Media Characterization for the Visualization of Secure Paths

Paul DiGioia
Institute for Software Research
University of California, Irvine
Irvine, CA 92697-3425
pdigioia@ics.uci.edu

Paul Dourish
Institute for Software Research
University of California, Irvine
Irvine, CA 92697-3425
pdourish@ics.uci.edu

1. INTRODUCTION

Discussions of security often rely on trust; specifically, who one can trust at any moment. In some sense, then, the notion of security relies on the identification of these trustworthy entities – and, conversely, the singling out of any rogue entities. The ubiquitous nature of both wired and wireless network access brings with it interesting implications for trust. We no longer consider ourselves as being simply online or offline, but might instead speak of being connected with varying degrees of trust. We might consider the *qualities* of the (usually multiple) connections and determine whether one connection is, for instance, more trustworthy when compared to the others. These deterministic qualities might include its speed, its medium (wireless or Ethernet), and its administrative owner. These connection details are often squelched – to the point that design goals, for the most part, strive to make these details transparent to the user. However, we feel that this is inappropriate as a definitive design rule, since ideas of what constitutes a secure connection often vary from one user to the next. Earlier work on the Impromptu project [1] has focused on the use of visual techniques for the purpose of usable security. Our approach is not concerned with making an application more secure in a purely mathematical sense; rather, we focus our efforts on using visual cues in the user interface to aid users in their decision-making process. In this poster, we present our extensions to the Impromptu interface which challenge the notion of connection transparency in the interface. Our design investigates the merits of abolishing this transparency, presenting the details of network paths in the user interface, and, perhaps, causing a bit of suspicion on the part of the user regarding these network qualities.

2. MEDIA CHARACTERIZATION

A natural starting point for the characterization of network qualities is the last-hop connection type. Users generally have a choice between wireless (e.g., WiFi, GPRS) and wired (e.g., Ethernet) for their connection to the backbone network. These connection types are generally considered to be shared mediums, with varying degrees of data visibility. For instance, an eavesdropper might easily collect traffic traveling over a wireless link, but may have a more difficult time with a switched Ethernet connection. In terms of security, then, a user may very well be interested to know the connection type(s) in use.

While it is generally assumed that the user of a system has the ability tell how their own machine is physically connected to the network, the connection method used by *others* in a shared system is usually not apparent. Applications which require the sharing of potentially sensitive information may present a problem to some users who are concerned about eavesdropping – not only at their own location, but also at the last-hop of the recipient's side of the link. Consider a shared workspace where users are able to distribute files amongst each other, as in the Impromptu interface (Figure 1). The files being exchanged may be sensitive – thus, users are naturally concerned about security and wish to ensure that no eavesdropping occurs. Despite the fact that a user may trust each of the other users of the shared system, the threat of undetected eavesdropping remains. Users – even collocated ones – have a vested interest in the connection type used by every other user of the shared system.

For instance, if someone is recklessly connected to the application over an unencrypted wireless link, this may affect the decisions of others regarding whether they choose make any sensitive information available to that user. If the system is set up in such a way that information is communally shared amongst all users of a system, this decision-making process might take a "weakest link" approach. Alice may freely share sensitive information with other (trusted) users of a shared workspace until she notices that Bob is using a connection type that Alice considers to be insecure. At this point, Bob is identified as using the "weakest link," and Alice may decide to remove her sensitive documents altogether.

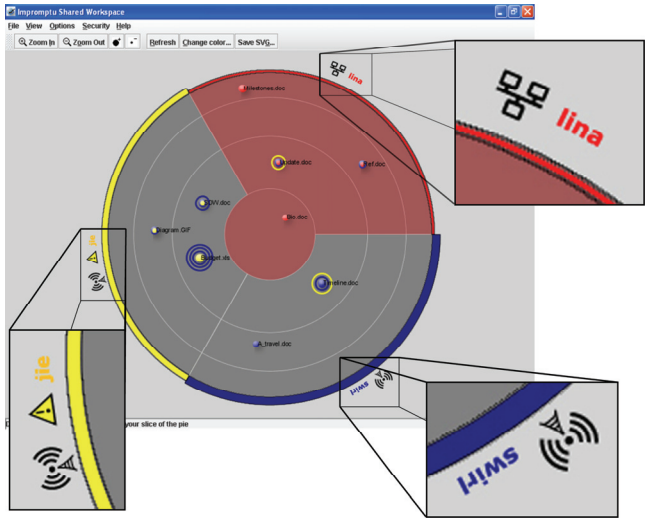

Figure 1 - Media Characterization icons in Impromptu

What is important here is that the type of connection – more particularly, the *conscious decision* of choosing one connection type over another – has repercussions not only for that single user, but for potentially all users of a shared system. By bringing attention to the connection methods of users, we intend to inform users of a shared system about their potential for information leakage. One such notification method might display the connection type of each user of a shared system, as in Figure 1. Here, the action of being connected through a wireless (or wired)

channel is represented by an icon situated adjacent to the username. Identifying the qualities of the various connection types that are employed in a shared system allows users to determine the proper security measures to suit their individual needs; we steer away from any *a priori* determinations of what security should mean to any particular user.

3. PATH CHARACTERIZATION

We now move from an evaluation of the last-hop connection type to the more holistic view of a path or channel between two endpoints. In evaluating the security of a particular channel, there are a number of qualities exhibited by that channel that one might consider. Static attributes of the individual links – such as speed, bandwidth, and administrative ownership of the link – are obvious starting points. These are convenient qualities to begin such an evaluation, as these qualities are immediately available for inspection, and rarely change over time. Freely available tools, such as Jacobson's pathchar [2], are able to identify the type of media (OC12, Ethernet, WiFi, etc.) used by each link in a channel between two nodes. Additionally, the owner or administrator of the links within a path can speak volumes about the security of a channel. As an example, a connection between a client in California and a server at CMU most likely should not contain a router located overseas. For international travelers with a corporate firewall, however, this case may be completely normal. The various organizational boundaries a path may traverse are thus interesting observations in the context of path characterization. Jacobson's traceroute tool reveals the identity of each hop within a channel between two end hosts. Given what we are able to discover about individual links (speed, bandwidth, link type, ownership, number of links), we can begin to think about classifying *channels* as collections of these individual links. We can use the characteristics of the entire path to determine whether or not the path should be trusted.

This classification allows us to detect changes in the channel. The sudden alteration of a channel (i.e., use of a different set of links) may be of interest from a security standpoint. Was a more efficient route discovered? Am I now traversing different organizational boundaries than I was before? Was I manually re-routed through a link within a compromised network? We can imagine being able to flag the changes in a channel on a good-or-bad scale, differentiating between those changes that could potentially affect the security of a channel (in which case the user should be notified) and those that cause no change in security.

While unplanned changes in the communication channel may raise or allay some concerns, it is conversely possible for users to *force* changes in the end-to-end path. Encrypted VPN tunnels created between two points on the network effectively change the topology of the network in such a way that the security of the links over which the VPN tunnel is created is no longer a concern. Thus, the path between two end hosts can be actively transformed by the user; this is, in fact, commonly used as a defense mechanism against untrusted links.

4. AN ARTIFACT-BASED APPROACH

Up until this point, we have focused on the attributes of links and paths in the immediate sense; that is to say, we can determine the attributes of a currently-existing path between two points. This focus is able to answer questions of where is my traffic currently going, or what path my traffic will take if I were to send a packet towards a certain destination. We can then look at the characteristics of the path, and determine whether or not the path should be trusted. That is, does the information about the middlemen in the path affect my decision to make the connection to the endpoint? This evaluation of the links in the path is accomplished by careful inspection of the path itself. Equally interesting, however, is the artifact-based approach, where artifacts which *have traversed* a path are used to tell an accumulated history of the path they have taken. Rather than examining the path itself, here, objects (such as messages, frames, or packets) are inspected for details about the path they have traversed. Switching from channel-based inquiry to artifact-based inquiry allows us to address such questions as what has happened in this case, and, where has this been? Additionally, given where this artifact has been, the user may determine whether it can be trusted.

Email messages serve as a particularly intuitive example of this artifact-based information gathering. Email messages already tag themselves with information about each hop that was taken during their journey from sender to recipient – these tags appear in the "received-from" entries in the header of the message. We can utilize this information to determine whether an email message was sent over a known-good path, or a questionable one, for example. Focus is not limited to the recipient of the message; the sender is equally interested in whether the message was delivered to its intended recipient without any modification.

5. FUTURE WORK

We are currently working on extensions to the Impromptu interface which address the issues detailed in this paper. Specifically, these features integrate various details about the network connections into the user interface. By bringing these details to the foreground of the user experience, our intention is to increase the user's awareness of potentially insecure network qualities, as well as unexpected *changes* to those qualities. Current features are discussed in [3]; additional features described in this poster are pending. In future work, we intend to thoroughly evaluate the features presented here. In particular, we are interested in their utility in everyday situations. Our intention is to examine how these features are employed to make decisions about various connectivity choices or encryption options.

6. REFERENCES

[1] de Paula, R., Ding, X., Dourish, P., Nies, K., Pillet, B., Redmiles, D., Ren, J., Rode, J., Silva Filho, R. Two Experiences Designing for Effective Security. Proceedings of the 2005 Symposium on Usable Privacy and Security. 2005, pp. 25-34. Pittsburgh, Pennsylvania

[2] Jacobson, V. Pathchar – A Tool to Infer Characteristics of Internet Paths. Presented at the Mathematical Sciences Research Institute (MSRI); slides available from ftp://ftp.ee.lbl.gov/pathchar/, April 1997.

[3] Rode, J., Johansson, C., DiGioia, P., Silva Filho, R., Nies, K., Nguyen, D., Ren, J., Dourish, P., Redmiles, D. Seeing Further: Extending Visualization as a Basis for Usable Security. To appear in Proceedings of the 2006 Symposium on Usable Privacy and Security, Pittsburgh, Pennsylvania.