The Sisterhood of the Traveling Packets

Matt Bishop
Dept. of Computer Science
University of California at Davis
Davis, CA 95616-8562
+1 (530) 752-8060
bishop@cs.ucdavis.edu

Carrie Gates
CA Labs, Inc.
1 CA Plaza
Islandia, NY 11749-7000
+1 (631) 935-2007
carrie.gates@ca.com

Jeffrey Hunker
Jeffrey Hunker Associates LLC
401 Shady Avenue, Suite B408
Pittsburgh, PA 15206-4493
+1 (412) 661-0106
hunker@jeffreyhunker.com

ABSTRACT
From a cyber-security perspective, attribution is considered to be the ability to determine the originating location for an attack. However, should such an attribution system be developed and deployed, it would provide attribution for all traffic, not just attack traffic. This has several implications for both the senders and receivers of traffic, as well as the intervening organizations, Internet service providers and nation-states. In this paper we examine the requirements for an attribution system, identifying all of the actors, their potential interests, and the resulting policies they might therefore have. We provide a general framework that represents the attribution problem, and outline the technical and policy requirements for a solution. We discuss the inevitable policy conflicts due to the social, legal and cultural issues that would surround such a system.

Categories and Subject Descriptors
K.4.1 [**Computers and Society**]: Public Policy Issues – *abuse and crime involving computers, ethics, privacy, regulation, transborder data flow*; K.5.2 [**Legal Aspects of Computing**]: Governmental Issues – *censorship, regulation*; K.6.0 [**Management of Computing and Information Systems**]: General – *economics*; K.6.5 [**Management of Computing and Information Systems**]: Security and Protection – *authentication.*

General Terms
Management, Design, Economics, Security, Standardization

Keywords
attribution, authentication, economics, security, traceback, trust

1. INTRODUCTION
In this paper we investigate the requirements behind an attribution system, defining first the general characteristics of "attribution" (Section 1). We then provide a more complete definition of attribution, discussing how this differs from the approaches typically used in the literature. From this, we design a framework for an attribution system (Section 2) that details three key components—the actors, the attribution vector and the policy negotiation system—and describes the system requirements. We then discuss the interesting political and social aspects of the framework. First is the feasibility of developing a full attribution system by analogy to other real-world developments (Section 3), then the economic incentives for creating such a system (Section 4) and finally we posit a likely path towards creating such a system (Section 5). We consider the limitations of attribution (Section 6) and provide some related work (Section 7). We conclude with thoughts on future directions (Section 8).

2. WHAT IS ATTRIBUTION
The Merriam-Webster dictionary defines "attribution" as [4]:

1. the act of attributing; especially: the ascribing of a work (as of literature or art) to a particular author or artist
2. an ascribed quality, character, or right

In general, attribution is desired to hold one (an individual, an organization, a nation) accountable for one's actions. Attribution has been a desired feature of networks (and, indeed, of data at rest as well, although we focus here on data in motion) for some time.

With respect to cyber-security, attribution has been defined as "determining the identity or location of an attacker or an attacker's intermediary" [5]. Defenders such as security professionals and governments have traditionally defined requirements for an attribution system, and made an underlying assumption that attribution in all cases is both necessary and good. Typically a combination of IP traceback schemes are used to determine the actual IP address from which a packet was generated—a scheme that was originally designed to determine the originating IP address for spoofed packets in a denial-of-service attack—and public key infrastructures (PKI) bind a particular individual to a particular message. There has been much work in IP traceback (see, for example, Savage et al. [2] and Burch and Cheswick [1] for early work in this area) and stepping stone detection (see Staniford-Chen and Heberlein [3] for early work in this area).

Within the academic literature, the term attribution (as well as accountability) is used without being defined. The literature generally assumes the Merriam-Webster definition, ascribing the attack to a person or, more commonly, to the originating computer.

Several characteristics shape how we think about attribution frameworks:

- Under some circumstances perfect non-attribution may be desirable, for example to whistleblowers or to websites that will not want to provide the identities of individuals visiting their site even if compelled by subpoena. These highlight the political and cultural aspects of attribution, because some cultures exalt the whistleblower, whereas other cultures condemn him or her. Further, some situations require false attribution, such as some forms of whistleblowing, or an intelligence agent being undercover and surfing to a terrorist website that requires attribution from its visitors.
- The target of attribution may differ depending on the need of the stakeholder. In some cases it might be necessary to attribute a message to a particular individual, while in other cases, only to a specific computer, IP address, or organization. For example, arresting an individual for participating in illegal activities requires binding the individual to the activity. If a nation state has been attacked, it needs to attribute the activity to another state, and not necessarily to the specific individuals who launched the attack.
- Given the possible stakes inherent in the use of an attribution system, the system must provide some indication of the degree of confidence that the attribution is accurate and correct. Returning to the case of the individual to be arrested for illegal activity, the attribution mechanism must provide sufficient evidence and rigor to validate the attribution beyond a reasonable doubt, the standard for a criminal conviction (at least in the United States). It is not sufficient to simply provide the attribution; the attribution must be one in which the user can have confidence.
- The logic of desired attribution is in a sense circular: the degree of attribution considered "adequate" depends on the purpose of the attribution, which includes the types of responses and actions that might be taken based on the degree of attribution. The required level of attribution necessary for a military response to a nation believed to be launching a distributed denial-of-service (DDoS) attack will likely be different from those required by an individual attempting to determine the legitimacy of a financial offer just received over the Internet. For the DDoS attack, if the attribution is not adequate to justify military action then the range of responses available to the target will be circumscribed. Alternatively, both senders and recipients may want or need absolute non-attribution (e.g., political dissidents operating under repressive regimes).
- Different senders and receivers may require different attribution policies. For example, a government web site might require attribution to the user level, but be willing to negotiate down to just an IP address should the user prefer to not provide personal identity. Conversely, a dissident web site needs to advertise its policy of not accepting any forms of attribution before a visitor accidentally provides some (correct) attribution information. We therefore need to determine what policies might be required, as well as the requirements for a negotiation system. Mechanisms for advertising policies also need to be devised, along with an examination of where policies should reside and be enforced—at the end points, intermediate routers, or somewhere else.

Throughout this paper, we assume that attribution is of interest to the stakeholders (who will be identified later). The nature of the attribution that each stakeholder desires may vary, but all are interested in either attributing something to the data, or in preventing the attribution of something to data.

3. AN ATTRIBUTION FRAMEWORK

We define "attribution" as the association of data (called a *characteristic*) with an entity (person, process, file, other data). For example, authentication is a mechanism for attributing an identity to an entity, and is thus an example of attribution. The time at which the data was sent is an attribute of interest in situations with temporal constraints. The route data takes over the network is an attribute that network administrators may find useful to know. This data also may imply how visible the data was as it goes from its source to its destination. Broadcast-style routes enable many more sites to see the data than do point-to-point routes.

The goal of attribution is to show that the characteristic associated with an entity has a particular value, or one of a particular set of values. The purpose for using attribution is generally that of accountability—in a cybersecurity context, it is generally used to determine who is initiating an attack (e.g., Wheeler and Larsen [5]) and therefore assumed to be good and desirable.

This view of attribution is overly limited. The side effect of providing attribution for an attack is that attribution must also be provided for non-attack traffic[1]. Furthermore, our concept of attribution includes interests other than that of the recipient; it encompasses the interests of senders, network perspectives, and other (possibly secondary) requirements. Our definition of attribution, therefore, is much broader in concept, involving multiple parties with multiple intentions, spanning geographic, cultural, social, legal, and national interests. To the best of our knowledge, no one has determined all of the requirements for an attribution system, including who the interested parties are, and what their requirements and incentives are.

Four aspects of attribution are relevant to our discussion.

First is the set of actors. We identify at least nine different entities that have an interest in attribution with respect to a message:

1. The sender of the message;
2. The organization associated with the sender;
3. The governments[2] of the country of the sender;
4. The ISPs over which the message transits;
5. The network backbone providers over whose backbones the message transits;
6. The governments of any intermediate nations through which the message transits;
7. The governments of the country of the recipient;
8. The organization associated with the recipient; and
9. The recipient.

[1] Unless someone is able to implement the "evil bit" defined in RFC 3514 [13] successfully.

[2] For example, in the United States, the state and federal governments are different. County, municipality, and other political subdivisions may also have their own interests.

Multiple parties shape if, and how, we can show that the characteristics associated with an entity have a particular value. Each of these entities has a distinct and different set of interests in attribution; understanding what attribution really means rests on understanding what these interests are, and under what circumstances the varying interests of different parties can *and cannot* be reconciled.

The second is *what* is being attributed. The interests and capabilities of the different parties define the characteristics of interest for attribution. We represent this by an *attribute vector* that lists the characteristics for which the values are requested, or lists the pairs of characteristics and their values.

The third aspect is *assurance*, namely the level of certainty associated with showing that the characteristic associated with an entity has a particular value, or one of a particular set of values. We refer to this as the *attribute assurance.* This level of certainty may vary. Authentication for a web site shows identity much less rigorously than does authentication for a passport.

The fourth component is a policy negotiation system that the actors use to negotiate an acceptable level of attribute assurance, or to determine that no such level is possible under the extant circumstances. Both the desired and achieved levels of attribution depend on choices made by many different parties involved in the creation, transmission, and receipt of a message. Senders, receivers, and other parties such as network operators but also governments, law enforcement agencies, and companies may have distinctly different desires in what they consider useful attribution; these desires may change depending on the circumstances.

Figure 1 provides an overview of our attribution framework.

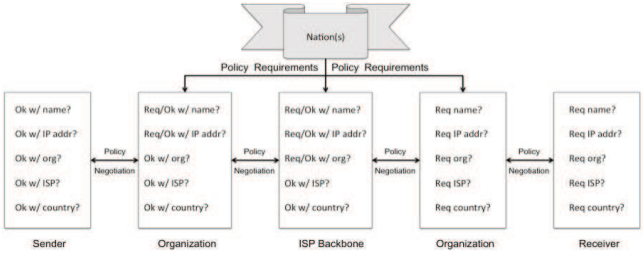

Figure 1. A General Attribution Framework. The attribute vector elements are represented by the characteristics in the boxes. The level of assurance is inherent in the negotiations.

In this section, we describe our framework in detail. The first three subsections discuss each of the above components of the framework. The final one presents system requirements necessary to support the framework.

3.1 Actors

Suppose the Cow Grain Company wants to become the supplier for UC Davis' student cafeteria. The two negotiate a contract over the Internet. The final exchange involves a signed contract, sent from and signed by the UC Davis Dining Director, then received, signed, and returned by the Cow Grain Company President. Correct attribution of both signers is critical because for business purposes, both the senders and the receivers must be certain their peer is the party who may legally commit the peer's company or institution.

In another scenario, a group of attackers launch a DDoS attack on a company that does all its business over the web. When the flooding begins, the company needs to have the flooding packets attributed to the originators of the attack. The originators of the attack, on the other hand, do not want those packets attributed to them. Here, the senders want non-attribution, but the recipients want attribution.

A third scenario has intelligence agents examining terrorist web sites. The web sites want to know who is looking at them, both to get information about potential recruits and to know if adversaries (e.g., intelligence agents) are examining the sites for information about potential attacks. Here, the senders (the terrorist web sites) want full attribution; the recipients (the intelligence agents) want their traffic non-attributable (not merely unattributed).

Finally consider two dissidents in a country with a repressive government who wish to communicate. As neither fully trusts the other, and both believe that the government may be monitoring the messages, neither wants attribution of any kind. Here, both the sender and the receiver want no attribution.

These four scenarios present cases where attribution requirements differ. The recipients may want:

- Perfect non-attribution, in which attribution is not possible; for example, the dissident scenario;
- Perfect attribution, in which the attributes of both the sender and recipient are known to both; for example, the business scenario shown above;
- Perfect selective attribution, in which the recipient wants the attributes known to some entities but not to others; for example, a recipient may care that her spouse knows she received a payment, but not her employer;
- Sender non-attribution, in which the recipient does not want to be able to attribute data to the sender; for example, a whistleblower such as "Deep Throat" in the Watergate scandal;
- Recipient non-attribution, in which the recipient wants to attribute data to a sender but does not want the sender to be able to attribute anything to the recipient; for example, the intelligence agent scenario; and
- Unconcern, in which the recipient does not care about attribution

Similarly, the senders may want:

- Perfect non-attribution; for example, a whistleblower;
- False attribution, in which the recipient can determine attributes of the message but the data, while consistent, is inaccurate; for example, the intelligence agent scenario above, with the agent wanting the terrorists to attribute her messages to an ally of the terrorists;
- Randomized false attribution, or false attribution without the consistency; for example, the intelligence agent scenario in which the agent repeatedly visits the web site, each time under a different identity; and
- Imprecise attribution, in which the recipient can eventually attribute data accurately and to the precision needed, but to do so takes too long (so the knowledge is useless or redundant) or costs more than the value of knowing the attribution.

As noted above, other stakeholders participate in determining type and level of attribution. The ISPs and backbones over which the messages travel may, or may not, add or delete attribution

information. For example, if the originating host's IP address is assigned using the NAT protocol, the firewall doing the NATing effectively eliminates the ability to attribute host origin behind the firewall. But the ISP can attribute IP origin to a subnet, here the one with the firewall connected to the ISP. In order to attribute further, the firewall would need to keep a time-stamped log of internal address assignments, and the ISP would need to record the time each packet left the firewall.

This highlights a central issue for ISPs and backbones to provide attribution. What is the financial cost? ISPs may want to provide attribution services only if they are profitable and the ISP is unlikely to be sued. This balance of profitability and liability is central to the business judgment about whether to provide any service.

Included in the liability issue are cultural and legal constraints. For example, privacy rules in the European Union are considerably more restrictive than those in the United States, so an ISP in the former would not be permitted (or be unable) to provide the attributions that the latter could provide. In some cases, this may be a choice. Anonymous remailers are a good example. Cypherpunk type I remailers provide limited non-attribution because keeping a list of the pseudonyms and senders enables one to derive attribution data. But a Cypherpunk type II remailer prevents this by using sophisticated cryptographic and traffic routing and fragmentation techniques.

Organizations are a different matter. As noted earlier, the organizations of interest are the sending organization, the receiving organization, their governments, and the governments of the countries through which the message transits. A message being sent from the United States to Russia over a network that transits North Korea may have questionable attribution information added. Thus, the attribution data from intermediate nodes, or attribution data that relies on intermediate nodes, is affected by the organizations controlling those nodes.

This brings us to two broad categories of communications.

3.1.1 Cooperating senders and receivers

Senders and receivers that co-operate provide attribution capabilities. Consider the case where both sender and recipient agree on a desired level of attribution, as well as specifically to the party to which the attribution applies. The simplest situation is where the sender organization and government are in agreement with this desired level of attribution.

This agreement requires carefully defined and commonly accepted attribution attributes, and a mechanism for negotiation among all of the parties to ensure agreement on the attributes to be communicated. So it is in all parties' interests to have a robust system to ensure the agreed upon level of attribution.

Backbones and intermediate nodes, however, have no generic incentive for co-operation. Thus, cooperating senders and receivers have to specify some attributes of the network path (for example, no packets can go through North Korea) to enhance or ensure the required attribution.

Cooperating entities with similar needs create new capabilities: mechanisms for either agreeing in advance on the desired level of attribution and the services needed to support the agreed upon level of attribution, or in having an efficient negotiating system. Furthermore, ideally there would be metrics for the trust placed in backbones and intermediate nodes. A policy based path routing would also be necessary to ensure the paths provided the appropriate support for attribution.

3.1.2 Conflicting senders and receivers

Senders and receivers with conflicting attribution needs create choices that either, or both, must make. Political dissidents in repressive regimes provide a scenario that contrasts with that for co-operating senders and receivers. The senders may not (and probably will not) want attribution; whether the recipients would agree to having their receipt of particular packets attributed back to the sender is less clear.

This is a situation in which sending governments (and possibly organizations) want attribution of the sender for repressive political reasons. Recipients, or the international community at large, will probably not want senders to have their messages attributed to them, though this prospect raises the concern that bogus or falsified messages are passed off as legitimate to the recipients.

Furthermore, without the cooperation of sending governments and organizations, creating a policy based routing system will depend on the technical specifications that establishes the policy based trust network, and the extent to which the trust network can in fact be trusted. In particular, the ability to embed the type of attribution and the attribution data in the transmission (at either the message or infrastructure level in a form that can be both assured and reported to the end points) will require an ontology of relevant attribution characteristics, cryptographic or other mechanisms to ensure the integrity and binding of the attribution information to the message.

The cryptographic mechanisms to do this exist. The definitions of relevant characteristics and the ontology of values do not. Nor can the current infrastructure embed this information into the packets; this would require a change, or extension, to the definition of some fields of IP packets. With current technology, the best approach would be to use techniques from IP traceback (discussed below) that embed routing information in the packets, and then use something like a reputation engine to assess trust, and apply that to attribution. Such a solution is unsatisfactory, for obvious reasons.

In this scenario, multiple choices exist. Politically dissident senders may simply choose not to use the Internet. Recipients may be less trusting of traffic without sender attribution—for example, how do recipients know that such traffic is not really government-sponsored disinformation? Intermediate nodes and backbones may cooperate with the sending government for reasons of their own, thus making the policy based trust network less reliable.

3.2 Attribution Vector

An attribution vector consists of a sequence of pairs. The first element of each pair is a characteristic for which a value is either present or desired. The second element is the value of the characteristic, if the vector is providing attribute information. If not, the second element is the distinguished symbol $\perp$, meaning that the value of this characteristic is either requested or not available.

Various types of characteristics will recur when attribution is requested. In practical terms, probably the most common characteristic will be the origin or source of a message to a person or organization. Here, "source" may mean originating user or IP address, the time at which the message was sent, and how the

message was protected in transit (for example, by encryption or access control bits), all of which are easy to provide (assuming all parties are trusted). It may also mean information beyond the current infrastructure's ability to supply, such as who originated (as opposed to sent) the information that was sent, the route that the message took (which gives information about who has access to read or alter it), and where geographically did the message travel (which may bear on delays or the appropriateness of the mechanism chosen to protect the message).

Underpinning the values in the attribution vector is the level of assurance of the values. Values supplied by untrusted sources are less credible than values supplied by trusted sources. For example, asking an ISP for assurances that a government intelligence agency did not read messages transiting that ISP would produce assurances of little meaning if the ISP were known to share its data with the government regularly. As already noted, the degree of confidence in attribution depends on its intended use, and possibly on the source of the values that are attributed.

The origination of the attribution is thus important. Typically one thinks of attribution as relating a packet back to an originating machine. That may be insufficient, and in fact misleading and meaningless. As an example, consider DDoS attacks launched by 'botnets'. Here, attribution back to the botnet provides little insight into the real source of the attack. Attribution may reside with the machine, the organization, and the human being. Attribution may also reside with the network. As an example, at least one major ISP (Rogers, a major ISP in Canada) has confirmed it inserts advertisements into packets responding to certain addresses [14]. These advertisements can only be attributed to the intermediate ISP or network.

Two other considerations affect how the attribute values are handled. The first is to whom the information is reported. Attribution is traditionally thought of as in the ability to determine, based on the interest of the recipient, where the message came from. But how is attribution handled in instances where (for example) one's spouse is an acceptable attribution recipient, but one's employer is not? More generally, one can consider attribution information as being reported to: 1) the recipient; 2) some central authority (e.g., a government or a set of governments) or 3) other intermediate nodes who, either for their own purposes or to forward on the information, find it of value to know what traffic is occurring between two different locations.

The last consideration is the characteristic of why the message was sent. Perhaps this is the most challenging information to attribute, but in many situations imaginable, it will be the most important aspect of attribution. An adequate answer, however, remains an open research question, especially because of the need to examine human motivations. Those are notoriously hard to determine by skilled investigators, let alone by an automated system.

3.3 Policy Negotiation

With nine different classes of actors potentially involved in the attribution, typically a policy negotiation will be required in order to establish an agreed upon attribution vector. Such an agreed upon attribution vector is *a policy contract*. In some cases the negotiations will not succeed; in others, the policy contract will achieve a semi-permanent basis. One can think of policy contract negotiations as a continuum: at one extreme is the oriental bazaar, where everything is constantly negotiated; the other extreme, religious canon, which changes only very slowly if at all. Which structure will predominate we cannot predict; however a policy contract negotiation system should first and always be workable and agreeable to all parties. Given this snap shot of the different goals and needs of the different parties with a stake in attribution, having defined who all of the players are and their needs, a full attribution system needs to have several features:

- *A common nomenclature of attribution vectors* (policy contract elements) provides a precise and mutually understood structure including a common language that each involved party can use to define the desired attribution state. The desired attribution state might include the length of the agreement, specified trust levels among network parties (particularly ISPs and backbones), and penalties for non-performance.
- *A system for communicating and negotiating the policy contract* among the different parties should be transparent, low cost, and made routine to the extent possible. No system that requires a complex legalistic structure to be negotiated in anything but a few rare cases will work for a commonly accepted attribution structure.
- *The ability to specify and communicate desired attribution states and levels of assurance* enables the parties to inform one another in advance of what they require the values of specific attribute characteristics to be in order to accept or reject messages, or continue or terminate policy contract negotiation. At a minimum the senders must be able to specify a level of attribution and the receivers must be able to communicate what levels of attribution it finds acceptable. For example a sender may require that messages not be attributable to its source; the receiver may require full attribution to the source
- *A verification system for ensuring that contracts are performed* will ensure that the entire policy contract negotiation mechanism is enforceable. The verification mechanism needs to provide consequences for those who follow, and fail to follow, negotiated contracts. For example, it might publicly note those who honor policy contracts and those who do not, by using a reputation-based system; or, it may impose a punishment system for violating agreed upon policy contracts, up to and including ostracizing those who breach them.

Policy negotiations themselves cannot violate existing policies. For example, a sender may already have as its policy that its identity never be attributable. Whether a negotiation can succeed under *existing* policies is a question of some import, especially because those policies may not be known when the negotiation starts. One possibility to ameliorate this is to provide a trusted storage mechanism for existing policies, which specify the framework for any further negotiations, or identifies specific types of policy negotiations that may take place between either wholly or partially anonymous parties.

As an example of the complexity of policy negotiations, a government web site might require attribution to the user level, but be willing to negotiate down to just an IP address should the user prefer not to provide his personal identity. Conversely, a dissident web site needs to advertise its policy of not accepting any forms of attribution *before* a prospective user accidentally provides it. We therefore need to determine what policies might be required and how they might be made known to other participants. Mechanisms for advertising policies need to be

devised, along with an examination of where policies will reside and be enforced (for example, in addition to policies at the end points, intermediate routers may also have policies that all transiting traffic must honor). An example from a different application would be the "negotiation" that takes place between a recipient with a telephone blocking calls that suppress the caller ID, and a caller whose telephone does not transmit the caller ID (clearly requiring some other mechanism to initiate communication, or simply the sender determining that communication is not possible). Finally, the mechanisms must be available to non-participants who wish to join the circle of negotiation in order to communicate with entities that require policy contracts.

In many cases, one party may act as a representative for a class of parties to determine a generic policy contract. This is akin to "class action lawsuits," in which a set of actors with a common interest authorizes one actor to negotiate on their behalf. In this case, the policy negotiation mechanisms must enable the binding of all parties, not just the negotiator, to the contract.

This leads to some specific system constraints that support policy negotiation.

First, there needs to be a trust network that enables actors to trust that other actors, and the network, will honor their commitments as negotiated in the policy contract. Networks cannot tag or alter packets of their own accord[3]; some entity must allow them to do so. Thus, signers of a policy contract must have some measure of trust in the other actors to provide attribute values, and to provide *acceptably accurate* values. This trust system might be tied to the verification system mentioned above, and function much as a reputation system would.

Next, a policy-based routing mechanism is needed to ensure that messages traverse networks and midpoints with appropriate attribution mechanisms and levels of trust. This is particularly important if messages are to be routed dynamically (as in today's Internet). Unless the actors do not care whether the attribution changes in transit, or the intermediate nodes cannot alter the attribute vector and do not add any attribute data of their own, the path that the message takes affects both the values in the attribute vector and the level of assurance of that vector (including the values).

This brings up the "superuser" or "Administrator" issue, in which one privileged user can override normal user controls. Traditionally, this mechanism is used to provide an escape to correct severe problems or failures. In high assurance systems, this omnipotent role is partitioned into a set of less powerful roles. A potential concern is defining the role of central authorities for overriding the policy-based trust network under defined circumstances; when do these circumstances constitute a "severe problem" or "failure"? What powers such a role should have in the systems implementing the policy negotiation, or indeed whether such a role should exist, is an open question. In theory, it should not exist because the actors in the negotiation can simply decide no agreement is possible. But in practice, other authorities (such as governments) may require such a role for non-technical reasons. If so, how such a role would be implemented across multiple jurisdictions is a difficult question, especially when the jurisdictions involved are those of different nations.

[3] Excluding errors

Other issues include the extent to which protocols to implement the policy negotiation system must be adopted. This depends in part on the goals of the system. If attribution is to be ubiquitous, then the protocols (or at least interoperable protocols) must be adopted. Several policy negotiation systems might exist, each supporting different types of attribute vectors or different levels of assurance for attribute vectors; in this case, the ability to map goals from one system to the other, and to create translation mechanisms to allow the respective protocols to interoperate, define the extent to which attribution information and trust may be shared.

In fact, none of these issues are unique to networked systems; the world of negotiating structures and mechanisms is well established in the non-technical world, and many mechanisms exist in the technical world to support negotiations. All of these issues have been resolved in various ways in the physical world (including a realization that, in some cases, negotiations are not feasible).

3.4 Discussion

The requirements for an attribution framework raise a number of interesting questions.

What constitutes "adequate" attribution, and who decides what is adequate? This itself is an issue of governance. Equally difficult is ensuring the governance guiding the answers reflects the changing needs of users, administrative domains, and other interested parties.

This also raises the question of theoretical assurance—how can we reason about attribution quality and types? Clearly, dual-valued logic fails, as there are gradations of attribution. A natural candidate for this type of analysis is fuzzy logic; indeed, although not specifically identified as such, much of the description above uses ideas from that field.

An interesting aspect of this governance is the issue of revocation—when can attribution be undone? In a centralized world, a central authority could direct all networks (specifically, intermediate nodes) to discard all attribution information, and not provide any of their own. But in a distributed world, as our Internet is today, this is not possible.

Another area is selective access to the attribute vectors. Mentioned earlier, this idea raises issues of negotiation. If multiple central authorities are involved in creating (or assuring) an attribution vector, will multi-jurisdictional co-operation depend upon limiting access to the vector and if so, how will that affect their actions?

Conflicts and ambiguities will undoubtedly arise. The negotiation system and supporting infrastructure must handle them appropriately. For example, attribution may be desirable for crimes and cyber attacks, and undesirable for political speech and whistleblowers. Actors may have different, conflicting goals, and hence the success of a governance system in resolving such conflicts will be a measure of its success, although the metrics to evaluate the degree of success will have to be developed. Note that this is far more than a technical problem; it delves into the political and cultural aspects of attribution, where our culture assumes that the ability to visit a dissident web site is good, but the governments of some countries would strongly disagree with this particular belief. In some cases, notably when a whistleblower reveals information his or her organization would prefer to keep

secret, the organization will want full attribution whereas doing so would be inimical to the whistleblower.

Finally, recall that under some circumstances a requirement for false attribution exists. In addition to the intelligence agent example given earlier, whistleblowers and political dissidents may desire no attribution or a false one when they visit a reporting or government web site which requires attribution. Under what circumstances should this capability be provided? If not always available, who should determine when circumstances warrant its use?

4. IS AN ATTRIBUTE FRAMEWORK FEASIBLE?

We propose a multinational framework for attribution that requires both new technical capabilities and new policy structures. The framework will have to operate robustly among mutually distrusting parties. Is designing and implementing this framework feasible?

Historical experiences such as the nuclear arms race of the Cold War show that a system of adequate attribution among mutually distrusting parties is possible. A number of lessons from the decades-long experience with non-proliferation and arms control apply directly to building a system of adequate network attribution. Banning certain nuclear tests and intermediate-range weapons had multi-lateral support for many reasons. While it was advantageous to ban these weapons in a mutually trusting environment, gaming strategies made complying with treaties in a mutually distrusting environment [6] potentially disadvantageous.

The U.S. government's position was "trust but verify," which was easier said than done. Technology and methods to verify arms reduction or verify and attribute nuclear testing were not available. Through diplomacy, multi-lateral cooperation, technology roadmaps, and a wide variety of processes and procedures, the parties developed means to verify and attribute. Technology alone was inadequate to address the problem. A clear understanding of verification and attribution objectives drove political and technical developments, allowing the parties to construct meaningful and enforceable treaties.

Organizationally, within the U.S. government no single agency was or is responsible for non-proliferation; major agencies play key roles set by presidential directives and coordinated by the White House.

Furthermore, the U.S. experience with non-proliferation demonstrates that non-proliferation is not a static concern; issues can emerge and mature over time. "Twenty years ago, the proliferation of WMD [weapons of mass destruction] was often an afterthought in discussions of the strategic environment. With the end of the Cold War and the reprioritization of US strategy, the profile of nonproliferation grew rapidly." [7] After 1989 the President created an NSC directorate, issued new policies and directives, while Congress passed legislation. New authorities, sanctions, and regulations were developed and the Defense and State Departments created new offices to deal with the new challenge. Internationally the United States created new multi-lateral organizations for coordinated action against WMD, reenergized existing institutions, and made nonproliferation a norm for international behavior and a factor in every major initiative. Non-proliferation went from being a relatively minor part of U.S. national strategy to become one of its most critical elements [8].

Achieving enforceable treaties requires considerable investment: decades of diplomacy and treaty negotiation, thousands of individuals working together in an international setting to develop technology and procedures, and continuous refinement of treaties and practices. Nuclear non-proliferation treaties can serve as examples for creating a managed system of attribution on the Internet. Nuclear non-proliferation and an adequate attribution system on the Internet have several challenges in common: promoting recognition of the problem, achieving international cooperation, developing policies and treaties, developing and enforcing laws, creating enabling technology, and constructing a culture of continuous improvement. At the heart of all of these issues is the need to attribute actions to an actor or party. A system of adequate attribution has the additional burden of providing non-attribution under defined conditions.

From the non-proliferation and arms control agendas, we learned that pursuing goals is a long-term process, and is dynamic as new concerns emerge. In the U.S. context, the policy framework has been flexible; the President establishes the overall policy context that lays out agency missions and weaves non-proliferation into international efforts. No "central agency" is responsible; rather, the goal has been to incorporate non-proliferation efforts into many kinds of bilateral and multi-lateral projects. We anticipate a somewhat similar structure for network attribution, but private interests have a distinct and significant voice in attribution systems that will require even greater flexibility than nation-to-nation non-proliferation talks.

5. ECONOMICS

As noted, successful multilateral frameworks have been developed to address needs like non-proliferation and weapons limitations, where verification – the cousin to attribution – is key. This gives us confidence that, if sufficient incentives are provided, a system of full attribution is feasible.

The motivation for weapons control frameworks is only indirectly economic; countries that have become party to agreements such as the Nuclear Non-Proliferation Treaty continue their support in part because they perceive that the US and the four other nuclear weapons states are carrying out their own responsibilities to seriously move toward eliminating nuclear weapons. Underlying these motivations for continued adherence is the shared desire to avoid bad outcomes—a mass nuclear exchange, nuclear weapons falling into the hands of terrorists, additional nuclear-capable rogue states, and a nuclear power making a tragic mistake [26].

Our intuition is that economics rather than national self-preservation will be the primary motivation for creating a full attribution system. Methods other than communicating across linked electronic networks exist that provide many needed attribution or non-attribution characteristics (e.g., sending a registered letter) but these techniques may be either cumbersome or inconvenient for many sender-recipient pairs.

In considering the role of economics in motivating a full attribution system, three questions must be answered.

1. Is sufficient underlying economic value created that, properly allocated, would provide sufficient motivation to justify creating a full attribution system?

Our intuition is that the economic flows from a full attribution system will be considerable (although we have no evidence to support this assertion), and that a variety of business models will

emerge variously trading off trust, traffic volume, cost and even side payments from other parties.

There is a substantial body of work demonstrating that trust and privacy have a real economic value [8,9,10,11,12]. Creating a market for attribution and non-attribution among the nine sets of participants would be a way of monetizing this value. Senders, receivers, and the intermediate organizations could make side payments in order to achieve the desired attribution outcome. The real value of this market, how such markets would clear, and how they would be governed, are speculative right now, but such markets appear to be conceptually attractive. Even if the value of trust and privacy is not (fully) monetized, it nonetheless contributes to the economic surplus (either in terms of the consumer or producer surplus).

2. What choices shape the economic incentive?

Backbones and intermediate nodes face a couple of different economic models for their businesses. For example, intermediate nodes and backbones could position themselves as the most trusted intermediate carriage points for traffic with attribution or non-attribution requirements. In this case, the rationale is that by being highly trusted these carriers would obtain more traffic (but this assumes that the market for attribution will in fact be significant). Alternatively, nodes could adopt a low cost strategy—make no guarantees as to the validity of the traffic crossing these nodes, but count on transmitting significant traffic at a low cost. A more venial instance would be for nodes to accept side payments (from governments or organizations) in order to corrupt or monitor their traffic, without the knowledge of other attribution system participants.

Governments and organizations also have to make choices as to how they are positioned in providing a trusted range of attribution choices. To cite a banking analogy, at one end of the spectrum are the trusted Swiss; at the other end would be countries like Nigeria.

Policy choices may shape the ultimate network economics. By treaty, international telephony provides payments to less developed countries to support their connection to the multinational network (in total, such reverse payments are on the order of 8-10 billion USD per year). The Internet has no such structure of reverse payments, but such a system might be a powerful incentive for select countries to provide and participate in a trusted attribution system. This payment structure deserves careful analysis.

Finally (and related to the previous paragraph) is the question of "who pays." The attribution system as outlined in this paper would require significant investment in multilateral capabilities that do not now exist. In the fully developed version of an attribution system as described these include:

- A common multilateral policy framework to formalize the cooperation, definitions, and collaborations necessary for attribution across administrative, jurisdictional, and national boundaries;
- Technical cooperation far exceeding the agreements in principle now extant. Such cooperation would fill important gaps, such as researching and recommending the best attribution techniques, and providing on-going support for a multilateral attribution capability;
- Negotiating structures (not just for senders and receivers, but all nine sets of parties involved) with defined terms for levels of attribution and non-attribution to be associated with each message; and
- Policy based trusted network routing across backbones and nodes. Ideally a formalized metric for trustworthiness would be developed and used as the basis for routing decisions.

All of these initiatives are necessary parts of the fully mature attribution system we have outlined.

We believe there to be sufficient underlying economic value to motivate a full attribution system. How that economic value is distributed will depend both on choices of firm strategy and public policy. These strategy and policy choices are closely intertwined with the third question: *how might a full attribution system develop?*

6. CREATING THE ATTRIBUTION SYSTEM

Networked systems can emerge and grow either organically or through unilateral action. The launch of ARPANet (funded by the U.S. Government) and the launch of electric power grids (funded by Thomas Edison) are instances of unilateral action. The emergence of fax machines exemplifies organic growth: organizations first bought fax machines to transmit documents between their offices, and only later began using them to connect with outside fax machines.

We hypothesize that attribution systems will grow organically, and appear initially as value added services. There is a rich menu of potential initial attribution services that could be positioned as value added offerings.

If a policy based routing mechanism were such as to allow sender-receivers to specify the exact route, then only a small set of nodes need offer value-added attribution services initially. In this case the entire full attribution system does not have to exist in order for combinations of the nine path entities or subsets of them to agree on a value-added package of attribution services.

Most likely and easiest to imagine would be initial offerings where strong attribution is provided to cooperating sender-receiver pairs, or even where strong attribution is provided only for portions of the path between cooperating senders and receivers. Such initial offerings would not require any negotiation or policy selections; the service could be offered on a take-it-or leave-it basis at first.

If messages were to be routed dynamically (i.e., route selection by senders was not provided) then organic growth of value added services would be far less appealing, because no one in the attribution system could guarantee that their services would be employed. In this case, a full attribution system might never develop unless it were imposed unilaterally by some joint network wide agreement. This development is unlikely.

If strong policy based routing selection were available, over time we would expect offerings of value added services to expand. Logical evolutionary steps going forward would include (in no obvious order):

- Negotiation or policy choice mechanisms between cooperating senders, receivers, and the intermediate path nodes;
- Expansion of attribution paths both geographically and functionally (more countries, more backbones, ISPs and organizations);

- Development of trust ratings for different path elements;
- Financial flows and payment/accounting regimes.

More challenging will be creating full attribution between conflicting senders and receivers. Two paths for expanding attribution to these cases exist. First, already described, would be to implement technical specifications establishing the policy based trust network. This may prove very difficult if not impossible to do.

An alternative (not mutually exclusive from the first option) would be to institute a sanctions regime or other external pressures to incent non-cooperating parties (particularly governments) to change their policies. The precedent for such multinational regimes exists: the G-8 Financial Action Task Force (FATF). The FATF comprises countries that have agreed to observe certain best practices for international financial transactions. FATF's goal is to make money laundering more difficult and more easily detected. The group develops best practices and standards and will not accept a new member until it has made progress in adopting them. Members who fail to live up to their obligations face sanctions from the financial community. Until 2002 the FATF also had an evaluation mechanism which yearly put some countries on the "Non-Cooperative Countries and Treaties" list (the "Blacklist"). While under international law the FATF Blacklist carries no formal sanctions, in practice a country on the Blacklist often found itself under intense financial pressure. Most larger countries with significant financial centers consider transactions involving a country on the FATF Blacklist to a suspicious activity, triggering greater regulatory scrutiny. This listing appears to have put pressure on those countries to cooperate in fighting money laundering; the list shrank from 15 to 3 at which point it ceased to be updated.

Unlike financial transactions, Internet traffic is mostly free of government oversight. A sanctions regime would require a new level of state cooperation and involvement in Internet management. Though difficult to create, such a regime is not beyond contemplation.

The types of organizations most likely to adopt attribution frameworks early are large law and business firms in North America, the United Kingdom, and Europe. They will use such a system to assure that business and legal documents are being sent to, and received by, the parties involved.

7. LIMITATIONS OF ATTRIBUTION

The attribution policies discussed so far create an interesting, and yet realistic, dichotomy. Consider the attribution policy of first origin.

This policy states that the network operators can trace coordinated entities back to their origin. The utility of this policy arises from distributed denial of service attacks using botnets, where the immediate origins of the messages are known (the bots); the policy requires the ability to trace back to the distribution points, or distributors, of the bots. In the context of tracing network attacks, the first origin policy is not merely reasonable; it is salutary, because it minimizes disruption and suspicion of those unwitting people and systems on which the botnet entities run.

Now, consider the same policy in a political context. A nation with repressive political policies discovers a large number of messages that poke fun at the government. The first origin policy allows the government to trace back to determine the origin of these coordinated entities (one or more messages). The ability for the dissidents (or ordinary citizens) to criticize their government anonymously no longer exists.

This leads one to ask the purpose of attribution. Attribution, or rather the lack of attribution, provides the ability to send messages without fear that the entities involved can be identified. Differing levels and types of attribution modify the level of fear, and the ability to send such messages, in various ways.

The ability to conceal the origin of messages affords the sender protection from reprisals. The example using political dissidents is one context in which this ability is critical. Another example is whistle blowing, in which a subordinate reports actions of a superior (or an equal) to an external authority, such as the press or regulatory or law enforcement. Extending this to an agency or country, the ability to deny attribution allows an attacker to place the target in a state of confusion, a tactic of warfare encouraged by Sun-Tzu, among others.

This ability also enables one to protect privacy. The right to be "let alone" enables one to live one's life without interference and without having to account for one's actions. As an example of the value of such privacy, consider someone who wants to learn to use the Python programming language. He goes first to the web site http://www.python.com. The pornographic images on that site indicate it is not the site where one may download the language interpreter, so he tries http://www.python.org, which is the correct site. But anyone observing his activities would see he visited a pornographic web site, and from that could (erroneously) conclude he was downloading pornography.

With privacy comes power. A lack of attribution enables entities to avoid taking responsibility for their messages. For example, experience with the anonymity that the Internet affords shows that it prevents those who are the targets of slanderous communication from identifying the sources, and taking legal or other actions to protect themselves.

The sources here may include the government. In many countries, people are tried (either in a court of law or in the court of public opinion) without being told who their accusers are or what specific offenses are alleged. So the lack of attribution that protects the individual also can harm the individual.

Further, the ability to trace messages enhances the ability to detect attacks at the non-cyber level, ranging from individual threats (for example, harassment) to societal threats (for example, terrorism and warfare). Thus, this point of view stems from a belief that providing attribution encourages social order and protects both individuals and society.

There is no right answer to the level of attribution that should be provided. This is a policy issue that must be decided somehow, either by a deliberate crafting of policy or by an acceptance of the existence of tools and services that can provide varying degrees of attribution.

Ultimately, there may be several distinct and separate Internets, each with a different level of attribution, and people who desire disparate levels may simply be unable to communicate. While disquieting, this mimics the non-cyber world perfectly. Two people may talk, but one may not believe the other's claims because the attribution of those claims is insufficient for the skeptic's purpose. That the speaker cannot provide the level of attribution that the listener desires interferes with communication, and in some cases simply cannot be overcome. So in this way the

use of attribution in cyberspace has the same effects as the use of attribution in "realspace."

8. RELATED WORK

The technical literature contains a large body of work on some forms of attribution, focusing primarily on IP-traceback (e.g., see [1,2,15,16]) as a scheme for achieving attribution. The goal of IP traceback is to determine the originating IP address of a packet, regardless of the stated originating IP address, which is intended to address attribution in the case of spoofed-source denial of service attacks. Thus this approach addresses only a subset of the attribution problem that we present. While there are other technical approaches to attribution, as surveyed by Wheeler and Larsen [5], these approaches also address only limited aspects of the attribution issue as a whole.

Pyun and Reeves [17] take a different approach to attack attribution—they provide a graph theoretic approach to the optimum placement of sensors on the internet. They find that monitoring only 5% of autonomous systems will provide attribution of packets with little ambiguity. They do not address, however, the policy issues surrounding attribution, nor do they acknowledge that attribution may not be desirable under certain circumstances. Thus their approach also addresses only one part of the attribution problem.

An early paper by Staniford-Chen and Heberlein [3] identified the issue of attribution in terms of accountability, and more specifically the need to hold attackers accountable. They focused on the problem of "stepping stones"—where an attacker chains his activity through multiple machines in order to obscure his origin—and developed a technique called "thumbprinting". A thumbprint is a summary of the contents of a connection, and can then be used to determine if connections between two different machines (e.g., between machines A and B, and between machines B and C) are similar or the same, thus potentially indicating stepping stone behavior. While this paper highlights the need for accountability, its solution also addresses only one part of the attribution problem. However, this paper does address determining the ultimate source of a message, regardless of any intermediate chaining, which is an identified weakness with our current framework.

Strayer et al. [18] provide an attribution architecture in their 2003 paper. More specifically, they provide an architecture that combines packet-tracing techniques (e.g., such as IP traceback approaches) with stepping stone detection (such as thumbprints) to provide integrated source attribution. Their architecture does not address any policy considerations or identify cases where non-attribution is preferable.

These papers do not address a framework that spans both the technical and non-technical realms, focusing instead on the technology required to implement their notions of attribution, and in some cases on the uses to which it may be put. Our framework expressly goes beyond this, focusing on the higher-level implications and requirements and eschewing many technical details because those depend on the precise characteristics of interest.

Daniels and Spafford [19] investigated the attribution issue, calling such systems Network Traffic Tracking Systems (NTTS). They identify several attributes for such a system—accuracy, precision, the ability to resist subversion, low overhead, low cost, scalability, real time tracking, privacy and control, and note that the attributes may be orthogonal to each other. Like us, Daniels and Spafford identify, albeit briefly, the need for discussion around the privacy issues and who controls access to the attribution information (e.g., government authorities, the final recipient, or others), and use the case of a whistle-blower to indicate the complexity of the issue. They provide three models of NTTS, the "Internet Model" being the one that most closely resembles our framework. They identify a number of practical issues—which are also applicable to our framework—in this model, and question if an NTTS is even desirable given the privacy implications.

This paper introduces many elements that our framework addresses. In particular, we assume that the desirability or undesirability (such as the level of privacy to be provided) can be determined using the policy negotiation system we introduce. Critical to our work is the observation that *all parties* must have a say in the attribution characteristics and assurances that they are willing to accept—even if only to say, "None!" In this way, we generalize the question that Daniels and Spafford ask, to encompass not only privacy but also other attribute characteristics of importance to the actors.

Much research has been done on negotiation systems; however, these tend to focus on negotiation strategies [20] or specific issues or environments such as network management issues [21,22], trust negotiation [23], and electronic commerce [24]. Closest to our proposed system is SCENS [25], designed to support data sharing but easily adapted to different forms of negotiation and environments such as governance. These works all focus on implementations, whereas our discussion focuses on goals and the properties needed to meet those goals. Many of these systems could be used to implement parts of our framework.

9. CONCLUSION

Adequate attribution deals not just with identifying the source of attack packets addressed to a particular recipient. An adequate attribution system needs to address a wide range of needs held by at least nine different types of network participants. An expanded—and necessary—view raises a number of policy and technical questions not answerable under current approaches. Such an expanded view of attribution points towards major policy reforms, as well as technical needs. Even so, such a system will not meet the needs of all parties. Conflicts will emerge between the various network interests; the extent to which such conflicts can be resolved satisfactorily will shape the extent to which some groups may choose, for instance, to stay off of the Internet. Defining what the needed policy structure is, what needs it will satisfy, and equally what needs it will not satisfy, is an important step forward in creating a security structure for networked systems. This will allow us to delimit what works for attribution, and what does not.

This attribution framework cannot be imposed upon a structure with control as fragmented as that of the Internet. Such attempts will fail miserably; the United States, North Korea, Iran, and Israel would be unlikely to agree on mechanisms and policies for attribution (or anything else, for that matter). And even where the technology provides added benefits (such as Secure DNS), the Internet community is slow to adopt it. Thus, this framework will grow slowly, piece-meal, in a modular fashion, with incompatibilities between interfaces being ironed out or accepted as reality, resulting in the development of "attribution networks" as discussed above. Perhaps legal issues, such as criminal and

civil liability, will drive the development of different corporate and governmental attribution policies and enforcement mechanisms, while the individualism rampant in other parts of the Internet will drive the non-attribution policies and mechanisms. Only time will tell—and the need for attribution and non-attribution ensures that, as the ancient curse claims, we shall live in interesting times.

To summarize, an adequate delineation of attribution involves:

- Multiple parties, not just the receiver attempting to identify the origin of the sender;
- Different needs and requirements for these parties, with negotiating possible; and
- Imperfect trust relationships between all parties.

Formalizing this structure requires consideration of the context of the messages sent, including content. Trustworthiness of the links is critical. There are opportunities to make improvements in the technology. However, we are still left with the situation where users are unable to use the Internet with the desired levels of attribution.

However, while the multinational policy and technical requirements for adequate attribution may appear daunting, positive historical examples of similar systems suggest that the framework outlined is feasible to create.

10. ACKNOWLEDGMENTS

The authors thank the paper's shepherd, Brian Snow, for his helpful comments and suggestions. Matt Bishop acknowledges the support of awards CNS-0716827 and CNS-0831002 from the National Science Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

The authors also thank the movie "Sisterhood of the Traveling Pants" for providing the inspiration for the title.

11. REFERENCES

[1] Hal Burch and Bill Cheswick. Tracing anonymous packets to their approximate source. In LISA '00: Proceedings of the 14th USENIX Conference on System Administration, pages 319–328, Berkeley, CA, USA, 2000. USENIX Association.

[2] Stefan Savage, David Wetherall, Anna Karlin, and Tom Anderson. Practical network support for ip traceback. SIGCOMM Comput. Commun. Rev., 30(4):295–306, 2000.

[3] S. Staniford-Chen and L. T. Heberlein. Holding intruders accountable on the internet. In SP '95: Proceedings of the 1995 IEEE Symposium on Security and Privacy, page 39, Washington, DC, USA, 1995. IEEE Computer Society.

[4] Meriam Webster. Definition of Attribution. In Meriam-Webster's Online Dictionary, 1651. Last Visited: 16 April 2009.

[5] David A. Wheeler and Gregory N. Larsen. Techniques for cyber attack attribution. Technical Report IDA Paper P-3792, Institute for Defense Analysis, October 2003.

[6] Jeffrey Hunker, Bob Hutchison, and Jonathan Margulies, Role and Challenges for Sufficient CyberAttack Attribution, Dartmouth College: The Institute for Information Infrastructure Protection (The I3P), 28 January 2008.

[7] Representative James R. Langevin, Representative Michael McCaul, Scott Charney, and Lt. General Harry Raduege, Securing Cyberspace for the 44th Presidency; A Report of the CSIS Commission on Cybersecurity for the 44th Presidency, Washington

[8] Hauberman, Bernard, Adan, Etan, and Ane, Leslie 2005. Valuaating Privacy. IEEE Security and Privacy Vol.3, Issue 5 (September-October 2005) 22-25.

[9] Acquisti, A. and Grossklogs, J. 2003. Losses, Gains, and Hyperbolic Discounting: An Experimental Approach to Information Security Attitudes and Behavior. Proceedings of the 2nd International Workshop on Economics and Information Security. www.cpppe.umd.edu/rhsmith3/agenda.htm

[10] Hann,I.H.et.al. 2003. Overcoming OnLine Information Privacy Concerns: A Comparison of Privacy Policies, Convenience, and Promotions. Working Paper (September 2003). Marshall School of Business, University of Southern California. www.rcf.usc.edu/~hann/publications_files/overcoming%20online%20information%20privacy%20concerns.pdf

[11] Kleinberg,J., Papadimitrou,C.H., andRaghaven,P. 2001. On the Value of Private Information. In Proceedings of the 85th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2001), J. van Bertherm, ed. Morgan Kaufman, 2001. 249-257.

[12] Chang,a.,Kannan,P.K.,Whinston, A. 1998. 'Goodies' in Exchange for Consumer Information on the Internet: The Economics and Issues. In Proceedings of the Thirty-First Hawaii International Conference on System Sciences Vol. 4. (Hawaii, USA January 6-91998). IEEEE, 533-542.

[13] S. Bellovin, "The Security Flag in the IPv4 Header," RFC 3514 (Apr. 2003). http://www.ietf.org/rfc/rfc3514.txt

[14] L.Weinstein, "Google Hijacked -- Major ISP to Intercept and Modify Web Pages" (Dec. 2007). http://lauren.vortex.com/archive/000337.html

[15] Alex C. Snoeren. Hash-based IP traceback. In SIGCOMM '01: Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications, pages 3-14, San Diego, CA, USA, 2001.

[16] Alex C. Snoeren, Craig Partridge, Luis A. Sanchez, Christine E. Jones, Fabrice Tchakountio, Beverly Schwartz, Stephen T. Kent, W. Timothy Strayer. Single-packet IP traceback. IEEE/ACM Transactions on Networking (TON), 10(6):721-734, 2002.

[17] Young June Pyun and Douglas S. Reeves. Strategic Deployment of Network Monitors for Attack Attribution. In BROADNETS '07: Proceedings of the Fourth International Conference on Broadband Communications, Networks and Systems, 2007, pages 525-534, 2007.

[18] W. Timothy Strayer, Christine E. Jones, Isidro Castineyra, Joel B. Levin, and Regina Rosales Hain. An Integrated Architecture for Attack Attribution. BBN Technical Report No.8384, December 31, 2003.

[19] Thomas E. Daniels and Eugune H. Spafford. Network traffic tracking systems: folly in the large? In NSPW '00: Proceedings of the 2000 workshop on New security paradigms, pages 119-124, Ballycotton, Ireland, 2000.

[20] Sheng Zhang, Song Ye, Fillia Makedon, and James Ford, "A Hybrid Negotiation Strategy Mechanism in an Automated Negotiation System," *Proceedings of the 5th ACM Conference on Electronic Commerce* pp. 256–257 (2004).

[21] T. Atdelzater, E. Atkins, and K. Shin, "QoS Negotiation in Real-Time Systems and its Application to Automated Flight Control," *IEEE Transactions on Computers* **49**(11) pp. 1170–1183 (Nov. 2000).

[22] Karl Czajkowski, Ian Foster, Carl Kesselman, Volker Sander, and Steven Tuecke, "SNAP: A Protocol for Negotiating Service Level Agreements and Coordinating Resource Management in Distributed Systems," *Job Scheduling Strategies for Parallel Processing* pp. 153–183 (2002).

[23] M. Winslett, T. Yu, K. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, and L. Yu, "Negotiating Trust in the Web," *IEEE Internet Computing* **6**(6) pp. 30–37 (Nov. 2002).

[24] Samuel Choi, Jiming Liu, and Sheung-Ping Chan, "A Genetic Agent-Based Negotiation System," *Computer Networks* **37**(2) pp. 195–204 (Oct. 2001).

[25] Song Ye, Fillia Makedon, Tilmann Steinberg, Li Shen, James Ford, Yuhang Wang, Yan Zhao, and Sarantos Kapidakis, "SCENS: A System for the Mediated Sharing of Sensitive Data," *Proceedings of the 3rd ACM/IEEE-CS Joint Conference on Digital Libraries* pp. 263–265 (2003).