

Securing Spontaneous Interactions

Tim Kindberg
Hewlett-Packard Laboratories, Bristol, UK
timothy@hp.com

Introduction

This paper describes techniques for securely associating devices in ubiquitous computing (“ubicom”) environments. In particular, an important characteristic of ubicom is *spontaneous* device association. Humans carry personal devices with them as they move from one ubicom environment to another. They may also acquire devices that are meant for use by visitors to the environment – e.g. smart whiteboards and pens. Often, devices become really useful only when they work together: e.g. Johnny borrows a smart pen from David and records digital ink onto the personal storage device on his belt; Mary sends a picture from her camera-phone to Robert’s digital picture frame while visiting his house; two teenagers who encounter one another in a shopping mall play a wireless game together with their game-phones or PDAs; two colleagues meet at a conference and transfer a document from one’s PDA to the other’s laptop.

Spontaneous interactions are potentially of great value to users. However, they are not yet a reality. One problem is security. Only a modest level of security is consistent with everyday, spontaneous computing. Nonetheless, security is important and achieving it is non-trivial. Referring back to the examples just given, spontaneous interactions occur in untrustworthy surroundings. The physical surroundings themselves may be untrustworthy – the conference where colleagues exchange a document is attended by business rivals. But even if the surroundings are familiar, they are nonetheless typically within wireless reach from unknown parties – whoever lives near Robert when Mary transfers her image; other people with wireless devices in the shopping mall.

Our research question is: how can we best enable humans to make secure associations between trusted devices in spontaneous circumstances?

Problems with standard approaches

Securing spontaneous interactions is non-trivial because, for spontaneity, the protocols must rely on minimal *a priori* data. Cryptographic techniques will not suffice by themselves. Those can achieve only the secure binding of a cryptographic key to certain electronic identifiers (network addresses, device names). But that cannot satisfy the user of what is actually required in spontaneous circumstances: that *this* physical device is securely associated to *that* physical device. For example, the two teenagers who have just met in the mall in general have no *a priori* knowledge of one another’s device names – which anyway may easily be spoofed or accidentally confused with other device names – and the two users have no cryptographic data in common. If one user network-discovers a device called “N-gage” (apparently the other device’s name), she can only wonder whether what she exchanges a fresh key with is actually the device in the hands of that person over there.

Relatively laborious solutions exist. For example, the users may think of a password or choose a challenge on the spur of the moment and communicate it secretly on a slip of paper or a device screen. But the value chosen is typically drawn from a relatively small set, so that entering it into a small device is not too inconvenient. This trade-off between security and convenience is best avoided.

Physically constrained channels

To overcome the problems identified above, we have utilised *physically constrained channels* in protocols to secure spontaneous interactions, specifically, in the form of lasers [3] and combinations of ultrasound and radio propagation [2]. A physically constrained channel is such that only principals in a certain physical context may transmit a message over the channel – or that only principals in constrained physical circumstances may receive messages over the channel. This is best illustrated by the case of a laser. Consider that one user shines a data-modulated laser light from her device onto the intended target device; and consider that the laser emits no light except onto the sensor where the user points its narrow beam. Given those receiving constraints, we can use the laser channel to send a secret key to the recipient. Conversely, given a send-constrained channel, such as we can create with combinations of ultrasound and radio propagations, we can send authenticating data. In each case, we achieve security with reference only to physical properties – no *a priori* values are required.

Stajano and Anderson [4] were the first to consider physically constrained channels for securing spontaneous associations in the form of direct electrical contact, and Balfanz et al [1] proposed them in parallel with us, in the form of infrared and audio.

Multimedia evidence

Although we have utilised them in our protocols, physically constrained channels require special hardware – even infrared transceivers are increasingly omitted from products – and are often hard to engineer. None of the radiative technologies provides absolute guarantees of physical constraint, because of refraction and reflection and the existence of sensitive receivers that can detect faint signals. Their guarantees are often good enough for everyday computing, but it would be preferable to eliminate such hardware issues. Physical contact would not be “leaky”, but it would require the devices to have special terminals, and to be held in close and sometimes awkward proximity.

More recently, we have begun to devise methods for securing spontaneous associations that exploit only physical indicators with which personal devices are commonly equipped: LEDs and displays, audio output and vibrators. Moreover, we now divide the problem of securing associations so that we need solve only a sub-problem we call *physical validation*, and can rely on well-established techniques for solving the remainder.

A good example of this new approach is the “harmony” protocol. Consider that two devices exchange keys – at least, that is the intention – using the Diffie-Hellman protocol. The harmony protocol answers the following questions: (1) Have these two devices in fact exchanged keys (there is a chance of accidental, if not malicious, mis-association)? (2) Given that the two devices appear to be associated, is there nonetheless a man in the middle?

The two devices play out multimedia streams and the users observe whether the two streams are harmonised. For example, the stream at one device plays a piano part and the stream at the other device plays a bass part. In general, each stream may consist of any perceptible events renderable by the devices concerned. If a user observes harmony between the streams, and if the devices do not indicate a man-in-the-middle attack, then the devices are securely associated.

The harmony protocol operates as follows. One device, the initiator, plays a stream of multimedia events and concurrently sends encrypted commands to the other device to cause it to play what should be harmonised multimedia events. If the receiving device cannot successfully decrypt the events, or if the resultant streams are not harmonised, then there is a failure to associate. The protocol uses traffic analysis to detect a man in the middle, who may be present even when the streams are harmonised. The analysis exploits the fact that a man in the middle has to relay the commands in a timely fashion, if harmony is to be maintained. Thus (under constraints imposed by the protocol) for every legitimate command issued by the initiating device, two packets traverse the wireless network if a man in the middle is present. We assume that the devices use a broadcast network such as 802.11. The receiving device monitors the traffic to detect unencryptable command-packets issued “just before” encryptable commands. If it finds a statistical correlation, it signals a probable man-in-the-middle attack.

Summary and future work

This paper has shown how a range of physical phenomena may be exploited to validate the exchange of keys between devices in spontaneous circumstances; that is, in the absence of *a priori* identifiers or cryptographic data. There are several practical concerns outstanding, to do with the integrity of “physically constrained” channels, and the human factors involved in multimedia comparisons. For the moment, we are concentrating on the question of human factors, while we build prototypes for trials. The outstanding question for the research community is: how should we evaluate and select from variants such as those we have suggested, to provide “good enough” security for what seems to be an important ubiquitous computing requirement?

References

1. Balfanz, D., Smetters, D.K., Stewart, P., and Wong, H.C.: Talking to strangers: authentication in ad-hoc wireless networks. Network and Distributed System Security Symposium; February 2002.
2. Kindberg, T., and Zhang, K. Validating and Securing Spontaneous Associations between Wireless Devices. Proc. 6th Information Security Conference (ISC'03), October 2003.
3. Kindberg, T., and Zhang, K. Secure Spontaneous Device Association. Proc. UbiComp 2003, Seattle, USA, October 2003.
4. Stajano, F., and Anderson, R. The Resurrecting Duckling: Security Issues for Ad-hoc Wireless Networks. In B. Christianson, B. Crispo and M. Roe (Eds.) Security Protocols. 7th International Workshop Proceedings, Lecture Notes in Computer Science, Springer-Verlag, 1999.