by the sensor web and as users who may need to access information collected or synthesized by the sensor web. In both cases, we need good ways to provide user interfaces. For subjects, we have the questions of how they receive notice (“you are under surveillance in the following ways”) and how they specify preferences (“please allow people in my work group to know my location except when I am on my lunch break.”) For users of sensor web, we need good ways to specify queries, to receive easily understandable answers to queries, and to specify and monitor policies.

² Note that here we are using policy in its technical sense as a set of rules, rather than in its more general sense as an instance of “public policy”.

Once these techniques are in place, we can conduct a variety of important experiments and analyses:

- **Experimentation on Policies:** In addition to examining our human interfaces to determine their clarity and effectiveness, we can use this system as a sociological testbed to see individual's comfort with different types of monitoring policies. Do people gradually acclimate to monitoring policies? Or do monitoring policies have an observable change on behavior and work habits? Is notice effective? Do people understand what is being collected, and what consequences it has?
- **Security Analyses:** To what extent are privacy safeguards designed into the sensor web vulnerable to attack? Are there single points of failure which if breached present a serious risk of large scale release of private information?
- **Masking Strategies:** How effective in practice are anonymizing masking strategies such as adding noise to data or presenting data in a purely statistical form?
- **Legal Analyses:** What is effectively done with technology, and what gaps remain? Where do we need additional regulation to protect individual privacy? What sort of legal regulation is well supported by technology and what laws are not? Are there laws which are effectively unenforceable because technology can not support those laws?

A fully completed analysis along these lines could yield a substantially deeper understanding of privacy issues in sensor webs.

Acknowledgements. I began thinking about this material as I chaired the US Department of Defense Information Science and Technology Study Group on Security with Privacy. While the members of that study chose to remain anonymous, the author is deeply grateful to them for many detailed discussions. Portions of section 2.1 of this paper was adapted from my commentary to that study. I would also like to thank the staff of DARPA, and the Information Awareness Office, for answering many questions with openness.

Material in Section 3 grew out of discussions I have had with Tom Kalil, Yutaka Miyake, Deirdre Mulligan, Adrian Perrig, and Paul Wright.

My thinking about this subject was stimulated by my students in my Fall 2002 Graduate Computer Science Seminar on Privacy.

I was supported during the writing of this study from grants from the United States National Science Foundation and the United States Postal Service.

The Tokuda Lab at Keio University, the Institute for International Policy Studies, and the National Science Council of Taiwan (through the International Computer Symposium) helped provide partial travel support to present this material in Japan and Taiwan. I am particularly grateful to Mr. Isao Hiroki, Professor Shiuhyng Shieh, and Professor Hide Tokuda, and for facilitating this Asian support.

While I benefited from a broad variety of advice, input and financial support for this study, the opinions in the report are my own. They do not necessarily reflect the opinions of any person or organization mentioned above, any funding agency, or of the US Government or any of its agencies.

References

1. D. Dyer. *Genisys*. In [8]
2. S. Haber and W. Stornetta. "How to Timestamp a Digital Document." *Journal of Cryptology*, Vol. 3, No. 2, pp. 99–111, 1991
3. G. Mack, B. Bebee, I. Shafti, G. Wenzel, B. Medary, E. Yuan. *Total Information Awareness Program (TIA) System Description Document (SDD), Version 1.1*. July 19, 2002. In [8].
4. A. Perrig and J. D. Tygar, *Secure Broadcast Communication: in Wired and Wireless Communications*, Kluwer, 2002.
5. J. Poindexter. *DARPA's Initiative on Countering Terrorism: Total Information Awareness*. In [8].
6. President's National Critical Infrastructure Board. *Draft National Strategy to Secure Cyberspace: Draft for Comment*. September 18, 2002. <http://www.whitehouse.gov/pcipb/>
7. D. Song, D. Wagner, and A. Perrig. "Practical Techniques for Search on Encrypted Data." In *Proceedings 2000 IEEE Symposium on Security and Privacy*.
8. *Total Information Awareness* (CD-ROM) distributed at DARPATech 2002, July 2002.
9. See URLs <http://www.darpa.mil/iao> and <http://www.darpa.mil/ixo>.