To Have and Have Not: Variations on Secret Sharing to Model User Presence

Quentin Stafford-Fraser
University of Cambridge
Computer Lab
quentin@pobox.com

Graeme Jenkinson
University of Cambridge
Computer Lab
graeme.jenkinson@cl.cam.ac.uk

Frank Stajano
University of Cambridge
Computer Lab
frank.stajano@cl.cam.ac.uk

Max Spencer
University of Cambridge
Computer Lab
max.spencer@cl.cam.ac.uk

Chris Warrington
University of Cambridge
Computer Lab
chris.warrington@cl.cam.ac.uk

Jeunese Payne
University of Cambridge
Computer Lab
jeunese.payne@cl.cam.ac.uk

Abstract

We address the problem of locking and unlocking a
device, such as a laptop, a phone or a security token,
based on the absence or presence of the user. We detect
user presence by sensing the proximity of a subset of their
possessions, making the process automatic and effortless.
As in previous work, a master key unlocks the device and a
secret-sharing scheme allows us to reconstruct this master
key in the presence of k-out-of-n items. We extend this
basic scheme in various directions, e.g. by allowing items
to issue a dynamically variable number of shares based on
how confident they are that the user is present. The
position we argue in this paper is that a multi-dimensional
approach to authentication that fuses several contextual
inputs, similar to that already adopted by major web sites,
can also bring advantages at the local scale.

Introduction

At major web sites, user authentication is already evolving
from a boolean evaluation of "password correct /
incorrect" to a multi-dimensional approach based on
machine learning, in which several contextual inputs are
fused to produce a continuous-valued confidence level [2].
The position we argue in this paper is that we can
fruitfully apply variations of this strategy at local scale
too, to decide when to lock and unlock a laptop, a
smartphone, or even—somewhat recursively—an

authentication token. Our approach relieves the user from having to memorise secrets and offers the benefit of continuous authentication.

Previous work by Desmedt *et al* [3], Peeters *et al* [8], and ourselves [11, 12] has discussed splitting a master unlocking secret into shares held by devices owned by a user. We present several new ideas that enhance and extend the basic model: dynamically weighted shares, negative shares and hierarchies of shares. These variations allow us to assign a more accurate rating to the trustworthiness of the contextual hint supplied by each device, thereby enhancing security while retaining better usability than traditional unlocking methods such as passwords and PINs.

Background

Establishing that a specific, conscious and approving user is present and is requesting authorisation for a particular action has traditionally involved testing for *something you know*, *something you have* or *something you are*. Greater security is often achieved by combining tests from more than one of these categories. However, each test incurs a cost, be it one of finance, inconvenience, privacy loss, or simple taxing of the memory. If the well-known and increasing problems with passwords make it less desirable to depend on *something you know*, and if the privacy, accuracy and convenience concerns about biometrics make it difficult to establish *something you are*, then it is worth exploring whether *something you have*, and the particular context in which you have it, may now usefully play a larger role in computer security, as it does in other parts of life.

The things we carry in a given context already enable us to perform many actions that other people could not. The obvious examples are key-rings, wallets, purses, and smartphones, but other possessions can also play a role: the wearing of particular items of clothing may be required to secure entrance to a formal restaurant, for example, and carrying a major credit card will help ensure a trouble-free exit afterwards. In many countries, if you wish to ride a motorbike on the public highway, you need a helmet and a particular piece of paper, as well the ignition key.

Some of these items, such as keys, may be an absolute requirement if you are to perform an action at all, while others simply increase the likelihood that society will consider the action to be legitimate. The wider context is also important: carrying your front door key will not authorise you to ride your motorbike. Wearing a motorbike helmet is highly unlikely to help you gain entrance to the formal restaurant, and might positively bar you from entrance to a bank. The possession of objects which are a positive indicator in one context may be counted negatively in another.

In the past, computer systems have been poorly equipped to detect the items that a user may be carrying with them, but the wide deployment of cameras, accelerometers and similar sensors, and of low-power radio systems, makes this approach increasingly viable.

The Pico Project:
A testbed for 'something you have'

The Pico project [12] is an interesting case study that tries to take the concept to the next stage: eliminating the need for passwords altogether, and relying solely on devices that the user carries with them. We provide a brief overview here; the project is described in detail elsewhere [12, 4, 13] and the original paper discusses related work.

The user's credentials are stored on the Pico device. The user scans the barcode of a service provider with their Pico to indicate their intention to perform an action (such as logging in to a web site, or making a payment). The Pico then uses the appropriate stored credentials to inform the service provider, over a secure channel, that the user has requested the action, which the provider can then authorise, for example by causing the the user's browser session to be logged in. In the case of an ongoing session, the Pico also provides *continuous* authentication: maintaining the session while the Pico is present, and closing it automatically as soon as the Pico is absent.

Possession of the Pico device alone, however, is insufficient to indicate *which* user currently possesses it. The Pico therefore relies on the proximity of partner devices known as 'Picosiblings', which may be items of jewellery, wristwatches, phones, key fobs, RFID tags in clothing, and so forth — items likely to be more closely attached to the user. The proximity to the Pico of a significant number of these greatly increases the likelihood that the user who possesses them is also present, and that the Pico has not been lost or stolen; you might leave your Pico on the seat of your taxi, but the situations in which you might simultaneously leave your earrings and your shoes there are much fewer. To use a Pico, then, you must also be in possession of several associated, nearby Picosiblings.

From a user's perspective, the Pico may be thought to have a 'comfort zone', in which it feels secure enough to give up its secrets. In fact, the Pico does not simply make a *decision* whether or not to authorise an action: it is actually incapable of doing so in the Picosiblings' absence, since they each possess part of the key needed to unlock those same credentials. This is an approach previously used by, for example, Desmedt *et al* [3]. A secret-sharing

algorithm, such as that of Shamir [10] or Blakley [1], can ensure that, although a user may possess n Picosiblings, at least k of them must be present for the Pico to be able to act. The Pico deletes the credentials from its working memory after using them, so k Picosiblings will be required again for recreating the master secret the next time.

Perhaps because of its simplicity, the secret-sharing concept is very powerful, and the basic k-of-n model for Picosiblings can be extended in a number of ways to model human behaviour more closely. The enhancements proposed below apply to many systems, and could be used to unlock many other kinds of devices, not just a security token like Pico: a personal computer or laptop, a smartphone, maybe even a car or motorbike. Without lack of generality we shall however retain the 'Pico' and 'Picosibling' model and terminology, since our project approaches these issues in a relatively pure and uncomplicated way.

Variable weightings

The mapping of shares to Picosiblings does not need to be a simple one-to-one relationship: different Picosiblings may have different 'weights'. If you *always* wear your watch, and nobody else ever wears it, then it is a very good indicator of your presence, and may be allocated a larger number of the shares required to unlock the Pico. A subcutaneous implant might be given even more, while something that you regularly leave at home or in the car might only have one share.

In the Pico project, there is a very specific example of when different weightings can be useful. The information stored on your Pico device can be backed up elsewhere and, in the event of the Pico being lost, stolen or

damaged, the backups can be restored to a new device. But the backups are encrypted in the same way as the data on the original Pico, so can only be decoded in the presence of a suitable number of Picosiblings. What happens if the Picosiblings, or a significant subset of them, have also been lost? To plan for this, you can grant a single Picosibling sufficient shares that it could, *singlehandedly*, unlock your Pico, and then deposit that Picosibling in a safe at your bank or with a trustworthy friend. In the event that you lose all your other worldly possessions, you can still restore an encrypted backup onto a blank new Pico and then use this special Picosibling to unlock the restored Pico and recover access to your credentials, as well as the ability to redefine a new set of Picosiblings.

Dynamic shares

Different Picosibling devices, then, can be given different levels of intrinsic importance. But the number of shares represented by a device could also depend on other factors, and might be dynamic, rather than being a fixed value that is either present or absent. Dynamically-varying numbers of shares can be used to give a more reliable assessment of the likelihood that the user is present.

Proximity is the first and most obvious example. If an estimate of the Pico-to-sibling range is available (based either on radio signal strength for a low-cost and low-energy example, or on more sophisticated distance-bounding technology for higher security), the system might use one share from a sibling which is moderately close, and two from one known to be much closer.

Ideally, this would be determined by each Picosibling, not by the Pico: when asked, the Picosibling only surrenders a number of shares corresponding to the proximity of the Pico that was asking. If Picosiblings are very dumb or passive devices, similar to RFID tags, the distinction could be provided by simple electronics. The sibling might be queried twice by the Pico, using two different radio frequencies, or two different power levels. Higher frequencies would normally degrade faster with distance and obstructions, and so could only return a share when the Pico was closer at hand. Passive Picosiblings which harvest power from the radio signal might only be able to transmit a number of shares dependent on the energy they had captured. These techniques are open to simple attacks from those able to boost the radio signals, of course, so should be taken as an illustration rather than an implementation suggestion, but they might still be appropriate for certain low-cost, less-critical scenarios.

More sophisticated Picosibling devices may also use the number of shares to indicate their own level of 'comfort' based on other factors:

- A Picosibling incorporating a biometric sensor might release multiple shares when confident of a good match, but fewer when it is less certain.

- A Picosibling normally worn on the body might stop giving out shares when it becomes completely stationary for a minute or so — usually an indication that the wearer has taken it off. A more sophisticated strategy might be to give out no shares when the device is stationary, one when it senses movement, and more shares, depending on confidence, when it recognizes the wearer to be the legitimate owner [6].

- A device that detects a particular geographic location may give up a number of shares dependent

not only on its location, but also on the quality of the GPS signal received.

- Some Picosiblings may be happier when in the presence of particular WiFi networks, indicating that they are in a familiar location.

In addition, some Picosiblings may need to be 'comforted' from time to time:

- A device incorporating a fingerprint reader may be willing to assert that the user is there, but only for a limited period since the last fingerprint scan.

- A device may need to contact a server periodically, to check that its authorisation to expose its shares has not been revoked. We use this strategy in our project to be able to remotely disable a lost or stolen Pico [12].

- Such time-dependent checks might also be expressed as a confidence level — a number of shares — which decays over a period: the proximity in time of a Picosibling's last check being comparable to the use of its proximity in space discussed above.

Some of these ideas are open to individual attacks, but the nature of these combinatorial systems is such that an attacker would typically have to employ multiple simultaneous approaches to be successful.[1]

Presence arithmetic
We can extend these ideas of variable and dynamically-changing numbers of shares by building more complicated expressions to detect user presence, based on real-world habits.

There are situations, for example, where the whole environment may be worth more than the sum of its parts. My car, and my wristwatch, may each be worth one share on their own. But if my watch is *in* my car, that's arguably a *much* greater indication that I am there too, and might warrant more than simple addition — this combination may be worth three or four shares rather than two. One way to accomplish this would be to split some shares in two and give one half of each share to each Picosibling: only when both are present can the shares be reconstructed and used. An alternative system might have *the car* also detect the proximity of the watch, and release more shares if the former can see the latter. Picos may not be the only things that need a 'comfort zone' - that concept could be extended to the Picosiblings.[2]

The opposite is also possible: danger zones. These are higher-risk situations in which a Pico should feel *less* comfortable than normal. Your suitcase could have a *negative share*, meaning that you need a larger number of Picosiblings when travelling in its presence than you might at home. A criminal baggage-handler finding a suitcase with Picosiblings in it would also then have a harder time using them. A restriction that depends on the presence of

---

[1]We have also tended to use optimistic language here, stating that a Picosibling may be willing to release more shares because it feels comfortable; it is just as valid to describe it as unwilling to release its normal complement of shares in a context when it feels threatened, thus providing greater security than a naïve k-of-n system.

[2]As a thought experiment, we might assume that Picos and Picosiblings are fundamentally the same, and extend the hierarchy both up and down. Some Picosiblings might need comfort zones, when k other 'SubPicosiblings' are nearby, and, going upwards, we might imagine a life-support machine incorporating a SuperPico so that it can only be switched off when the Picos of k other family members are in the vicinity.

a negative influence will often be weak, because a knowledgeable attacker may find it easy simply to remove that influence, for example by tipping the contents out of the suitcase. But it does reduce the likelihood that somebody who gets hold of your Pico by picking your pocket in the check-in queue will then be able to use it simply by standing close to you. Negative shares can also be dynamic: your car might start to emit them if it knew it was in a less salubrious part of your town.

Implementing a negative-share system (a challenge to the cryptographers in the audience) requires that the code and execution environment on the Pico be trusted, to ensure that all negative shares are taken into account and not simply ignored, but this is already necessary for several other aspects of the system.

In general, there is an assumption that communications between the Pico and siblings are suitably protected; Peeters's dissertation [7] discusses relevant cryptographic protocols. Once such protocols are extended to handle negative shares, attackers should not be able to distinguish negative from positive shares, otherwise they could attempt to silence or hide the negative ones. Additionally, Picosibling hardware must incorporate enough tamper-resistance to resist cloning.

Classes of shares

Modelling human habits even more realistically may require greater subtlety than can be achieved simply by adding up plusses and minuses. Consider, for example, the case of a fashion-conscious celebrity with hundreds of pairs of shoes. A shoe might be an excellent place to embed a Picosibling, but few people will be wearing more than two of them at any one time. The number of shares accessible to a burglar entering the celebrity's wardrobe might therefore be hundreds of times the number needed to unlock their Pico.

Various solutions can be suggested for this: shoes which will release their shares only when a foot of the right size is wearing them, for example, or when a human of the correct weight is standing in them.

But a more general solution would be for Picosiblings to be divided into 'classes'. For a particular user, the classes might include footwear, jewellery, millinery, gadgets, and implants. A surfeit of shares in one class would have no effect if shares from the other classes were missing: all the shoes in the world are of no use if you don't have a suitable tiara to go with them, and tiaras are kept in the safe, not the wardrobe.

We can model this quite nicely with the secret-sharing ideas as well, by adopting a two-stage process. First, the master key needed to unlock the Pico's credentials is split into m shares, one for each class, such that the shares from j-of-m classes are needed to recreate it. (We may not need to insist on an item from *every* class, just a reasonable subset.) Then, each of those class-specific shares is split up in the same way such that k_C of n_C of these sub-shares are needed to reconstruct the share for class C.

We can then require, for example, any combination of three of the following sets: {two shoes}, {one hat}, {one phone}, {three items of jewellery}.

Conclusion

We have shown how 'something you have' can be extended to 'several things you have' and then, through secret-sharing, to more sophisticated models of user presence by considering variable numbers of shares, dynamically-adjusted numbers of shares, negative shares, and classes of shares.

It remains to be seen which combinations of real-world situations should map onto Pico comfort zones or danger zones, and then whether the components described above are useful ways of modelling these. Such a system may also become very user-specific, which is not a problem if the user is willing and able to understand its operation. But can we generalise such models to provide a useful default template configuration for certain classes of user, and if so, how would those classes be defined? By gender or generation, by culture or nationality, by level of expertise, or perhaps by occupation?

Location privacy will need to be protected. The specification for the Picosiblings protocol [12, section 4.1] requires that eavesdroppers be unable to infer long-term pseudonyms, but it won't be trivial for actual implementations to provide this feature.

Another challenge, as we make a more complex and secure system, is whether we can also make it understandable. The Pico system's increased security and general ease of use may be of little consolation to somebody locked out of their house in the small hours of the morning just because they left their tiara in a taxi, especially if they don't understand *why* that was the

critical factor. We must not only represent user behaviour with such a system, but also provide a good mental model of how we do so, and communicate that effectively [9].

There are opportunities for valuable psychological and ethnographic user studies to be done in this area, but the underlying k-of-n model — being easily pictured in terms of having 'enough letters to guess the word', or 'sufficient jigsaw pieces to see the picture', or 'at least three of the starship's officers present to initiate the auto-destruct sequence' — may be easier for users to grasp than some other systems, at least if we don't go overboard with hierarchical levels of classes of shares. Usability will, however, always be crucial and this is why in our project all design ideas are subject to validation through user studies. We like inventing innovative and potentially more secure schemes, but we'll only carry them through to the next stage if the prototypes we build and test on non-geeks convince us that they actually make life easier for regular humans.

Acknowledgments

References

[1] Blakley, G. Safeguarding cryptographic keys. In *Proceedings of the 1979 AFIPS National Computer Conference*, vol. 48, AFIPS Press (1979), 313–317.

[2] Bonneau, J., Herley, C., van Oorschot, P., and Stajano, F. The past, present and future of password authentication on the web. (In submission).

[3] Desmedt, Y., Burmester, M., Safavi-Naini, R., and Wang, H. Threshold things that think (t4): Security requirements to cope with theft of handheld/handless internet devices. In *Symposium on Requirements Engineering for Information Security* (2001).

[4] Jenkinson, G., Spencer, M., Warrington, C., and Stajano, F. I bought a new security token and all I got was this lousy phish—Relay attacks on visual code authentication schemes. In *Proceedings of Security Protocols Workshop 2014*, Bruce Christianson et al., Ed., LNCS, Springer (2014). (To appear).

[5] Krause, F. Designing secure & usable picosiblings: An exploration of potential pairing mechanisms. Master's thesis, University of Cambridge, 2014.

[6] Mantyjarvi, J., Lindholm, M., Vildjiounaite, E., Makela, S., and Ailisto, H. Identifying users of portable devices from gait pattern with accelerometers. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'05)*, vol. 2 (2005).

[7] Peeters, R. *Security Architecture for Things That Think*. PhD thesis, K.U. Leuven, June 2012.

[8] Peeters, R., Kohlweiss, M., and Preneel, B. Threshold things that think: Authorisation for resharing. In *Proceedings of iNetSec 2009 — Research Problems in Network Security*, J. Camenisch and D. Kesdogan, Eds., vol. 309 of *IFIP Advances in Information and Communication Technology* (Zurich,CH, 2009), 111–124.

[9] Peeters, R., Kohlweiss, M., Preneel, B., and Sulmon, N. Threshold things that think: usable authorization for resharing. In *SOUPS*, L. F. Cranor, Ed., ACM (2009).

[10] Shamir, A. How to share a secret. *Commun. ACM 22*, 11 (Nov. 1979), 612–613.

[11] Stajano, F. The resurrecting duckling — what next? In *Proceedings of Security Protocols Workshop*, B. Christianson, B. Crispo, and M. Roe, Eds., vol. 2133 of *LNCS*, Springer (2000), 204–214.

[12] Stajano, F. Pico: no more passwords! In *Proceedings of the 19th International Conference on Security Protocols*, SP'11, Springer-Verlag (2011), 49–81.

[13] Stajano, F., Jenkinson, G., Payne, J., Spencer, M., Stafford-Fraser, Q., and Warrington, C. Bootstrapping adoption of the Pico password replacement system. In *Proceedings of Security Protocols Workshop 2014*, Bruce Christianson et al., Ed., LNCS, Springer (2014). (To appear).

[14] Toader, C. User authentication for pico: When to unlock a security token. Master's thesis, University of Cambridge, 2014.