Computer Science and Artificial Intelligence Laboratory

Technical Report

MIT-CSAIL-TR-2007-034 June 13, 2007

Information Accountability

Daniel J. Weitzner, Harold Abelson, Tim Berners-Lee, Joan Feigenbaum, James Hendler, and Gerald Jay Sussman

massachusetts institute of technology, cambridge, ma 02139 usa — www.csail.mit.edu

Information Accountability

[1]Daniel J. Weitzner, [1]Harold Abelson, [1]Tim Berners-Lee, [2]Joan Feigenbaum, [3]James Hendler, [1]Gerald Jay Sussman

[1]Massachussetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory, Cambridge, MA USA
[2]Yale University, Department of Computer Science, New Haven, CT USA
[3]Rensselaer Polytechnic Institute, Troy, NY USA

djweitzner@csail.mit.edu, hal@mit.edu, timbl@csail.mit.edu, joan.feigenbaum@yale.edu, hendler@cs.rpi.edu, gjs@mit.edu

I. Introduction

Ease of information flow is both the boon and the bane of large-scale, decentralized systems like the World Wide Web. For all the benefits and opportunities brought by the information revolution, with that same revolution have come the challenges of inappropriate use. Sensitive personal data disclosed, corporate secrets revealed, copyrighted material distributed without permission, confidential records shared among organizations in violation of regulation and policy -- these breaches of established social norms and laws have become part of everyday life in the Information Age. Such excesses and abuses in the use of information are most commonly viewed through the lens of information security. They are seen as consequences of unauthorized access, the result of information "escaping" beyond an appropriate boundary. Accordingly, enormous effort in current information technology research and development is devoted to inventing more reliable methods for restricting access to information.

These efforts notwithstanding, access restriction alone is inadequate for addressing information misuse on the Internet. An exclusive reliance on access restriction leads to technology and policy strategies in which information, once revealed, is completely uncontrolled. It's like focusing all one's attention on closing the barn door and ignoring what might happen to the horses after they've escaped. The reality is that even when information is widely available, society has interests in whether or not that information is used appropriately. Information policies should reflect those interests, and information technology should support those policies.

Even when access restriction can be perfectly and completely achieved, there are significant cases where policies implemented purely as *ex ante* (up front) controls are too rigid to faithfully reflect societal needs. One example is copyright control and the need to accommodate fair use. Another might be in determining whether law-enforcement agencies can permissibly share information about targets of investigations. Traditionally, law and policy address such complexity through enforcement mechanisms where control is imperfect and exceptions are possible, but where violators can be identified and held accountable. Information technology infrastructure should do likewise.

This paper argues that debates over online privacy, copyright, and information policy questions have been overly dominated by the access restriction perspective. We propose an alternative to the "hide it or lose it" approach that currently characterizes policy compliance on the Web. Our alternative is to design systems that are oriented toward *information accountability* and *appropriate use*, rather than *information security* and *access restriction*. In a world where information is ever more easily copied and aggregated, and where automated correlations and inferences across multiple databases can uncover information even when it has not been explicitly revealed, accountability must become a primary means by which society addresses issues of appropriate use.

Our goal is to extend the Web architecture to support transparency and accountability. When information has been used, it should to possible to determine what happened, and to pinpoint use that is inappropriate. This requires augmenting Web information with data about provenance and usage policies, and creating automated means for maintaining that provenance and interpreting policies. Transparency and accountability can be supported by a set of technical mechanisms we call *Policy Awareness*. Policy Awareness is a property of information systems that provides all participants with accessible and understandable views of the policies associated with information resources, provides machine-readable representations of policies in order to facilitate compliance with stated rules, and enables accountability when rules are intentionally or accidentally broken. Our understanding of the dilemmas of information policy in the age of the Web is built upon the work of numerous legal academics[1] and writers[2]. Our proposed public policy and systems architecture framework is an effort integrate these insights into a comprehensive framework of law and technology that, while not immediately providing answers to all legal

or technical design questions, sets out a direction in which we are likely see the Web and other large-scale systems evolve toward greater accountability.

Transparency and accountability make bad acts visible to all concerned. This visibility alone does not guarantee compliance. Then again, the vast majority of legal and social rules that form the fabric of our societies are not enforced perfectly or automatically, yet somehow most of us follow most of the rules most of the time. We do so because social systems built up over thousands of years encourage us to do so and often make compliance easier than violation. Social rules are known to us in a transparent manner, though often in the tacit dimension of our environment[3]. For those comparatively rare cases where rules are broken, we are all aware that we may be held accountable through a process that looks back through the records of our actions and assesses these actions against the rules. Augmenting systems with Policy Awareness is an attempt to bring to the Web and other information architectures at least to the level of transparency and accountability that we have in other arenas where human interaction is governed by social rules.

II. Challenges to the Information Hiding Approach to Policy Compliance

Personal privacy, copyright protection, and government data mining are examples of the sorts of areas that present significant policy challenges for the information society. It has been more than a decade since stable policy regimes in these areas have been upset by the rise of network computing and the World Wide Web. Yet society still seems far from robust technical or public-policy solutions to these problems. In each of these cases, we believe that excessive reliance on secrecy and up-front control over the flow of information has resulted in policies that fail to meet social needs and technologies that stifle information flow without actually resolving the problems for which they are designed.

A. Privacy

Information privacy rights, at their core, seek to safeguard individual autonomy as against the power that institutions or individuals may gain over others through use of sensitive, personal information[4]. Policy makers worry that sensitive, and possibly inaccurate, information will be used against people in financial, employment or healthcare settings. Democratic societies also worry that citizens' behavior will be unduly restrained if they fear they are being watched at every turn. They may not read controversial material, or they may feel inhibited from associating with certain communities for fear of adverse consequences.

Society has typically sought to safeguard information privacy rights by limiting the collection of, or access to, personal information[5]. The privacy of email messages or telephone calls is protected by prohibiting access to and disclosure of the stored message data or interception of the audio signal except under very limited circumstances. Personal bank balances are kept secret from all except authorized bank personnel. Medical records are protected from public view. In countries with comprehensive data protection laws (including the EU, Canada, Hong Kong, and Australia) no personal information can be collected about individuals at all without their affirmative consent.

Today, however, with the proliferation of personal information on the World Wide Web and the increased analytic power available to everyone through facilities such as Google, preventing access to or collection of a given piece of data is often without real privacy benefit.[6] Worse, many privacy protections, such as the lengthy privacy policy statements we receive online or in health or financial services contexts, are mere fig leaves over the increasing transparency of our social and commercial interactions. Either the same information is available elsewhere on the Web in public, or it is possible to infer private details to a very high degree of accuracy from other information that itself is public.[7] What's more, in the case of publicly available personal information, people often make this data about themselves available intentionally, not just by accident. They may not intend that it be used for every conceivable purpose, but they are willing for it to be public nonetheless.

With personal information so widely available, privacy protection through information hiding and access control is largely ineffective. As a case in point, the growth of electronic commerce on the Web over the second half of the 1990's sparked a concern about consumer privacy that led to an emphasis on Web site privacy policies, and infrastructure like the World Wide Web Consortium's Platform for Privacy Preferences (P3P). Today, most consumer Web sites post privacy policies, and popular browsers make use of P3P for the purpose of restricting the placement of cookies that violate self-regulatory norms agreed to by the online advertising industry. However, P3P did not succeed as effective privacy management tool beyond the control of cookie usage.

Much has changed about the Web since P3P was designed over ten years ago. The design target for P3P in the mid 1990s was the largely bilateral relationship between an individual consumer and a Web site. At most a third party advertiser was also involved. Today far more complex multi-party exchanges of information are the norm. P3P does not today suffice as a comprehensive technological support for privacy on the Web because of growing tensions in the "notice and choice" model of Web privacy. A fully-implemented P3P environment could give Web users the ability to make privacy choices about every single request to collection information about them. However, the number, frequency and specificity of those choices would be overwhelming especially if the choices

must consider all possible future uses by the data collector and third parties. Consumers should not have to agree in advance to complex policies with unpredictable outcomes. Instead, they should be confident that there will be redress if they are harmed by improper use of the information they provide, and otherwise they should not have to think about this at all.

Consider this scenario: Alice has a three-year old child with a severe chronic illness that requires long-term expensive treatment. Alice tries to learn all that she can about the illness, buying books about the disease online, searching on the Web using AOL, and participating on online support parent support chat rooms. One day Alice applies for job and is rejected. She suspects it's because a background check somehow learned about her Web activities and flagged her as high risk for expensive family health costs.

Stories like this are often put forth as cautionary tales to support the need for Web privacy policies. Did the book store assert that the titles of Alice's purchases would be kept confidential? Did AOL promise never to release information about Alice's online searches? Did the chat service guard against lurkers in the chat room recording the names of participants? A policy regime based on data hiding would focus on these potential acts of data release, perhaps even taking the position that it should be Alice's responsibility to inform herself about these Web sites' privacy policies before using the services. That focus is misplaced: the *actual* harm was caused not by the disclosure of information by the bookseller or AOL or the chat service, but by the decision to deny Alice the job, i.e., by the inappropriate use of that information. In fact, it is quite conceivable that Alice actually wants to be identified as someone who has an interest in this particular illness. Forcing her to hide in order to protect herself against improper information usage significantly limits her ability to exercise her right to freedom of association. Instead, we should instead be able to live in online environments which provide transparency of information usage and accountability to rules that limit harmful uses of personal information. In sections III and IV we discuss the legal and technical framework to enable accountability to privacy rules.

B. Copyright

Copyright control for information on the Internet poses challenges similar to those of privacy. It presents similar opportunities for information policy to take better advantage of transparency and accountability, thereby coming to rely less on secrecy and access restriction. As with privacy, the increased flow of copyrighted information brings enormous benefits, but there are also associated risks – here the risk that copyright holders will lose control over their works altogether and suffer massive infringement.

In the copyright context, information hiding commonly takes the form of digital rights management (DRM) systems, where information flows through the Web in encrypted form, and decryption keys that unlock the information are granted pursuant to licenses that users negotiate. As with privacy, it has proved extremely difficult to keep information reliably locked up, and efforts at up-front control over the information flow results in user frustration and substantially imperfect security. The recording industries have long believed that music sold online must have copy protection in place in order to ward off the possibility if large-scale illegal copying of their digital assets. Yet it has become clear that illegal file-sharing on the Internet won't be stopped by inconveniencing honest consumers. There are enough sophisticated copyright infringers who continue to populate illegal Internet services with stolen content regardless of the technical limits music publishers may erect. Recently, Apple's CEO Steve Jobs wrote[8] that DRM has not worked and is never likely to. Soon afterwards, Apple changed the way it sells music online by offering a more expensive version unencumbered by technical copy protection measures that have typically been present in online music sales. There is one catch, however, to Apple's elimination of DRM. These unlocked tracks have embedded inside them the name and other personally-identifying information of the purchaser. That way, if he or she shares the purchased music with their hundred million closest friends on MySpace or elsewhere, there is a chance the purchaser can be held accountable.

Experience with online copyright has shown, in a pattern also similar to the privacy dilemma, that attempting to realize social policy goals through access control mechanisms can become extraordinarily complex. With copyright, it's been common for access controls to go beyond the statutory rights reserved to copyright holders and impose additional restrictions. The notable example of in the US and many other jurisdictions is the doctrine of Fair Use. Fair uses, by their very nature, are unauthorized by the copyright holder, and involve the determination by a court, after the fact, of whether what would otherwise have been copyright infringement was in fact permissible.

Some legal scholars maintain that fair use requires that users should not need to request permission from copyright holders in advance, or must get access even when permission has been requested and refused, requirements that seem fundamentally inconsistent with access control.[9] Several proposals have been advanced for surmounting this inconsistency, including taking big steps towards accountability architectures through systems that permit users to override access controls, in which case the access is logged, presumably with something like the accountability architecture we describe below.[10] Switching from access control architectures to accountability architectures does little to reduce the complexity of fair use

determinations, but it at least provides a framework where fair use arguments can be addressed.

One of the great promises of the Web is that it can be much more than just a distribution network for consumer content, but rather, a vast arena for collaboration, where creators can build on each other's work. The appropriate infrastructure to support this goal is a architecture not of access control but of accountability. One example is provided by the Creative Commons organization, which promotes a family of copyright licenses that encourage the flow of information but retain certain rights for the work's creator.[11] For example, some licenses might proscribe commercial redistribution of licensed works; some licenses might permit distribution but prohibit making derivatives; some licenses might permit making and redistributing derivatives, but only under the terms of the original license.

Creative Commons has also developed a Rights Expression Language that recasts these licensing terms in machine readable form, thus enabling licensed works to participate in accountability architectures. For example, there is a browser plug-in that identifies when a visited page is under a Creative Commons license, and alerts the user with information about the licensing terms.[12] This presents the technical challenge of extending the architecture to provide support for accountability when content under different licenses is mixed, a task that is just beginning to be explored.[13]

The Creative Commons approach does not seek perfect enforcement of the licenses up front, as DRM does. In fact, the Creative Commons organization takes the position that its licensing terms are incompatible with DRM systems, since the terms require redistribution of licensed works in a way that preserves fair use. Rather, the Creative Commons architecture recognizes the value of having information flow around the Internet quite freely but still seeks to impose certain restrictions on how the information can be used. While the ultimate enforcement of these conditions flows from the force of law, Creative Commons also illustrates how technical tools, such as those the make licenses visible and allow for machine computation of the effect of overlapping licenses, can come together to make it easier for users to comply with complex rules in a dynamic information space.

Aside from meeting a particular set of open access goals, this approach to copyright rules has is actually much closer to the operation of copyright law in the offline world (books, movies, etc.) than is the DRM-controlled world that some publishers have turned to on the Internet. It is a system of online information accountability that does a far better job of reflecting the traditional relationship between legal rules and human behavior than do more restrictive copyright technologies. The force of law behind the Creative Commons licenses is no less forceful than other online copyright regimes, but it takes an approach that is better tuned to the fluid nature of information on the Web.

C. Surveillance and Data Mining

Recent government uses of advanced data mining techniques are yet another example of the deficiency of access control and collection limitation approaches to privacy compliance on the Web. Data mining holds out the promise of being an important new component of terrorism detection and prevention, and perhaps even criminal investigation, but raises at the same time a whole new category of privacy challenges.[14] The power of data mining technology lies in its potential to bring to light non-obvious investigation targets or identify terrorist threats through inferences drawn on decentralized data sets spread around the Web, *i.e.* around the world. Traditional public policy tools appear inadequate to the task of protecting basic privacy rights in this situation. The vast majority of data through which government agencies are mining is either available for access with minimal legal process, or, in many cases, is already public. Laws that limit access to information will not protect privacy here because so much of the data is publicly available. Taken item by item there is little legal basis for controlling access to this information, but the intrusive power comes from the aggregation of large numbers of data sources. To date, the law has not developed any way to address this privacy loophole.

Current technical investigations of the impact of data mining on privacy have generally focused on limiting access to data at the point of collection or storage. Much effort has been put into the application of cryptographic and statistical techniques to construct finely tuned access-limiting mechanisms.[15] Yet for all this emphasis on access restriction, the reality is that the Web is making it increasingly difficult to limit access to data, while at the same time making it increasingly easy to aggregate data from multiple information sources, and to do searching and inferencing based on these aggregations.

The case of airline passenger screening activities by law enforcement and national security agencies, for example, illustrate the growing complexity of information handling and transfer. Society may be prepared to accept and even expect national security agencies to use very aggressive data mining techniques over a wide range of information in order to identify potential terrorism risks. However, we consider it unacceptable to use the same information with the same powerful analytic tools in order to investigate domestic criminal activity. Therefore, we need rules that address permissible uses of certain classes of information, in addition to simply access and collection limitations. Our initial research in this area[16] has shown that transparency of data access and tools which provide accountability to clearly-stated usage-oriented rules can be a viable alternative to failed access control techniques.

III. Information Accountability as a Legal Framework

The information accountability framework has the advantage that it more closely mirrors the long-standing relationship between law and human behavior than do the various efforts at forcing policy compliance through access control over information. Upon reflection we realize that the rule of law in society operates largely in the background of our activity. Overall, compliance with the law is very high in democratic societies, yet in the vast majority of cases there are no mechanisms that *force* us to comply with legal or social rules.[17] Rather, we comply because rules are *generally known* and social institutions tend to *make compliance easier than violation*. As a final backstop, if we are tempted to break rules, we are aware that there is a rich web of institutional and social forces that make it more likely than not that our rule breaking will eventually be discovered.

There is a significant paradox associated with protecting privacy or other information policy values through increased transparency. Must we collect *more* information, much of which may be personal and sensitive, in order to protection privacy? In many cases it is only by making better use of the information that is collected, and by retaining what is necessary to hold data users responsible for policy compliance that we can actually achieve greater information accountability. We should, of course, to the maximum extent possible protect the data that is stored for accountability purposes. Creating new points of failure in an already insecure Internet environment is not an attractive strategy, but in most cases the data needed to ensure accountability is already collected somewhere. Our goal is to be sure that this data is organized, annotated and available to be used for better policy compliance and accountability.

A. Models for regulating large-scale information systems

Transparency and accountability lie at the root of many legal regimes generally regarded as successful. In particular, two legal systems responsible for regulating large-scale information environments in the United States – the Fair Credit Reporting Act and the Securities Act – illustrate that it is possible to achieve substantial control over how information is used without the tight, upfront control sought by policy/technology designs inspired by the traditional computer security model.

Fair Credit Reporting Act

One of the earliest large-scale uses of personal information in the private sector is the consumer credit reporting system. This system generates credit risk assessments for use in lending and employment decisions. From early on, policymakers recognized the need to protect individual privacy, assure accuracy of the data, all while maintaining enough flexibility in the operation of the system to allow a thorough analysis of consumer credit data based on the maximum amount of useful information possible.

The Fair Credit Reporting Act[18], the statute that regulates the credit bureaus, their users and data sources, has balanced these three regulatory goals by constructing the rules such that the credit bureaus are permitted to collect very wide range of information. Privacy is protected not by limiting collection of data, but rather by placing strict rules on how the data may be used. Analysis for the purpose of developing a credit score is essentially unconstrained, but the resulting information can be used only for credit or employment purposes. It cannot, for example, be used for marketing or profiling. Strict penalties are imposed if these use limitations are breached. Data quality is assured by giving all consumers the right to see the data held about them (transparency). And, if a user of the data makes a decision adverse to the consumer (denial of a loan or rejection of an employment application) that decision must be justified with reference to the specific data in the credit report upon which the decision was based (accountability). If the consumer discovers that this data is inaccurate, s/he may demand that it be corrected. Again, there are stiff financial penalties imposed against the credit bureau if the bureau fails to make the appropriate corrections.

Taken together, these provisions of the fair credit statute are an example of an information regulation regime that achieves is goals by setting *usage* rules, not *access* or collection rules. It relies on transparency and accountability, as opposed to any effort at up-front enforcement of the rules. There is no guarantee that all data in the credit system is accurate or even that it will be used according to the rules. However, there is a robust feedback loop built in to correct errors and strict penalties when rule violations are discovered.

Some may be surprised by the idea that the consumer credit bureaus should be used as a model for any future regulation. It is the case that people come into contact with this system often when there is a problem with its operation. However, given the scale of the system, comprising data about most adults in the United States, collecting data from hundreds of millions of sources and providing analysis to hundreds of thousands of users, and the high stakes decisions the system is responsible for supporting, this system is an important model for regulation of large scale, decentralized information systems.

Securities law

A large part of the commercial economy of the United States rests on a foundation of transparency and accountability. Four times per year tens of thousands of public companies are required to file detailed information about their financial status. The Securities and Exchange

Commission sets rules specifying exactly what information must be disclosed and maintains copies of all of these documents for public examination. Were we not so accustomed to this system, one might characterize the transparency requirements as nothing less than radical: disclosure is required for everything from sensitive financial data, executive compensation details, and current threats to the company's economic prospects.

The only up-front enforcement that goes on is the requirement that the report itself be filed. The SEC generally does not look at the content or check the accuracy of the filings at all. Instead, accountability to substantive rules and duties of financial management is achieved because shareholders, their lawyers, or in very rare cases government prosecutors will examine SEC filings of companies whose financial position is in doubt for some reason. If a company's filing is either inaccurate or incomplete, there are then serious penalties. As with the Fair Credit Reporting Act, this is a regulatory system that relies on thorough transparency and after-the-fact accountability for behavior that is discovered to be illegal.

IV. Technical Architectures for End-to-End Information Accountability

What technical architecture is required to support information accountability? We have demonstrated in the case of privacy and copyright that compliance and accountability challenges arise from the extreme free flow of information. Thus we can no longer reliably achieve policy goals by merely limiting access to information at one point on the Web. Exercising control at one point in a large information space ignores the very real possibility that the same data is either available or inferable from somewhere else. Thus, we have to engineer Policy Aware systems based on design principles suitably robust for Web-scale information environments. Here we can learn from the design principles that enabled the Internet and the Web to function in a globally-coordinated fashion without having to rely on a single point of control.

The need for information accountability illustrates in the policy compliance arena what Internet architects learned about communications protocols as they successfully deployed highly decentralized network services at a scale not previously achievable with traditional, centralized network design. The shift in emphasis from access control to accountability in data privacy is analogous to the shift from careful step-by-step relay protocols to end-to-end[19] protocols in communications networks. In these networks responsibility for correct behavior belongs at the end points of the communication system, rather than at the intermediate relay points. Effort to ensure correct behavior at the relay points, beyond that needed to obtain performance, is usually wasted. Much of credit for the robust behavior of the Internet is due to this design choice for the protocols. The analogous idea here is that the purpose of privacy is to protect holders of personal information from adverse consequences of the misuse of that information. Efforts to ensure perfect access control at each step of an information transfer may not be as effective as making sure that adverse consequences of misuse can be mitigated.

Applying the lessons of end-to-end design to the goal of building policy aware systems, we must then ask how we can build accountability features into every node of the Web at which information is used, and how will these highly distributed accountability functions will come to be integrated into system-wide accountability.

A. An initial Policy Aware architecture

Information accountability on the Web will emerge from the development of three basic capabilities: policy-aware audit logging, a policy language framework, and accountability reasoning tools.

Policy Aware transaction logs: In a decentralized system each endpoint will have to assume the responsibility of recording information usage events that may be relevant to current or future assessment of accountability to some set of policies. These logs will become the basis of assessing policy accountability either in real time or at some point in the future when such an assessment is needed. A policy-aware transaction log will initially resemble traditional network and database transaction logs, but also include data provenance, annotations about how the information was used, and what rules are known to be associated with that information. Policy-aware logging, not unlike TCP/IP packet routing is a relatively dumb operation that is oblivious to the semantics of the subject of log entries. What matters is that events are logged. Other tools will be responsible for analysis and action based on the content of the logs. A number of fundamental questions must be answered about logs, however: what information should be kept and what discarded? How will the logs themselves be secured? How will integrity be assured?

Policy Language Framework: Assessing policy compliance over a set of transactions logged at a heterogeneous set of Web endpoints by a diversity of human actors requires some common framework for describing policy rules and restrictions with respect to the information being used. We consider it improbable in the extreme that the entire world would ever agree on a single set of policy language primitives. However, drawing on semantic web techniques including ontologies and rules languages, we believe it will be possible for larger and larger overlapping communities on the Web to develop a shared policy vocabulary in a step-by-step, bottom-up fashion. Perfect global interoperability of these policies is unlikely but that is not a fatal flaw. Just as human societies learn to cope with overlapping and sometimes

contradictory rules, especially when jurisdictional boundaries are crossed, so too will policy aware systems be able to develop at least partial interoperability. Partial decidability is also fact in today's legal systems and it will be so online as well.

Policy Reasoning Tools: Accountable systems must assist users in seeking answers to questions such as: Is this piece of data allowed to be used for a given purpose? Is a string of inferences permissible for use in a given context, depending on the provenance of the data and the applicable rules. It seems likely that special purpose reasoners, based on specializations of general logic frameworks, will be needed to provide a scalable and open policy reasoner.[20] An initial application of these reasoners has been implementation of policy aware access control[21] that enable standard web servers to enforce ruled-based access control policies specifying constraints and permissions with respect to the semantics of the information in the controlled resources and elsewhere on the Web.

B. Implementation and Deployment Strategy

A critical consideration in making the Web policy-aware is whether this would require significant re-architecting of basic Internet and Web protocols, or whether transparency and accountability can be layered on top of existing infrastructure.[22] One possible approach to doing the latter is through collection of *accountability appliances*[23] that are distributed throughout the Web and communicate using Web-based protocols, both among themselves and with other Web resources. The accountability appliances would serve as proxies to data sources, where they would mediate access to the data and maintain provenance information and logs of data transfers, and they would present client-facing services to browsers, for presenting accountability reasoning in human readable ways, and allow annotation, editing, and publishing of the data and reasoning presented.[24]

Accountability appliances must be able to reason about the restrictions on the use of information, and how those restrictions must be combined when corresponding pieces of information are combined. In addition, the appliances will need to interact, because assessing policy compliance sometimes entails reasoning about long chains of data transfers after the fact. These interactions present substantial engineering challenges: individual appliances should be constrained so that each operates mostly independently of the others, otherwise the interactions between the appliances will cancel out the advantages of distribution.

We are still exploring that challenges associated with reasoning over heterogeneous policy expressions. If the language of data restrictions is unlimited in expressive power then the problem of tracking and combining restrictions will be computationally infeasible. But there are useful policy subsets whose rules can be implemented in a straightforward way. One way to model this is through algebraic expressions involving properties of the data, the organizations, and the trail of data transfers, which can be formalized as a *Data-Purpose Algebra* to guide automated inference.[25]

Even once we solve the reasoning challenges of policy aware systems, practical questions remain regarding the deployment of these technologies on the Web and in enterprise systems. The only scalable deployment path is an incremental one. Some enterprises (government agencies, for example) and certain online communities (such as scholars who use Creative Commons) could be early adopters of policy aware systems because of their pressing need to manage their own data in an accountable manner. It is not necessary for the entire world to adopt this approach right away but as these communities depend on data from others, the users for whom accountability is a priority will encourage or require other data producers to support policy aware systems. Experience developing the World Wide Web suggests that it will be essential to provide easy-to-use, royalty-free technical standards that can be implemented widely and with minimal burden. Our design efforts are taking into account the need for a step-by-step deployment path.

V. Conclusion: Beyond information hiding

In 1967 Alan Westin published his landmark study *Privacy and Freedom."*[26] Still in the age of mainframe computers and relatively small numbers of data bases, it set the stage for thinking about privacy over the next three decades. In the book, Westin presents what has now become a classic definition of privacy:

> "Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others."

Westin's work remains essential today for its identification of the role of privacy in a free society. However, advances in communication and information technology and the ease of data searching and aggregation have rendered this definition incomplete as a framework for information policy and information architectures that seek to be policy aware.

In its place, information accountability through policy awareness, while a departure from traditional computer and network systems policy techniques, is actually far more consistent with the way that legal rules traditionally work in democratic societies. Computer systems depart from the norm of social systems in that they seek to enforce rules up front by precluding any possibility of violation, generally through the application of strong cryptographic techniques. In contrast, the vast majority or rules we live by have high levels of actual compliance without actually forcing people to comply. Rather we follow rules because we are aware of

what they are and because we know there will be consequences, after the fact, if we violate them, and we trust in our democratic systems to keep rules and enforcement reasonable and fair. We believe that technology will better support freedom by relying on these social compacts rather than by seeking to supplant them.

Acknowledgements

Work conducted at the MIT and RPI with support from National Science Foundation grants: the Transparent Accountable Data Mining Initiative (award #0524481) and Policy Aware Web project (award #0427275). Weitzner thanks UMD Mindlab for supporting this work and for providing an office away from the office in which an important part of this work was launched.

Notes

[1] Daniel Solove described limitations of the "secrecy paradigm" in his 2004 book, The Digital Person. As we will do here, he stresses the importance of accountability.

[2] David Brin wrote the Transparent Society (1998), proposing that the only solution is to maximize transparency for all.

[3] The Tacit Dimension. Michael Polanyi. Doubleday & Company, 1966.

[4] There are numerous definitions of privacy. For the purpose of this paper, our chief interest is in understanding privacy rights as related to collection and use of personal information, as opposed to other privacy protections that seek to preserve control over one's bodily integrity, for example.

[5] For important exceptions to this pattern, see the discussion of post hoc legal regimes in section IV.

[6] Consider publicly accessible data such as property records and voting rolls, and the fact that individuals are putting more and more data about themselves on the Web through blogs, social networking sites, and information sharing sites for photos and videos.

[7] L. Sweeney. *k*-anonymity: a model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10 (5), 2002; 557-570

[8] Steve Jobs, Thoughts on Music (6 February 2007) http://www.apple.com/hotnews/thoughtsonmusic/

[9] Erickson and Mulligan elaborate on the challenges for use poses for DRM in J. S. Erickson and D. K. Mulligan, "The Technical and Legal Dangers of Code-Based Fair Use Enforcement," 92 *PROC. IEEE* 92(6): 985-996 (2004).

[10] See T. K. Armstrong, "Digital Rights Management and the Process of Fair Use," *Harvard Journal of Law & Technology,* 20(1), (Fall 2006).

[11] At http://creativecommons.org/ (visited May 23, 2007).

[12] See the description of mozCC at <mozcc.mozdev.org/> (visited May 23, 2007).

[13] See Harvey C. Jones, *XHTML Documents with Inline, Policy-Aware Provenance*, MIT Dept. of Electrical Engineering and Computer Science Masters Thesis, June 2007.

[14] Protecting America's Freedom in the Information Age. Markle Foundation, 2002. http://www.markle.org/downloadable_assets/nstf_full.pdf. Creating a Trusted Network for Homeland Security: Second Report of the Markle Foundation Task Force, 2003. http://www.markle.org/downloadable_assets/nstf_report2_full_report.pdf

[15] For example Agrawal D. and Srikant, R. Privacy preserving data mining, Proc 2000 ACM SIGMOD Conference on Management of Data, 2000, 439-450

[16] Weitzner, Abelson, Berners-Lee, Hanson, Hendler, Kagal, McGuinness, Sussman, Waterman, Transparent Accountable Data Mining: New Strategies for Privacy Protection,; MIT CSAIL Technical Report MIT-CSAIL-TR-2006-007.

[17] R. Ellickson, Order Without Law (1991)

[18] Fair Credit Reporting Act, 15 U.S.C. § 1681

[19] Jerome H. Saltzer, David P. Reed, and David D. Clark; "End-to-end arguments in system design;" in ACM Transactions on Computer Systems, vol.2 no.4 (November 1984) pp. 277--288.

[20] Berners-Lee, Connolly, Kagal, Scharf, Hender, N3Logic : A Logic For the Web (2006) http://dig.csail.mit.edu/2006/Papers/TPLP/n3logic-tplp.pdf

[21] Weitzner, Hendler, Berners-Lee, Connolly, Creating the Policy-Aware Web: Discretionary, Rules-based Access for the World Wide Web In Elena Ferrari and Bhavani Thuraisingham, editors, Web and Information Security. IOS Press, 2005. http://www.w3.org/2004/09/Policy-Aware-Web-acl.pdf

[22] Steven M. Bellovin, David D. Clark, Adrian Perrig, Dawn Song (Eds), "Report of NSF Workshop on A Clean-Slate Design for the Next-Generation Secure Internet," GENI Design Document 05-05, July 2005

[23] Note early proposals for a 'privacy appliance.' Lunt, Teresa (2003). "Protecting privacy in terrorist tracking applications." Presentation to the Department of Defense Technology and Privacy Advisory Committee, September 29, 2003. www.sainc.com/tapac/library/Sept29/LuntPresentation.pdf

[24] This aspect of the accountability and transparency perspective is very closely related to the issue of maintaining provenance for scientific data. See M. Szomszor and L. Moreau, "Recording and Reasoning over Data Provenance in Web and Grid Services," in *Proceedings of the International Conference on Ontologies, Databases and Applications of Semantics (ODBASE'03)* 2888, pp. 603-620, Catania, Sicily, Italy (2003) and Golbeck, G. and Hendler, J. A semantic web approach to the provenance challenge, Concurrency and Computation: Practice and Experience, 2007.

[25] Hanson, Kagal, Sussman, Berners-Lee, Weitzner, "Data-Purpose Algebra: Modeling Data Usage Policies", IEEE Workshop on Policy for Distributed Systems and Networks 2006 (POLICY 2006).

[26] Alan Westin, *Privacy and Freedom,* New York, Atheneum Press, 1967. p.7.