Putting the Human Back in the Protocol
(Transcript of Discussion)

Bruce Christianson

University of Hertfordshire

Hello, everyone, and welcome to the 14th International Security Protocols Workshop. I'm going to start with a quotation from someone who, at least in principle, is in charge of a very different security community than ours:

> Our enemies are innovative and resourceful, and so are we. They never stop thinking about new ways to harm our country and our people, and neither do we.

It occurs to me that if we replace the word "country" by the word "system", and the word "people" by the word "users", then we have a pretty fair description of the current state in our own little security community, the security protocols world.

Our theme this year (may I have the envelope please, Mildred) is "Putting the Human back in the Protocol". We're very used to saying, almost as an afterthought, "Of course, Alice and Bob aren't really people. Really Alice and Bob are programs running in some computer environment." But computer systems exist — the reason we build them is — in order to enable people to interact in accordance with certain social protocols. This means if we're serious about system services being end-to-end (and if we're not, then we can spare the rest of the world an awful lot of what we do) then, at least at some level of abstraction, the end points, Alice and Bob, are human. This has certain consequences.

We're also very fond of saying that when we say, "believes" — for example when we say that Alice believes a particular key is fresh, Alice believes that what Bob says is true — we don't really *mean* believes, we mean squiggle, where squiggle is some mathematical predicate that satisfies certain axioms[1]. But, of course, if the end-points are human, then squiggle does map on to certain beliefs that are actually held by humans. The problem is that currently they're not beliefs about anything useful, indeed they're not the kind of belief that a rational person would willingly hold consciously for any great length of time. The dependent problem from that is that when we invite users to participate in security protocols, they are inevitably the weak link. Usually this is attributed to the fact that humans are untrustworthy, unreliable, and unable to do cryptography in their heads. We've all read Kevin Mitnick's book[2].

[1] Such as KD45.

[2] The Art of Deception, Wiley, 2003.

B. Christianson et al. (Eds.): Security Protocols 2006, LNCS 5087, pp. 1–2, 2009.

The trouble is, in ad hoc environments, and particularly in the context of pervasive computing, very frequently the human element of the system is the only one with any real understanding of what the security requirement is. I'm not saying the human knows what's actually going on, but they have some idea of what should be happening, and perhaps more importantly, of what shouldn't. The difficulty is that when we try to program systems to interact with humans by popping up a box which says, do you want to accept this certificate or not, we're asking the wrong question. We're asking questions in terms of the abstractions which we currently use to explain and analyse protocols, and humans are well-known to be very bad at this kind of logical thinking.

But suppose that you take a logical puzzle — does conclusion C follow from premises A and B — and rephrase it in a context where there's some kind of transaction, and the question at the end becomes, is Alice treating Bob fairly or is she cheating him? Now it turns out humans are very good at solving problems posed in this form, even humans who have not spent many years doing post-graduate courses in computer science and philosophy. So the question whether humans are really the weak point is something that I think we should re-examine.

A common objection is also that humans can't do cryptography, but with the increase of personal devices, and the emergence of the pervasive, ubiquitous computing environment, humans are typically now surrounded by a cloud of little devices that can do cryptography perfectly well, and with whom they have an extremely intimate relationship. For "personal" think "unshared", or at least potentially not shared with anybody not trusted.

So the question isn't just, how can we put the human back in the system? This isn't an HCI problem, it's not something that interface people can deal with. The question is, how can we put the human right back in the protocol, how can we align the interests of the human with the protocols that serve them, at all levels. Currently if you look at the deployment of middleware, it's a major research question to try and identify whether two end-points are the same end-point at different levels or not. That's definitely the wrong question to be asking. Perhaps we should instead devote more energy to determine whether Alice is talking to the correct stranger, as my student Jun Li rather nicely puts it[3].

This is a workshop and not a conference. The rules are similar to those of a Quaker meeting, it's OK to interrupt the person who's giving their testimony, but please make sure that your motives are pure, or will at least bear peer scrutiny. We have a few more PhD students than usual, including some from far-off exotic places, which is very nice to see. Please participate, and don't worry if you get it wrong.

These workshops usually descend into chaos at some point, so this year we decided we'd try and just get it over with. Accordingly Matt Blaze has very kindly volunteered to be our first speaker. [Laughter]

[3] Towards a "Localization of Trust" Framework for Pervasive Environments, PhD Thesis, University of Hertfordshire, 2008.