

Attribution Requirements for Next Generation Internets

Jeffrey Hunker
Jeffrey Hunker Associates
Pittsburgh, PA 15232
Email: hunker@jeffreyhunker.com

Carrie Gates
CA Labs
New York, NY 10022
Email: carrie.gates@ca.com

Matt Bishop
University of California, Davis
Davis, CA 95616-8562
Email: bishop@cs.ucdavis.edu

Abstract—The notion of attribution is usually tied to identity: the ability to associate the originator of something with that data. This notion is both a simplification of the myriad aspects of attribution, and a masking of many different types of attribution under the rubric of “identity.” Current efforts at attribution of data in the network focus either on traceback of packets or signing data at various network layers, and management of the keys involved. As the Internet evolves into the next generation network, and people implement testbeds to facilitate that evolution, they can design support mechanisms for other forms of attribution. What types of attribution should those mechanisms support? This paper presents requirements for attribution that will be useful both in the next generation infrastructure and in the data it manages.

I. INTRODUCTION

Attribution, defined as the “assigning or ascribing of a character or quality” [1], is a central problem in computing. It raises technical issues such as from which host or networks packets or data originated. It also raises legal and other related problems such as how one can prove (to an appropriate legal or other standard) that a specific individual created, or sent, a document over the network. As the world moves towards electronic commerce, the latter will increase in importance because paper contracts and receipts will become a thing of the past, and the *only* legally binding evidence of a contract or actions will be electronic.

Several papers have explored different facets of attribution. Wheeler and Larsen [2] and Hunker, Hutchison and Marguiles [3] discuss attribution of attacks in cyberspace. Strayer *et al.* [4] consider an architecture to support this, but do not focus on requirements beyond those needed for the architecture. Bishop, Hunker, and Gates [5] extend this idea to consider a general framework for the problem of attribution in the next generation Internet, and outline the technical, economic, and policy issues that constrain solutions to the problem. This paper presents generic policy requirements for attribution, and why they arise.

In this paper, *attribution* is “the binding of data to an entity.” So, for example, determining the identity of the sender of a message is attribution—binding data (the identity) to an entity (the sender). Similarly, attributing a delay in forwarding a packet to a particular network binds data (the length of the delay) to an entity (the particular network).

We now present background necessary for the requirements.

II. BACKGROUND

Previous work on attribution in computer security rests on two basic assumptions: first, that the ability to attribute identity, or the property of interest, is always beneficial; and second, that the needs of the various stakeholders are closely enough aligned that one can assume the needs of one (such as the security analysts) will satisfy all. Neither is in fact accurate.

Consider the ability to attribute a property (identity, for expository purposes) in the scenario of a user accessing a web site. The system administrator of the web site wants to know who the user is, so the site needs to be able to attribute identity to the user (sender of the messages). Now, suppose the web site is set up to enable anonymized traffic. In this case, the web site would not want to be able to attribute identity to the users; the *anon.penet.fi* incident, in which a Finnish court ordered such a site to reveal the identity of the anonymous posters, provides an example of the drawbacks. Thus, we have two different situations: in the first, the recipient (the web site) wants sender (user) attribution, and in the second, the recipient explicitly does *not* want sender attribution.

The different types of attribution are:

- *Perfect attribution*, in which the binding of the data to the entity is known;
- *Perfect non-attribution*, in which the binding of the data to the entity is unknown and undiscoverable;
- *Perfect selective attribution*, in which the binding of the data to the entity is known to some set of entities, and unknown and undiscoverable by other entities;
- *Imperfect attribution*, in which the binding of the data to the entity can be discovered, but to do so takes long enough that the knowledge is useless or redundant, or cost more than the value of knowing the attribution;
- *False attribution*, in which the binding of the data to the entity appears to be known, but the attribution is incorrect but consistent over time;
- *Randomized false attribution*, which is false attribution without the consistency over time; and
- *Unconcern*, in which an entity does not care about the binding of the data to the entity.

In the first scenario, the system administrator requires perfect attribution of identity to the user. In the second, *anon.penet.fi* wanted perfect non-attribution of a message

to an email address, but in fact had only perfect selective attribution, because the site manager could reconstruct the correct attribution. In the third, both the user and the bank want perfect attribution of identity to each other. In the last, the counterintelligence agent wants either false attribution (if her goal is to present a false but consistent identity across visits) or randomized false attribution (if she wants to appear as different visitors for each visit).

As we are considering attribution in a network, we consider only those entities involved in the sending and receiving of a message. The entities are the sender and recipient, which may be the individual, the associated account, the sending process, the computer, the individual's organization, or some other appropriate entity; the ISP and network backbone providers that carry the message to its destination; and the political entities with jurisdiction over all these entities.

The requirements for an attribution system for the Internet, which follow, are not yet supported to the level of assurance most users of attribution would desire. So these requirements are intended for the next generation Internet, and for other networks for which attribution is to be a feature.

III. REQUIREMENTS

Consider a packet on a network. We call the point of origin the "sender" and the point of destination the "recipient." In some cases, for example an email, the "sender" of interest is the human who created the message. In other cases, for example a web page, it may be the server process. In still other cases, the actual originating entity may *not* be the entity of interest; when one discovers packets coming from a botnet, the specific botnet is of less interest than the perpetrator who installed the botnet. The sender may even be an intermediary entity, for example an ISP that injects advertisements into packets responding to certain addresses [6].

REQUIREMENT 1. The sender and recipient must be defined.

Given the possible stakes inherent in the use of an attribution system, the entities relying on the attribution must know how confident they can be in it. So the system needs to provide some indication of the degree of confidence that a user can have in the attribution being correct. The means of reporting the attribution must enable metrics to assess the accuracy of the attribution. In general, the desired attribution depends on three interrelated factors: the desired confidence in the attribution, the nature of the actions for which attribution is desired, and the intended purpose of the attribution. The attribution assurance desired will depend on the consequences that might follow from that attribution. For instance, authentication for a web site requires less assurance than does authentication for obtaining a passport. We discuss this further in Section IV-C.

REQUIREMENT 2. The levels of attribute assurance must be specified or determinable by the attribution framework, and must be measurable.

The attribution mechanisms must know what characteristics to report. The sender and recipient may require that certain

characteristics be attributed to one another; a third entity may require different characteristics attributed to them. Thus:

REQUIREMENT 3. The policy requirements of entities must be specified, and the attribution framework in its full form should allow for specifying a range of possible attribution policies.

As shown above, multiple parties shape whether or how entities show the characteristics associated with an entity have a particular value. Each party 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.

REQUIREMENT 4. The attribution system must include the different actors, other than the sender and recipient, that have an interest in attribution.

Consider how these parties interact. Suppose Alice sends a message to Bob. Bob wishes to attribute the message to Alice. Alice's ISP also wishes to attribute originating time and router data to the message for Bob's ISP. Moreover, the backbones that carry some of the packets wish to share attributable data about the packets among themselves. So, there are multiple attributions by multiple actors. Further, the layer of abstraction at which the entity is viewed changes the attribution. The sender attribution of the message has the value "Alice". But the attribution of origination time and router data will be tied to the *packets* that make up the message, and not the message itself. Thus, the attribution of the message by the sender is at layer 7, but the attribution of the packets by the ISP here is at layer 3. This leads to the next requirement.

REQUIREMENT 5. The attribution framework must allow for attribution to originate and be received by actors at different layers of the (network) stack.

We represent the attribution values for a particular entity using a vector. Each position in the vector corresponds to a characteristic the value of which is to be attributed. A distinguished value must exist that means that, for this particular entity and characteristic, the value is unknown (whether due to lack of information or relevance).

Attribution as used today is concerned with *identity*: who originated the packet, who signed the contract, and so forth. But other characteristics may also be of interest. For example, the time at which the message was sent indicates whether certain deadlines are met. The time at which a message entered and exited various networks provides information about delays, which may relate to contractual requirements for quality of service. Whether the message was protected by encryption or access control bits, and where geographically the message traveled, may bear on the ability of an adversary to read, alter, or block the message. The originator of the attribution, and the user of that attribution, must have a common set of values which they can meaningfully supply and interpret.

REQUIREMENT 6. The definition of, and specific values for, the values of the elements of the attribution vector must be

well defined.

In the above example, the recipient of the attribution of the message and the recipient of the message are the same (Bob). Some other attribution information is reported to other entities involved in the transportation of the message from Alice to Bob (specifically, the ISPs and backbones). It is unlikely that Bob (or Alice) wants the ISPs or backbones to see the message. This emphasizes that the content of *the attribution* is quite different than the contents of the *entity* about which information is being attributed. It also emphasizes that the attribution framework must support entities other than those involved in the communication, for example a government agency measuring use of the network or a court that has authorized the recording of attribution information. The framework must allow these recipients of attribution information to be identified.

REQUIREMENT 7. The entity or entities that receive the attribution report must be specified.

One important characteristic is that of why the message was sent. Perhaps this is the most challenging information to attribute; in many situations, it will be *the* most important aspect of attribution. How to produce an adequate answer remains an open research question, especially because of the need to examine human motivations. Motivations are notoriously hard to determine by skilled investigators, let alone by an automated system.

REQUIREMENT 8. There must be a means to specify the characteristics giving the reason that the message was sent.

Many 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 (which does 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. Due to the financial cost, ISPs may want to provide attribution services only if they are profitable and the ISP is unlikely to be sued. Legal constraints may also shape attribution; privacy rules in the European Union being considerably more restrictive than those in the United States, an ISP in the former would be unable to provide the attributions that the latter could provide. Financial and legal considerations are central to the business judgment about whether to provide any service.

Other organizations than ISPs also affect attribution. Consider a message being sent from the United States to Russia over a network that transits North Korea, which may add questionable attribution information. Thus, the attribution characteristics from intermediate nodes, or that relies on intermediate nodes, are affected by the organizations controlling those nodes.

REQUIREMENT 9. There needs to be a means to define, and then specify, the requirements/interests of ISPs, backbones, and other parties.

Senders and recipients that co-operate provide attribution capabilities. Agreement on a desired level of attribution, and to the party to which the attribution applies, requires carefully defined and commonly accepted attribution characteristics, 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 have no generic incentive to co-operate. Thus, cooperating senders and recipients may have to specify some attributes of the network path (for example, no packets can go through North Korea) to ensure the agreed upon attribution.

REQUIREMENT 10. There must be a structure for *efficiently* defining policies for the special case of cooperating senders and recipients.

With many different actors potentially involved in the attribution, a policy negotiation will be required in order to establish an agreed upon attribution vector. We call this agreed-upon attribution vector 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; at the other extreme is the religious canon, which changes very slowly if at all. Which structure will predominate cannot be predicted; a policy negotiation system must first and always be workable and agreeable to all parties. Given the different goals and needs of the different parties with a stake in attribution, and having defined who all of the players are and their needs, a full attribution system should have several features.

REQUIREMENT 11. A common nomenclature for attribution vectors must be defined.

These policy elements provide 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.

REQUIREMENT 12. A system for communicating and negotiating the policy contract must be created.

The policy negotiation system should be transparent, low cost, and made routine to the extent possible for all parties involved. No system that requires a complex, legalistic structure in most cases will work for a commonly accepted attribution framework. Further, the entities involved need to communicate their desired attribution characteristics (and values) and the requisite assurance they require in order for a message to be sent and received. At a minimum, the sender must be able to

specify a level of attribution and the recipients must be able to specify the levels of attribution it finds acceptable. The desires may be incompatible; for example, a sender may require that messages not be attributable to the sender but the recipient may require full attribution to the sender. In this case, the policy negotiation system must attempt to resolve the conflict, reporting that the conflict cannot be resolved if it cannot be.

Entities other than the sender and recipient may affect the resolution of the negotiation, and indeed cause it to break down. Consider a government that requires perfect attribution of the sender in all electronic communications. The sender and recipient both require perfect non-attribution with respect to the sender of the message. But the government's requirement that the identity of the sender be attributed to the message violates this requirement. Thus, even though the sender and recipient in the negotiation agree on the content and level of assurance of the attribution, a third party (the government) does not, and thus the agreement cannot be completed.

REQUIREMENT 13. The policy negotiating system should allow relevant actors to specify and communicate desired attribution states and levels of assurance.

The policy negotiation system must not allow unwanted accidental outcomes, in which the attribution that the entities agree to is not in fact what is desired. A dissident web site needs to advertise its policy of *not* accepting any forms of attribution *before* a prospective user accidentally provides it. If the policy negotiation accidentally requires the user to submit attribution, not only will the web site refuse to accept the message, but also any observers monitoring traffic to the web site will discover the identity of the user.

Policies need to be advertised so that entities can discover each others' policies. In particular, there needs to be an authoritative (distributed or centralized) repository of policies, so the policy negotiation system can determine what those policies are. Then parties can determine if communication is possible, and if so how to initiate it. The earlier examples of senders and recipients with incompatible policies, or whose communications involve other entities with incompatible policies, show situations in which communication must be initiated out of bands, with entities relaxing their policies (for communication to occur) or a denial being transmitted out of bands, because entities will not relax their policies.

These properties suggest specific requirements to support policy negotiation systems.

REQUIREMENT 14. A trust network must be defined 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; some entity must configure them to do so. Thus, signers of a policy contract must have some level of trust that the other actors will provide acceptably accurate attribute values. This trust system might function much as a reputation system would.

This mechanism will ensure that the entire policy contract

negotiation mechanism is enforceable. The enforcement 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. It might impose social punishments for violating agreed upon policy contracts, up to and including ostracizing those who breach them. Under certain circumstances and given an appropriate environment, it might even propose legal action (which, of course, humans would then have to take).

REQUIREMENT 15. There must be a flexible verification mechanism for ensuring that policy contracts are satisfied.

The final requirement arises from the realization that any attribution policies must exist throughout the message's (or packets') travels. While an "attribution wrapper" is technically possible, using a digital signature mechanism to bind the attribution vector to the message, the intermediate nodes may have to add additional data (creating multiple attribution labels), complicating the key management problem, or could simply delete the attribution vector and appropriately alter any indicators in the now-unsigned packet. Hence trust in attribution is based in part upon the route used.

We note that an Internet-wide public key infrastructure (PKI) will enable end and intermediate entities to create digitally signed attribution vectors, but such a PKI does not now exist. Such a PKI would be a forest and not a tree, as governments will be unable to agree on a single entity to be the root. Further, cross-certification between governments with hostile relations, such as North Korea and South Korea, is equally unlikely to occur. Such a structure also does not deal with the *deletion* of signed attribute information. An alternate approach, requiring that the architecture of the Internet be changed to support attribution, is equally infeasible for similar, and other, reasons. Thus:

REQUIREMENT 16. A policy-based routing mechanism must be defined to ensure that messages traverse networks and midpoints with appropriate attribution mechanisms and levels of trust.

IV. GOVERNANCE ISSUES

The discussion so far has raised a number of issues that require one or more entities to resolve them. These are issues of governance. Governance need not imply centralized authority. How collective issues are resolved may not be the choice of any system designer. It may simply evolve, much as today's Internet evolved from a collection of various independently managed and interconnected networks

A. Who Makes the Decisions?

The obvious way to make key attribution framework decisions such as dispute resolution is to empower a privileged user, or set of privileged users, to override normal user controls. Computer systems, and other systems, use this mechanism to correct severe problems such as failures—the entity is referred to as the "superuser" or "Administrator,"

depending on the system. Where the threat of such users is deemed too great, power can be divided in two ways. Either several such users are defined, and a threshold number must agree to act, or several users with more limited, and different, powers are defined. What these powers are in a policy negotiation system is an open question, as is whether such users should even exist. In theory, they are unnecessary; actors can simply decide that no agreement is possible. But in practice, some entities may require such a role for non-technical reasons, for example when law enforcement requires sender attribution for a set of messages. Implementing such a role across multiple jurisdictions, especially multiple nations, is a very difficult question.

In order to support decision making, common or at least interoperable protocols must implement the policy negotiation system. A weak version of a common protocol framework would specify a common default form of attribution, leaving open the opportunity for more sophisticated attribution options in some networks. Alternatively, 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 is unique to networked systems. Negotiating structures and mechanisms exist everywhere, and the analysis of their functionality is a well established topic in the non-technical world. Many mechanisms exist in the technical world to support negotiations. All of these issues have been managed in various ways in the non-technical world—including the realization that, in some cases, negotiations are not feasible.

B. Economic Factors

Economic considerations will shape the development and behavior of any attribution framework. An attribution framework can create new economic value, and it is important for system designers to understand some of the potential economic incentives at work on different actors.

Clearly trust and privacy have economic value. Similarly, to some entities, attribution has an economic value. Lawyers would seem to find a system that provides perfect attribution useful in supplanting paper with electronic filings, as they could then demonstrate the signers of such e-documents to an acceptable level of certainty.

Entities desiring attribution (including non-attribution) might be willing to pay for such a feature. Intermediate entities such as backbones and ISPs could market their ability to provide certain attribution requirements. They could also adopt a “low cost” strategy in which they make no guarantees about attribution, and count on the low cost attracting large amounts of traffic. In fact, they could even take payments on the side from organizations to monitor traffic, or corrupt it, without the other entities knowing it; such actions might be criminal, or actionable in a court of law.

C. Attribution and Privacy

We define the *attribution privacy of entity A with respect to entity C* to be the binding of A to a characteristic being kept secret from C. For example, Alice sends a message to Bob with perfect attribution, but the attribution that the message Bob received is from Alice is kept secret (in whatever way) from Cathy. Then Alice has attribution privacy with respect to entity Cathy. Whether the content of the message is secret or not is immaterial to the attribution privacy of Alice—what is kept secret from Cathy is the attribution of the origin of the message being Alice.

Classical cryptography can provide message secrecy or attribution of origin, but not both simultaneously. Suppose Alice and Bob possess a shared secret key. Alice sends Bob a message, and Bob can decrypt it because both possess the secret key. But Bob cannot prove to a disinterested third party (a “judge”) that the attribution of the message origin has value Alice. Bob could have created the message himself, since he possesses the same secret key as Alice.

An *attribution privacy violation of entity A with respect to entity C* occurs when the attribution from A is known to C, but A desires that the attribution be kept secret from C.

Again, suppose Alice sends a message to Bob. Bob can attribute the message to Alice. Alice has attribution privacy with respect to Cathy if Bob cannot demonstrate to Cathy that the message is provably attributable to Alice. All Cathy can accept is Bob’s word that the message is attributable to Alice. Alice’s attribution privacy is violated if Bob sends a message to Alice, but does not want anyone (including Alice) to be able to provably attribute the message to him. If Alice can provably attribute that to Cathy, then Bob’s attribution privacy has been violated. Whether or not Cathy knows the content of the message is immaterial to whether Bob’s attribution privacy has been violated.

These examples do not specify if entities other than Cathy know that the message is attributable to Alice. But there are cases where it *does* matter if entities other than Cathy can attribute the message to Alice. For example, Alice sends a legal document to Bob and Cathy. Alice wishes Bob and Cathy to attribute the message to Alice, but not be able to prove this attribution (i.e., provably attribute this message) to anyone other than themselves. Or, Alice sends Bob a document but does not send it to Cathy. Alice wishes that Bob can demonstrate to Cathy that the message’s sender is Alice, but that Bob cannot demonstrate it to anyone else.

In both of these examples, if either Bob or Cathy can demonstrate to anyone else that the message is attributable to Alice, then Alice’s attribution privacy has been violated.

When A’s attribution privacy extends to every entity—that is, the binding of A to a characteristic is kept secret from every entity—then A is said to have *complete anonymity*. So, Alice has complete anonymity under the following condition: when she sends a message to Bob (or anyone else), no one can demonstrate that the message is attributable to Alice. A good example of complete anonymity is a cryptographic voting system such as Scantegrity, in which ballots are posted to the

web in such a way that the person who cast the ballot can verify the ballot is present, but that person cannot prove to anyone else which ballot is theirs (this prevents vote selling as well as protecting the secrecy of the ballot).

Attribution privacy raises several issues for which the implications for attribution system requirements are uncertain.

1) *Form of Association*: There appear to be two alternative ways to associate an entity and a characteristic, for example Alice or Bob with a message.

The association may be based on something that can be shared and once produced is no longer in Alice's control—for example, a notary public's seal, which is under the control of the notary public and not under Alice's control. Specifically, Alice sends a message to Bob, so Bob can now share with Cathy the "something" that provides the attribution back to Alice. Alice may not want the "something" to be shared with Cathy (whether or not the message itself is shared), so Alice's attribution privacy is violated.

The association may be based on a "quality" that cannot be shared without Alice's consent. For example, Bob has the attribution of Alice to a message, but cannot share this attribution with Cathy. Cathy has to trust that Bob is telling the truth when he says that the message is attributed to Alice. Alice could lie and tell Cathy that the message is not associated with Alice. In this case, Cathy has to decide whether to believe Alice or Bob. Alice's attribution privacy is not violated even by Bob's disclosure of the message to Cathy, since whatever "quality" provides the attribution of the message to Alice cannot be shared.

2) *Perfect Selective Attribution and Privacy*: We present as a hypothesis that attribution privacy violations can occur only when perfect selective attribution is desired. If Alice sends a message, and wants everyone to know that the message is attributable to Alice, then there can be no attribution privacy violation. If Alice has perfect anonymity (i.e., Alice's attribution is secret from everyone) then unless there are technical violations, no attribution privacy violation can happen. It is only when Alice does not want Cathy to be able to attribute to Alice the message sent to Bob (and Bob can attribute the message to Alice) that there is the potential for attribution privacy violations.

3) *Default Attribution Baseline*: Does the current baseline for attribution on the Internet make the issue of attribution privacy violations less significant? That is, Alice says that "Cathy can't actually perfectly attribute this message to Alice"—but Cathy replies, "c'mon, we know that from a practical perspective most likely the message came from Alice"?

Stated differently, does the default attribution baseline of the network (if one is specified) shape the practical consequences of the importance of attribution privacy and attribution privacy violations? It is unclear exactly what the default attribution baseline for the existing Internet is, if in fact one exists, but in a future system it seems important to choose whether such a default baseline should exist.

V. CONCLUSION

This work defined attribution as the binding of a characteristic (or data) with an entity (person, process, file, other data, etc). 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. While attribution has typically focused on the identity of the sender, other characteristics (such as the time at which a message entered a particular network) can also be of value as attribution. Thus, a more generalized framework for attribution would provide valuable capabilities for the users of any future network.

Beyond the mechanisms for associating attributes and their values with data looms the larger question of what these mechanisms should provide as a matter of policy. This paper explored the policy requirements for attribution that future networks may require. However, much of the emergence of the attribution framework is outside these specific requirements. A system of governance, answering the question "who decides" should there be conflicts among different actors, is needed. This governance system need not have a central authority, although that certainly is one approach. Social and economic factors will also be powerful factors shaping the attribution system.

Indeed, beyond the technical specification and construction of an attribution framework, there needs to be a focus on how an attribution framework will emerge. The answers to those questions will not be purely technical, and because of that, may represent the most challenging part of the work ahead in launching an attribution system. But ultimately, a framework allowing for attribution in its many forms must be central to a future network design.

VI. ACKNOWLEDGEMENTS

The authors were supported by the GENI Projects Office project 1776 (in turn supported by NSF Grant No. CNS-0940805). Matt Bishop was also supported by NSF Grants No. CCF-0905503 and CNS-1049738. Any opinions, findings, and conclusions or recommendations expressed in this material do not necessarily reflect the views of BBN Technologies, the GENI Project Office, or the National Science Foundation.

REFERENCES

- [1] J. A. Simpson and E. S. C. Weiner, Eds., *The Oxford English Dictionary*, second edition ed. Oxford, UK: Clarendon Press, 1991.
- [2] D. A. Wheeler and G. N. Larsen, "Techniques for cyber attack attribution," Institute for Defense Analysis, Technical Report P-3792, October 2003.
- [3] J. Hunker, R. Hutchison, and J. Marguiles, "Role and challenges for sufficient cyberattack attribution," Institute for Information Infrastructure Protection, Dartmouth College, Hanover, NH, USA, Tech. Rep., January 2008.
- [4] W. T. Strayer, C. E. Jones, I. Castineyra, J. B. Levin, and R. R. Hain, "An integrated architecture for attack attribution," BBN, Cambridge, MA, USA, Technical Report 8384, Dec. 2003.
- [5] M. Bishop, J. Hunker, and C. Gates, "The sisterhood of the traveling packets," in *Proceedings of the 2009 New Security Paradigms Workshop*. New York, NY, USA: ACM, Sep. 2009, pp. 59–70.
- [6] S. L. Stirland, "In test, Canadian ISP splices itself into Google homepage," <http://www.wired.com/threatlevel/2007/12/canadian-isps-p/>, Dec. 2007.