

Authenticating Ubiquitous Services: A Study of Wireless Hotspot Access

Tim Kindberg

HP Labs

Long Down Avenue, Bristol BS6 5ND, UK

Tim.Kindberg@hp.com

Chris Bevan, Eamonn O'Neill, James Mitchell,

Jim Grimmett and Dawn Woodgate

University of Bath, Bath BA2 7AY, UK

{C.R.Bevan, E.O.Neill, J.R.Mitchell,

J.W.Grimmett, D.Woodgate}@bath.ac.uk

ABSTRACT

This paper concerns the problem of phishing attacks in ubiquitous computing environments. The embedding of ubiquitous services into our everyday environments may make fake services seem plausible but it also enables us to authenticate them with respect to those environments. We propose physical and virtual linkage as two types of authenticating evidence in ubiquitous environments and two protocols based on them. We describe an experiment to test hypotheses concerning user responses to physical and virtual linkage with respect to fake Wi-Fi hotspots. Based on our experience we derive an improved protocol for authenticating spontaneously accessed ubiquitous services.

Author Keywords

Phishing, ubiquitous services, Wi-Fi, authentication.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g. HCI): Miscellaneous. H.1.2 [User/Machine Systems]: Software Psychology. D.4.6 Security and protection.

General Terms

Experimentation, Human Factors, Security.

INTRODUCTION

Phishing attacks, in which the attacker dupes a user into exposing personal data by masquerading as another party, are familiar to us from email and the web. They can be mounted effectively [4]. Ubiquitous computing, in which services are embedded in our everyday world, brings its own phishing problems. This occurs where a service seems to originate from a situation or environment that users are likely to trust, but in fact originates from an attacker. In this ubiquitous case, there are two respects in which the

apparent origin of the service is manifested. First, as with email or web phishing, the attacker makes the service's *content* match the user's expectations of whatever party they are likely to trust. Secondly, a new aspect of the ubiquitous phishing attack is that the *embedding* of the services in their environment may seem to be correlated with a trusted setting [8].

In a recent real-life example of a ubiquitous phishing attack [2], official-looking "traffic violation tickets" were placed on the windscreens of parked cars. The tickets gave a URL of a site which they claimed had photographic evidence of the "violation". Users who browsed to this site were subjected to a malware attack. Although most users will have browsed to the site later using their PCs, in the near future this could be even more embedded in the situation where the ticket was encountered, with the ticket acting as a "physical hyperlink", e.g. by bearing a barcode or NFC chip, and the "parking violation service" being accessed on a mobile device when the ticket was discovered. Arguably, the embedding of the ticket in the parking situation added to the attack's effectiveness compared to an email conveying the same information: whatever a person's beliefs about whether they were illegally parked, the presence of the ticket in the parking situation and the claim of a violation convinced some people to check the malevolent service.

A second example of the ubiquitous phishing problem involves a more routine example of a ubiquitous service: public Wi-Fi (802.11) network provisioning. It is a *de facto* ubiquitous service, since its limited physical reach embeds it into physical settings such as cafés and it is in widespread use. It is also subject to phishing attacks, not just in principle but in practice [3]. In this paper we investigate the efficacy of methods for preventing Wi-Fi phishing attacks. The following two forms of attack can be mounted:

Simple spoofing. In this case, the attacker has a laptop or other mobile device with a Wi-Fi interface but no (other) network connection to the open internet. She sets her laptop's wireless into infrastructure mode with an SSID designed not to arouse suspicion (or to be similar to an existing nearby access point's SSID such as the café's SSID) and acts as an access point. When a victim tries to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UbiComp 2009, Sep 30 – Oct 3, 2009, Orlando, Florida, USA.

Copyright 2009 ACM 978-1-60558-431-7/09/09...\$10.00.

connect to this fake access point, the attacker can install malware and attempt to obtain personal information such as credit card details. However, the attacker cannot connect simultaneously to the open internet and so cannot provide this connection to the victim.

Man in the middle. This attack adds to simple spoofing a simultaneous connection to the open internet, which may be via the genuine Wi-Fi network or could be, for example, via a 3G cellular interface. The attacker relays all traffic between the victim's machine and the internet. The attacker can eavesdrop and interfere with traffic without the victim being aware, unless it is protected by end-to-end encryption. That may not be so if the victim mistakenly accepts invalid SSL certificates used by the attacker.

These types of phishing attacks extend to ubiquitous services in general. As with public Wi-Fi access, a key characteristic of ubiquitous service access is its spontaneity [7]. Ubiquitous services will be embedded in many environments in our everyday world and in many cases will be made accessible to visiting users. It is not practical or desirable to expect those users always to be pre-configured with keys or certificates needed to ensure secure, authenticated service access. New ways have to be found of conveniently bootstrapping security at the time of spontaneous access without relying on trusted third parties, which will not exist or be reachable in many cases.

In this paper we present protocols for checking the authenticity of a spontaneously accessed Wi-Fi network using two types of evidence: *physical linkage* through a physical artefact in the environment where the service is embedded, and *virtual linkage* in the form of an interactive protocol that involves the user's device and information presented via the physically linked artefact. Through an experimental evaluation, we study users' perceptions of the strengths of the different types of linkage and their authentication value. Furthermore, we give an account of the evidence and rationales that users brought to bear in their evaluations. Using lessons learned from this analysis, we propose a more effective combination of physical and virtual linkage for authenticating spontaneous wireless access services.

RELATED RESEARCH

The research community has explored phishing attacks in recent years. In 'Why Phishing Works' [4], Dhamjia et al presented the principal design mechanisms that are exploited by phishers in order to mount a successful attack. A variety of visual design elements were examined, including browser-based security indications that were specifically designed to indicate potential fraudulence. Many such elements were found to be ignored or misinterpreted – even by participants who were reported to be highly I.T. competent. More recently, but in a more ubiquitous setting, users' trust behaviour in a Wi-Fi 'hotspot' was investigated [8]. In that study, visual design elements within the body of a website were explored, but

with the added manipulation of 'locativity' – the degree of relevance of an image or other design element to the immediate physical environment of the user. It was found that presentation of a clearly non-local design element significantly reduced user acceptance. This suggests that users are sensitive to material that is relevant to the embedding of a ubiquitous service, i.e. to its authenticity.

Embeddedness is also addressed by Riegelsberger [11] and Fogg [5], who suggest that the perceived distance (in time and space) between the parties is significant to trust-based transactions. They refer to this distance as degrees of 'disembeddedness'. As it increases, the perceived risk of opportunism rises. Fogg [5] suggests that providing a link between a website and the physical 'bricks and mortar' from which it originates ('reembedding') is a useful way to reduce this distance and increase trust.

On the systems side, the closest problem to Wi-Fi authentication that has been tackled is creating a secure spontaneous connection between two devices. The two devices are authenticated to one another, in a way that involves some kind of physical evidence of authenticity. Many proposed solutions use a two-stage process in which:

1. devices establish a shared secret key over the insecure broadband wireless channel; secret keys are generated so as to be unique to the pair of principals, e.g. as a result of using Diffie-Hellman key exchange.
2. the resulting shared secret key on each device are compared using a more secure auxiliary channel, in order to expose any man in the middle.

Since the keys may be very large, a secure hash of the key may be used instead. Variants of the Short Authenticated String (SAS) [16] and MANA [6] protocols have been used for authentication, both of which reduce the amount of information compared. This has benefits where users perform the comparison themselves or where the information transmission capacity of the devices is limited. However, the values that the user compares are random digit-strings chosen by the system, and we are investigating the effects on perceptions of security of comparing intelligible information chosen by users.

The auxiliary channels employed are typically physically constrained by range and direction, often requiring line-of-sight, auditory perception or short-range infrared communications [1]. The act of comparison may be carried out by the user or by the devices themselves. Examples of artefacts used in authentication include 2D barcodes [10] and coloured lights [13] – which Roth et al. used to authenticate wireless access points. A more extensive survey of device pairing techniques is given in [14].

Uzun et al [15] conducted a usability study, examining error rates, effectiveness and user perception of a range of user interactions in secure device pairing schemes. The schemes involved variations on comparing check codes between the devices or copying check codes between the devices. They



Figure 1. Bertorelli's café showing display running Interlock and leaflets on tables.

found interaction effects between complexity, error rates and perceptions of security. When examining user perceptions of the two schemes with the lowest error rates, the comparison-based method was considered to be “easy” but less “secure” compared to the copy-based method, which users described as “hard but professional” and the method they would most like to use for pairing devices.

EVIDENCE, PROTOCOLS AND HYPOTHESES

In this paper we concentrate on Wi-Fi phishing attacks. We assume that venues secure their access points with strong passwords so that they are indeed trustworthy. However, even if a connection is made to a bona fide access point, an attacker may also attempt to eavesdrop and inject malicious responses. But those attacks can be defeated by establishing a secret key with the bona fide access point and using it to encrypt and authenticate all communications. The most powerful of the authentication protocols below (Interlock) achieves that.

To defeat Wi-Fi phishing attacks, our goal is to provide users with salient evidence that allows them to discriminate between a Wi-Fi network that is genuinely provided by an establishment such as a café and one that is not. We are investigating evidence of authenticity that takes the form of physical phenomena rather than purely cryptographic properties. One consideration is how much a user trusts the café as a provider of a network for some particular transaction type over Wi-Fi (for example, credit card purchases). However, within the scope of this paper our focus is on authenticity of the network, i.e. with how much confidence does the user believe that a particular Wi-Fi source (which they identify at first simply by its SSID) is provided by the café or is at least sanctioned by it?

Our research concerns the value of different types of evidence for authenticity with respect to four questions:

1. What is the actual value of the evidence for evaluating authenticity?

2. How convincing do users find the evidence, and what is the relationship between this and the answers to (1)?
3. What factors concerning the user's understanding and dispositions affect their answers to (2)?
4. How usable are protocols that involve the evidence?

To motivate the types of evidence under consideration, we first describe two protocols that we have devised and implemented for authenticating Wi-Fi networks: Interlock and Synchronisation. They both operate between a client program called a supplicant on the user's mobile device (which is secure against malware attacks) and a large display attached to the wall of the café (see Figure 1). The wall display is securely connected by cable to a computer that is securely maintained by the café. Both protocols assume that the user has first selected and connected to a network. Neither the user nor the system has been authenticated to the other. The user does not have to bring any type of credentials to the premises, since the network is public.

Interlock protocol. This is a variation on an earlier protocol in which two people securely associate their devices [9]. That in turn is based on the interlock protocol of Rivest and Shamir [12]. First, the user's device makes an encrypted 802.11 connection to the access point. We have implemented this using simplified EAP-TTLS authentication, allowing any user to join the network by establishing an ephemeral Diffie-Hellman session key with the access point. Once the link-layer negotiation is complete, the user's connection is encrypted with this session key, but no authentication of either the user or network has taken place. The user, Alice, interacts via the supplicant on her laptop with an avatar on the café's display. Alice chooses one of a set of N phrases offered on the laptop, and is asked to transmit the phrase in two stages (halves) to the avatar on the display but, importantly, not to transmit the second half until the avatar indicates on the display that it has received the first half. When the avatar receives the second half, it presents the complete phrase on the café's display. Alice is asked to compare it with the phrase she sent from her laptop. If they differ, then a man in the middle has been detected and the network is fake.

The way this works is that the underlying protocol encrypts the phrase with a key derived from the link layer key, splits the ciphertext into two parts – the first containing the first half of each block, and the second containing the second half of each block – and sends them in the two stages known to Alice. A man in the middle, Mallory, cannot decipher the first half of the ciphertext by itself, since decryption requires complete ciphertext blocks. Yet he must forward a message so that the avatar will acknowledge receipt. By the nature of Diffie-Hellman key exchange, the key shared with the genuine access point is different from the one shared with Alice. So Mallory cannot simply forward the ciphertext. He must guess Alice's message and encrypt that with the key shared with the genuine access

point. Even if Mallory knows all the N possible phrases that Alice could have chosen, his chance of going undetected is $1/N$. This can be made negligible for a suitably large value of N . Moreover, N is made sufficiently large so that Mallory does not have time to encrypt all N phrases with the one-time key, so that he cannot infer the phrase by comparing the ciphertexts.

Synchronisation protocol. This protocol displays a dynamic image that changes simultaneously on the user’s laptop display and the café’s display. In our particular implementation, the image is of up to nine coffee cups randomly arranged in a 3x3 array. Each image lasts for a random interval between 1 and 5 seconds. Alice is asked to compare the sequences of images between the laptop and display and to verify that they are synchronised, i.e. change in lock-step. If they are not synchronised, then the network is not genuinely provided by the café. This protocol is weaker than Interlock: if the image sequences on the two displays are synchronised, then the network is highly unlikely to be simply spoofed, since the attacker would have to capture the sequence from the café’s display on a hidden mobile camera and play the image sequence back to Alice on her laptop. A video of the image sequence would be distinguishable from the original by its lack of fidelity. Also, image processing, with the goal of synthesising the original image sequence, can be made infeasible: in a full implementation, the coffee cups would be replaced by unpredictable images drawn from photo-sharing sites. On the other hand, a man in the middle could forward the images from the genuine café network over the fake network to Alice, and so go undetected.

The two protocols provide two types of evidence of authenticity. One is how the evidence relates in a simple physical sense to the service’s setting via some artefact in that setting (in this case a large display fixed to the wall), and the other is how that artefact is involved in the user’s interaction when authenticating the network. We refer to these types of evidence as *physical linkage* and *virtual linkage* respectively:

Physical linkage. The physical artefact that is associated with the premises (e.g. the café) and the physical circumstances of that association.

Virtual linkage. The type of user interaction (and underlying protocol) involving the user’s device and information provided by the physically linked artefact.

In Interlock and Synchronisation, the physical linkage is constant: it is the display fixed to the wall of the café. However, the user interaction and underlying protocols differ between these two protocols. Not only does the user engage in different steps, but the protocols yield different actual security values, as we have explained: only Interlock has value against a man in the middle attack.

These concepts led us to construct three further configurations for this study that would enable us to

	Virtual linkage		
	Password	Synchronisation (cups)	Interlock (avatar)
Physical linkage	Leaflet	n/a	n/a
	Poster	n/a	n/a
	Display	Display	Display

Table 1. Authentication conditions for Bertorelli’s Wi-Fi investigate our research questions by varying physical linkage as well as virtual linkage. We introduced two further physically linked artefacts: leaflets distributed on the café tables; and a poster attached to the wall near the display. We also introduced a *Password* protocol that operates on both the paper-based artefacts (leaflets and poster) and the display:

Password protocol. The artefact provides a password. Alice is asked to enter the password into the supplicant on her laptop. If it is not accepted, then the network is not genuine. This protocol has no actual security value, since any spoofed web page can be programmed to accept any password and respond as though it were genuine. However, it enables us to test how users reason about types of linkage and evidence by comparison with the other two protocols.

The five combinations of physical and virtual linkage are shown in Table 1. Only the Password protocol can be run using the leaflet and poster but it can also be run with the display. The display further supports the Synchronisation and Interlock protocols. It is the only one of the three artefacts that can do so, since they are dynamic.

In our study, the configurations using the Password protocol increased in the degree of physical linkage: the leaflets were not attached to any part of the café; the poster was attached to the wall professionally but without a frame, hence less robustly than the display which is firmly bolted to the wall and cabled. The configurations on the display increased in virtual linkage: the Password protocol is weaker than the Synchronisation protocol, which is weaker than the Interlock protocol. A portable screen-based device such as a tablet PC could also have supported a variety of protocols in a comparable format to the leaflet, but we did not explore this configuration for two reasons. First, we would risk confounding the evaluation of different types of virtual linkage with likely doubts about physical linkage. Secondly, it is implausible that a café would employ a free-floating but expensive device, whereas cafés are more likely to have large fixed displays. All three artefacts and the web pages shown to the user were designed and branded to a common template of a professional standard.

We hypothesised that users would *perceive* greater degrees of physical and virtual linkage to imply greater authenticity:

Hypothesis 1: Increasing physical linkage \Rightarrow greater confidence in authenticity

Hypothesis 2: Increasing virtual linkage \Rightarrow greater confidence in authenticity

We conducted an experiment to test these hypotheses, gathering data concerning the users' perceptions of and reasoning about the different types of evidence provided and the other considerations that users brought to bear.

METHODS

The study was conducted within an actual café on the campus of the University of Bath, U.K. We named our establishment *Bertorelli's*, a fictional café which we realised temporarily within the existing café (see Figure 1). Using the café provided some ecological validity, particularly with respect to spatial layout and the artefacts such as tables and a bar that were present. However, we acknowledge that this validity is limited, especially as, for practical reasons, the study was conducted while the café was closed to the public. We used the three types of artefact mentioned in the previous section: a 42" display that was already fixed to the wall of the café; a poster of comparable size (A1) that we fixed to the wall near the display; A5 size leaflets placed on tabletops throughout the café. All existing branding on the wall and tables was replaced with our own professional "*Bertorelli's*" branding.

In addition to the five conditions of Table 1, we added a sixth, Direct Connection, as a control condition. This condition involved neither an artefact nor a protocol: after connection to the network, the user was presented with a *Bertorelli's* web page that simply invited her to click a button to "connect to the internet".

We implemented six networks, one per condition, and provided a Windows laptop for the participants to connect to them. We implemented the networks with separate access points, and gave them neutral SSIDs *Bertorellis1*, *Bertorellis2*, ... *Bertorellis6*. No other networks were present. We supported the usual public Wi-Fi connection procedure using a "captive portal", i.e. a web page that the user is initially forced to visit. That is, after selecting and connecting to any network from the list of "available wireless networks" and attempting to browse, the user was redirected to *Bertorelli's* captive portal giving instructions appropriate to the corresponding condition. This would be insecure in an actual implementation, since an attacker could download malware via the captive portal; but we did not want to distract the participants with the unfamiliar 802.1X supplicant software. All six networks behaved as we have described in terms of the protocols used, except that we skipped the key-establishment step for Interlock. In a deployment of Interlock, that step would be transparent to the user. It is also worth noting that, as with many other public Wi-Fi networks, no encryption keys were associated with any of the networks at the point of connection. Windows refers to such networks as "unsecured".

Participants

There were 28 participants (20 male, 8 female, age range 18-70 years, mean = 25), recruited by email. Half the participants were sourced from staff and students within the university. The other half consisted of a diverse group of non-student local people.

Procedure

The study was carried out with individual participants. We began by compensating for different levels of knowledge of Wi-Fi networks. The participant was given an instruction sheet which included a short explanation of the potential threats involved in unsecured Wi-Fi use, including a man-in-the-middle attack. An experimenter then instructed the participant in how to connect to a network until she was comfortable doing this by herself. The experiment then proceeded in two phases, which were recorded on video.

Phase 1. The participant was informed that only one of the six networks genuinely belonged to *Bertorelli's*, and was asked to evaluate the authenticity of each. She connected to each of the six networks and followed the corresponding instructions in turn, in pseudo-random counterbalanced order. The experimenter encouraged the participant to discuss aloud her thoughts regarding the authenticity of each of the wireless networks as she worked through them.

Phase 2. A semi-structured interview was conducted with the participant. The first question involved rank-ordering the six networks on a scale of authenticity. The remaining questions concerned other aspects of the security of the network but are out of scope for this paper. Participants used a graphical interface to make their rankings by dragging icons that represented each of the networks around a computer display: both absolute and relative scores were thus captured. Again, the experimenter encouraged the participants to verbalise their reasoning.

We transcribed the video tapes for all but one of the participants (one was mislaid). We categorised participant statements with respect to physical and virtual linkage and other factors discussed below. For this purpose, the authors independently proposed coding schemes based on salient phenomena in the transcripts. The results were analysed and converged for coverage and consistency, through several iterations. Finally, each transcript was independently coded by two of the authors using the resulting scheme. Any inconsistencies were then resolved through further discussions.

RESULTS AND DISCUSSION

We begin by examining our hypotheses using the quantitative data on authenticity scores. Then we give an account of the qualitative data from the users' comments, and discuss the relationship between the two.

Testing our Hypotheses

Immediately after each experimental condition in Phase 1, participants were asked, with responses of 1-6 on a Likert




Figure 2. Normalised mean scores for Genuine (Single-condition and Cross-condition).

scale: “Was this wireless network provided by Bertorelli’s Café? (fake – genuine)?”. At the beginning of Phase 2, the participants were asked: “Using the terms genuine and fake, how would you rank the access points you have tried today?”, thereby evaluating all six conditions at the same time. Our main quantitative measure of authenticity is the responses to these questions. Looking at both sets of results allows us to see if our participants’ statements are consistent, i.e. whether they changed their opinion of a particular condition later, having seen all six.

Single- and Cross-condition results are shown in Figure 2, where the results have been normalised for comparison. The Single scores were on a six-point Likert scale. The Cross scores were on a scale of 0-1300, which the participants used for free-form comparison by manipulating icons representing the networks to obtain the desired ranking along a horizontal axis of 1300 pixels. Any two icons that overlapped in the free-form comparison were adjusted to be equal at their mid-point.

The Single-condition scores were analysed across all six conditions using a one-way repeated measures ANOVA. There was a significant effect for network authentication type [Wilks’ Lambda = .25, $F(5,23) = 13.619$, $p < 0.01$]. A multivariate partial eta squared value of .75 suggests a large effect size. Similarly, for the Cross-condition scores, ANOVA showed that there was a significant effect for network authentication type [Wilks’ Lambda = .27, $F(5,23) = 12.70$, $p < 0.01$]. A multivariate partial eta squared value of .73 suggests a large effect size.

Pairwise comparisons (repeated measures t-tests) were performed on the scores between successive condition pairs in which first physical linkage and then virtual linkage increases. The results are consistent between the Single-condition and Cross-condition analyses. Our hypotheses were that confidence in authenticity increases with *de facto* physical and virtual linkage. The experimental results suggest that these hypotheses held initially for both physical and virtual linkage. There is a significant pairwise effect in the participants’ scores between the Leaflet and Poster conditions, reflecting increased physical linkage. Similarly,

there is a significant effect between the Password/display and Synchronisation/display conditions, reflecting increased virtual linkage. However, there was no significant difference in the Password protocols on the poster and display. Similarly, there was no significant difference between the Synchronisation and Interlock protocols. We shall refer to these n.s. comparisons as “plateau” effects. We now examine these results in the light of our qualitative data.

Analysis of Participants’ Rationales

In this section we examine the data from the coded transcripts in order to answer the following questions: (1) how well do the concepts of physical and virtual linkage, as defined above, cover the types of evidence that the participants actually adduced in their evaluations of authenticity, and what specific types of linkage did they mention?; (2) What is the explanation for the plateau effect observed in the quantitative data?; (3) What other emerging factors are relevant to solving the ubiquitous phishing problem?

In addition to coding references to physical and virtual linkage, we coded for other factors of possible significance to understanding our participants’ backgrounds, perceptions and rationales. We coded the remarks of all 27 of the 28 participants for whom we had transcripts. References to “all participants” in the following analysis refer to these 27.

Physical and Virtual Linkage

Variations on both physical and virtual linkage were frequently mentioned by the participants. Taking physical linkage first, all but three of the participants felt that some aspect of the physical circumstances of the association between the artefact and the café was relevant to authenticity, for at least one condition. These aspects fell into four categories: physical attachment, legitimacy, visibility and accessibility.

Physical attachment is the most basic meaning that falls under physical linkage: how firmly is the artefact attached to some part of the café? Increased degrees of physical attachment were taken as more convincingly authentic. Examples of typical remarks were that *the screen looks bolted to the wall, physically bolted to the fabric of the building*, whereas the leaflet was *just floating around*. The relationship between the poster and the display was more equivocal. Several observed that the poster, like the display, was *fixed*, although one pointed out that it didn’t have a frame – unlike another poster in the café. Several also noted the cabling of the display.

Legitimacy refers to the look-and-feel of the artefacts, specifically whether they were “legitimate” with respect to provision by the café. Words such as “legal”, “official” and “branded” featured. As one said, *once again it’s branded and kinda looks consistent with everything else in the café*. Sometimes this convinced the participant, as in the one who said of the leaflets *yeah ... that looks official, that looks*

right. But consistent look-and-feel was generally taken to be a minimal requirement rather than proof of authenticity: *Well, if it wasn't the genuine one, it still could have said it was Bertorellis couldn't it?*

Visibility is that of the artefact itself, and whether a perpetrator might be noticed when installing it, particularly by the management and employees. The display was rated highly in this respect: *any kind of place wouldn't ... have screens and stuff that were interacting with computers and things without staff and management being very aware of it ... you couldn't scam that.* The poster was *more noticeable than the flyers because it's on the wall, and it's quite large, which makes it similar to the screen [in authenticity].* That remark was also typical of what participants said about the leaflets/flyers: *[they] could be surreptitiously slipped in.*

Accessibility refers to whether a supposed attacker had access to the part of the café where the artefact lay. Some observed that this rendered the leaflet, which was left on public tables, less valuable as evidence of authenticity than the poster or display: *anyone can go round and put in stuff.* However, it was not clear that the poster and display, near one another and with a door and the bar a few metres on either side, were felt to be in a space that was wholly restricted to staff only. One participant referred to the bar by contrast: *If they had a screen up, and the screen was clearly behind the bar, or somewhere where it is clearly belonging to Bertorelli's.*

With respect to virtual linkage, every participant mentioned some aspect of the interaction between the laptop and the artefact, for at least one condition. These references fell into three main categories: generalized references to interactivity, and two specialisations of this concept, synchronisation and causality.

Interactivity, or equivalently “interaction”, “information exchange”, “transaction” or “dialogue”, was a term used by about half the participants for a protocol between the display and the laptop. We include here non-specific mentions of whether there was any interactivity and, if so, how complex it was. Interactivity was generally referred to as providing evidence for authenticity. As one participant put it when making a specific comparison, *because [the poster] doesn't provide any interactivity, I wouldn't trust it as much as the screen method.* Another wanted *some sort of interaction which gives credence to it being a legitimate network.* Nine participants mentioned the complexity of interactivity as indicative of authenticity: *because it's a pretty complex interaction, you tend to think it was the real network.*

Synchronisation refers to specific mention of what happens exclusively in the protocol of that name: the agreement in time and content between the display and the laptop. About half the participants referred to this as evidence of authenticity. For example, one said: *cos it's changing quite quickly and it's staying exactly in sync ... I think it shows that you are directly connected to the thing.*

This was one of the few factors that brought out understanding of the man-in-the-middle threat: *potentially the information is being broadcast on two channels simultaneously and all I have to say is “yes it looks the same”.*

Causality refers to specific mention of cause and effect within interactivity, which was mentioned in both directions between the user and the display. Sometimes participants referred to feedback from the display: *I knew I was connected [to the right network] when I saw the results.* And *[display-based protocols] feel more secure because there's more feedback.* Others were particularly interested in whether they were prompted from the display: *doesn't seem like it's been verified by them if I don't have to follow some kind of instruction.* Whether the data involved in the dialogue was personal to the participant was also an important factor for some: *The message I chose was displayed [on the display], so I'm definitely connected to the right router.* And conversely: *but [password on display] didn't even give me a chance to contribute myself ... so I'm not liking this one.*

The participants considered physical and virtual linkage as relative rather than absolute evidence for authenticity. None said attacks against any of them were impossible. Rather, they weighed the difficulty of faking them. All but three participants mentioned the ease or difficulty of particular attacks. Broadly speaking, ease was taken to be indicative of dubiousness, and difficulty was taken to be evidence for authenticity. Sometimes this was a question of resources or effort. For example, printing and distributing leaflets was taken to be easier and less expensive than printing and mounting a poster. Other times the estimation of difficulty was expressed in terms of probability. That particularly applied to the Synchronisation protocol (the difficulty of getting the timing just right) and the Interlock protocol (the difficulty of guessing the chosen phrase). In very few cases were the participants estimating feasibility in a technical sense; rather, they appealed to factors that were familiar to them, such as the costs of printing and the need to choose unguessable passwords.

Reasoning about the Plateau Conditions

We now consider, in the light of our analysis of participants' rationales, explanations for the equivalence between the Password protocol run with the poster and with the café's display and, on the other hand, the Synchronisation and Interlock protocols on the display.

Password/Poster and Password/Display. The coded transcripts provide little explicit evidence for the perceived similarity between these conditions. Only five participants expressly compared them and in all cases they said that they were the same, without giving a clear reason.

However, it is straightforward to explain the perceived similarity based on the breakdown of participants' perceptions of physical linkage from the analysis of the

transcripts. Taking the four aspects of physical linkage in turn, we begin with physical attachment. The display was more firmly attached to the wall than the poster, although arguably not by much, especially compared to the difference between the leaflet and the poster. With respect to legitimacy, the poster and display were identically branded and in more or less the same location on the café's wall. Their location made them equivalent in both other aspects of physical linkage: visibility and accessibility.

Interlock and Synchronisation. The similarity in rankings between these two protocols is borne out by two marked effects in the participants' rationales: an inability to discriminate between two highly unfamiliar protocols and, to a lesser extent, a perception that these password-free protocols were equivalently insecure in respect of data protection as opposed to authenticity (see below).

About half the participants made statements to the effect that the two protocols, both run over the display, were similar in some sense. These people fell into two broad classes: participants who appealed to their intuition, and those who argued from an idea of the protocols' complexity. Taking the intuitive response first, this seemed to stem from an inability to discriminate between them. Typical statements were *I'm going back to these two ... dunno why ... I just like them* and *[these two conditions] make it harder to break in, and [these two conditions] are ok, the rest no* Others picked on complexity as a common factor between the protocols: they felt that both these conditions were too complex to fake, but again were unable to discriminate between the two. For example, *[Interlock and Synchronisation] are the ones I completely trust. It would take a lot of resources to fake them. And: They are different. They look like they are trying hard to make it secure. That's my impression.* But this same complexity made them opaque to some: *[they] were a bit more confusing because you had to think a bit.*

Several participants picked up on a common absence of what they termed 'security' between these two protocols. One specifically distinguished security from authenticity: *If it was [Interlock or Synchronisation], ... I know it's provided by the location, but I know it's not secure.* Another said: *I suppose [Interlock is] similar to [Synchronisation]. I don't think it's any more secure for being able to list a number of questions and a number of answers.* This point is picked up below: some participants may have marked both these password-free protocols down, on the grounds that they considered them insecure rather than inauthentic.

Other factors in participants' assessments

Physical and virtual linkage accounted for many of the participants' utterances as they explained the reasons for their ranking scores with respect to authenticity. This is not surprising, in that the only evidence offered to the participants fell into a combination of these categories of linkage, by definition. What was not obvious before the

experiment is the different aspects of linkage that the participants mentioned, as in the sub-categories that we have described.

Three other factors emerged from the participants' deliberations that deserve separate discussion here: the unfamiliarity of the protocols; perceptions of the use of passwords; and the participants' overall conception of security as opposed to authenticity in particular.

Unfamiliarity. Sixteen participants used words such as "strange", "novel" and "weird" to describe one or more of the protocols. In particular the participants often remarked upon the unfamiliarity of the Synchronisation and Interlock protocols. They are unlike anything the participants had experienced before.

Passwords. The Password protocol also sometimes caused consternation. Many remarked that it is the opposite of how they normally expect to use passwords: the password is public rather than private, and is not used for the reason people normally use passwords, i.e. to protect their personal data. The password protocol had no actual value for authenticating the network but only a few expressed that failing correctly: *you're just typing in something, it could be another access point that accepts the exact same password ... It doesn't mean that it's Bertorelli's.* Some (mis)interpreted the password as providing authentication of customers to the system, not the other way around. As one put it: *so you've got to authenticate yourself to the ... machine.* About half the participants wondered how often the password changed: it seemed to them that the arrangement would be more secure in some sense if the password changed frequently. Some preferred the display to the poster for that reason: *The beauty of the one on the screen is that they could change the password every 24/48 hours.* One theory was that this would make life harder for an individual attacker, who would, they reasoned, have to change the attack to accommodate new passwords. A less common theory involved an actual sense in which the password does authenticate the user's circumstances: only someone with physical access to the café could know the password, and 'therefore' the number of people who could mount an attack was relatively constrained (compared to all attackers across the internet).

Security vs authenticity. The participants were asked specifically about authenticity in the questions discussed in this paper, and yet a wider conception of security often emerged strongly when they answered them. To our participants, perhaps unsurprisingly, security primarily concerns protection of their data rather than authentication *per se*. As one put it: *I think security is ... degrees of security to my vulnerability to some kind of crime really ... access to my data ... [my] address book at one end to my debit card details at the other.* Several specific forms of attack were mentioned but the most common was eavesdropping. Nineteen participants either referred to some form of eavesdropping or the need for encryption to

protect them against it. Recall that all of the networks were unencrypted, or “insecure” in Windows terminology. (Although a complete implementation of Interlock would encrypt all traffic, we did not advertise this fact to our participants.) Even some of those who professed little technical knowledge could see the difference between security and authentication in the case of wireless connections: *The fact that data I'm inputting into my computer that's going to their Wi-Fi connection might potentially be being accessed by someone. I don't know enough about the technology to know more. It feels less secure than if it was a wired Ethernet connection. I might be using their actual system, but it doesn't mean it's secure.* However, we cannot be sure that authenticity and eavesdropping were clearly separated in all participants' minds. Indeed, the confusion over the password protocol, and the desire by many for the password to be changed, suggests that some at least may have thought the password provided a degree of *security* and that may have featured in what was supposed to have been their assessment of *authenticity*.

Discussion

To summarise, we have identified physical and virtual linkage as distinct factors in methods for testing the authenticity of Wi-Fi networks, and established that both types of linkage feature in users' rationales when comparing the methods. The results of our experiment tend to confirm our hypotheses that confidence in the authenticity of a network increases with the strength of both physical and virtual linkage. We verified our hypotheses for some comparisons between methods but not others. We have put forward explanations for the cases where ratings were not significantly different. Unlike the leaflet and poster, the poster and display are not in fact so different in their physical linkage to the café, according to the aspects of physical linkage that users consider. And, in the case of virtual linkage, users find it hard to compare two unfamiliar and relatively complex protocols, although they find them both more convincing than the simple Password protocol.

Participants raised two concerns as they rated the methods. The first, which we have not mentioned so far, is usability. One aspect is the time taken to run through the Synchronisation and Interlock protocols. The other aspect is the effort involved. While several users found the Synchronisation and Interlock protocols to be engaging, and even fun, making comparisons between near and far screens (the user's laptop and the café's display) was generally felt to be difficult. We address this point below.

We have already mentioned the participants' other main concern, which is about security as opposed to authenticity: they still feel vulnerable using public Wi-Fi, even if they are sure it is authentic. Wi-Fi is familiar to most people. Many have Wi-Fi networks at home and are educated about the danger of eavesdropping and the need to enter a secret key to protect their networks. Our experimental networks

were open (unencrypted), as is often the case with free public Wi-Fi networks. Some participants mentioned the warning that Windows gives when connecting to such a network. Some of our participants were rightly concerned that, even if they knew which network was authentic, they were still vulnerable to attacks. Others observed that use of public Wi-Fi networks always involves a trade-off: convenience has to be balanced against what type of activity the network is used for. For example, the risks when browsing an online newspaper were felt to be significantly lower than when banking online. And some felt that it would be worthwhile trusting a public Wi-Fi network if the benefit was sufficiently great. For example, some would not normally use open public Wi-Fi to read email, but they said they might if they were abroad and unable to access email any other way.

Trust is an important and arguably necessary aspect of human life, but it makes us vulnerable. Wi-Fi phishing attacks are a reality. We expect widespread deployment of new ubiquitous services to increase the risks substantially. Do our protocols protect Wi-Fi users satisfactorily, and are they extensible to ubiquitous services more generally?

The only protocol in the experiment that actually protects users against a man-in-the-middle attack is the Interlock protocol. Fortunately, our experiment has shown us that users also find this protocol convincing. However, our experience with a naïve implementation of it has taught us three lessons that have led us to design a better version. First, Interlock is somewhat complicated, and asking the user to send the ciphertext in two stages is confusing and hard to explain to the non-technical. Secondly, some participants were uncomfortable about the use of a publicly visible display; in particular, privacy is a concern since users can be easily observed as they interact with it. Several suggested that they would prefer an interaction between their laptop and some kind of token given to them personally. Thirdly, there are the usability issues that also derive from using a public display. The display may not be clearly visible to all and, even if it is, comparing text or images between a remote display and a laptop can be challenging. A few spotted another problem stemming from the public display that we had ignored to simplify the experiment: if, say, three people try to authenticate the network at the same time, then three distinct avatars will have to communicate with three users from the same public display without confusing them.

In ongoing work, we are addressing these deficiencies by simplifying the user interaction and by substituting the user's mobile phone for the public display. By sending the first half of the ciphertext once the user has chosen a phrase but not exposing this to the user, we can represent to the user just the sending of the second half as the (single) act of transmitting “the message”, and reduce the number of user steps by one without changing the underlying protocol. By using a mobile phone instead of a public display, we address the other two weaknesses of our implementation.

The user communicates with her own avatar on her own phone, which she can place next to the laptop screen to make visual comparison easier. No public display is involved. To create physical linkage, the user picks up a token from a point that is strongly physically linked to the café, in all four senses identified above: attachment, legitimacy, visibility and accessibility. This could for example be a card handed over by a member of staff. The simplest form of this token is a 2D barcode which the user reads with her mobile phone, thus linking it securely to the venue and obtaining a page displaying the avatar for carrying out the Interlock protocol. The telecommunications link between the mobile phone and the Interlock service is sufficiently secure for this application.

Finally, why persist in an investigation of the Interlock protocol, and not, say, the SAS/MANA protocols [16, 6], using what we have learned about physical linkage? As mentioned in the related work section, they involve the comparison of system-chosen data. What distinguishes Interlock from other protocols is that the user can choose – and indeed can compose – the data to be compared. As we have also mentioned, an important factor for some participants was whether the data involved in the interaction were personal. In ongoing work, we are testing the hypothesis that the user’s ability to compose the text to be compared increases the user’s confidence in the access point’s authenticity, without sacrificing de facto security.

CONCLUSION

We have proposed two new concepts for authenticating ubiquitous services: physical and virtual linkage. Our experiment tended to confirm that users rate Wi-Fi networks as more authentic the stronger they are physically or virtually linked to the establishment. It also demonstrated the variety of ways in which users think about physical and virtual linkage when assessing the authenticity of networks. We evaluated two authentication protocols, one of which, Interlock, provides protection against man-in-the-middle attacks so long as it is tied to an artefact that is sufficiently physically linked to the establishment. Based on our experience with a naïve implementation of Interlock, we have devised an improved version using physically linked tokens and mobile phones rather than public displays. In future work we will evaluate this new method.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers, our shepherd Urs Hengartner, Dolce Vita café, and Danaë Stanton Fraser and Tim Jay for contributions to the initial discussions. This research is funded by the UK Engineering and Physical Sciences Research Council, grant EP/C547683/1 (Cityware: urban design & pervasive systems).

REFERENCES

- Balfanz, D., Smetters, D.K., Stewart, P., and Chi Wong, H. Talking to strangers: Authentication in ad-hoc wireless networks. In Proc. Network and Distributed Systems Security (NDSS 2002), (2002).
- BBC News: Parking ticket leads to a virus. <http://news.bbc.co.uk/2/hi/technology/7872299.stm>.
- Bugzilla@Mozilla – MITM in-the-wild. https://bugzilla.mozilla.org/show_bug.cgi?id=460374.
- Dhamija, R., Tygar, J. D., and Hearst, M. Why phishing works. In Proc. CHI 2006. ACM Press (2006), 581-590.
- Fogg, B. J. Persuasive Technology: Using computers to change what we think and do. San Francisco: Morgan Kaufman. (2002).
- Gehrmann, C., Mitchell, J. and Nyberg, K. Manual authentication for wireless devices. RSA Cryptobytes 7(1), (2004), 29-37.
- Kindberg, T., and Fox, A. System software for ubiquitous computing. In IEEE Pervasive Computing 1(1), (2002), 70-81.
- Kindberg, T., O'Neill, E., Bevan, C., Kostakos, V., Stanton Fraser, D., and Jay, T. Measuring trust in wi-fi hotspots. In Proc. CHI 2008. ACM Press (2008), 173-182.
- Kindberg, T., Zhang, K., and Im, S. Evidently secure device associations. HP Labs tech report HPL-2005-40.
- McCune, J.M., Perrig, A., and Reiter, M.K. Seeing-is-believing: using camera phones for human-verifiable authentication. In IEEE Security and Privacy, (2005), 110-124.
- Riegelsberger, J. and Sasse, M.A. Trust builders and trustbusters: the role of trust cues in interfaces to e-commerce applications. In Proc. E-Commerce, E-Society, and E-Government 2001. Kluwer, London (2001), 17–30.
- Rivest, R., and Shamir, A. How to expose an eavesdropper. Communications of the ACM, 27(4), (1984).
- Roth, V., Polak, W., Rieffel, E., and Turner, T. Simple and effective defenses against evil twin access points. In Proc. ACM Conference on Wireless Network Security (WiSec), (2008).
- Saxena, N., Ekberg, J-E., Kostianen, K., and Asokan, N. Secure device pairing based on a visual channel (extended abstract). In Proc. IEEE Symposium on Security and Privacy, (2006).
- Uzun, E., Karvonen, K., and Asokan, N. Usability analysis of secure pairing methods. In Financial Cryptography and Data Security, (2008), 307-324.
- Vaudenay, S. Secure communications over insecure channels based on short authenticated strings. In Proc. Advances in Cryptology – CRYPTO 2005, (2005), 309-326.