Handout 6 (Zero-Knowledge Proofs)

Zero-knowledge proofs (short ZKP) solve a paradoxical puzzle: How to convince somebody else that one knows a secret without, revealing what the secret is? This sounds like a problem the Mad Hatter from Alice in Wonderland would occupy himself with, but actually there some serious and not so serious applications of it. For example, if you solve crosswords with your friend, say Bob, you might want to convince him that you found a solution for one question, but of course you do not want to reveal the solution, as this might give Bob an advantage somewhere else in the crossword. So how to convince Bob that you know the answer (secret)? One way would be to come up with the following protocol: suppose the answer is *folio*. Then look up the definition of that word in a dictionary. Say you find

> "an *individual* leaf of paper or parchment, either loose as one of a series or forming part of a bound volume, which is numbered on the recto or front side only."

Take the first non-article word in this definition, in this case *individual*, and look up the definition of this word, say

> "a single *human* being as distinct from a group"

In this definition take the second non-article word, that is *human*, and again look up the definition of this word. This will yield

> "relating to *or* characteristic of humankind"

You could now go on to look up the definition of the third non-article in this definition and so on. But let us assume you agreed with Bob to stop after three iterations with the third non-article word in the last definition, that is **or**. Now, instead of sending to Bob the solution *folio*, you send to him *or*. How can Bob verify that you know the solution? Well, once he solved it himself, he can use the dictionary and follow the same "trail" as you did. It the final word agrees with what you send him, he must infer you know the solution earlier than him. This protocol works like a one-way hash function because *or* does not give any hint as to what was the first word was. I leave you to think why this protocol avoids article words.

After Bob found his solution and verified that according to the protocol it "maps" also to *or*, can he be entirely sure no cheating is going on. Not with 100% certainty. It could have been entirely possible that he was given *or* as the word but it derived from an entirely different word—this seems very unlikely, but is at least theoretical a possibility. Protocols based on zero-knowledge proofs will produce a similar result—they give an answer that might be erroneous in a very small number of cases. The point is to itterate them long enough so that the theoretical possibility of cheating is negligibly small.

By the way, the authors of the paper "Dismantling Megamos Crypto: Wirelessly Lockpicking a Vehicle Immobilizer" who were barred from publishing

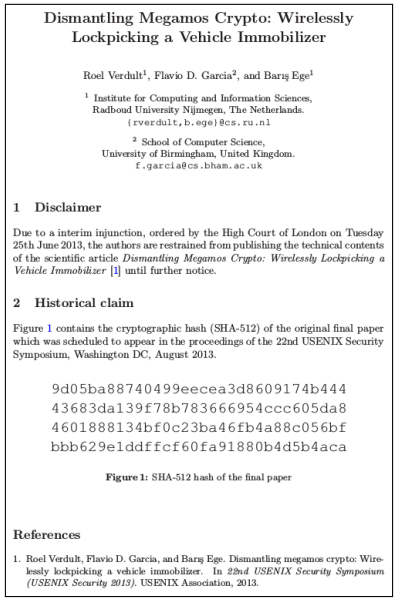

Dismantling Megamos Crypto: Wirelessly Lockpicking a Vehicle Immobilizer

Roel Verdult[1], Flavio D. Garcia[2], and Barış Ege[1]

[1] Institute for Computing and Information Sciences, Radboud University Nijmegen, The Netherlands.
{rverdult,b.ege}@cs.ru.nl

[2] School of Computer Science, University of Birmingham, United Kingdom.
f.garcia@cs.bham.ac.uk

1 Disclaimer

Due to a interim injunction, ordered by the High Court of London on Tuesday 25th June 2013, the authors are restrained from publishing the technical contents of the scientific article *Dismantling Megamos Crypto: Wirelessly Lockpicking a Vehicle Immobilizer* [1] until further notice.

2 Historical claim

Figure 1 contains the cryptographic hash (SHA-512) of the original final paper which was scheduled to appear in the proceedings of the 22nd USENIX Security Symposium, Washington DC, August 2013.

```
9d05ba88740499eecea3d8609174b444
43683da139f78b783666954ccc605da8
4601888134bf0c23ba46fb4a88c056bf
bbb629e1ddffcf60fa91880b4d5b4aca
```

Figure 1: SHA-512 hash of the final paper

References

1. Roel Verdult, Flavio D. Garcia, and Barış Ege. Dismantling megamos crypto: Wirelessly lockpicking a vehicle immobilizer. In *22nd USENIX Security Symposium (USENIX Security 2013)*. USENIX Association, 2013.

Figure 1: The authors of this paper used a hash in order to prove later that they have managed to break into cars.

their results used also a hash to prove they did the work and (presumably) managed to get into cars without a key. See Figure 1. This is very similar to the method about crosswords. They like to prove that they did the work, but not giving out the "solution". But this also shows what the problem with such methods are: yes, we can hide the secret temporarily, but if somebody else wants to verify it, then the secret has to be made public. Bob needs to know that *folio* is the solution before he can verify the claim that somebody else had the solution first. Similarly with the paper: we need to wait until the authors are finally allowed to publish their finding in order to verify the hash. This might happen at some point, but equally it might never happen (what for example happens if the authors loose their copy of the paper because of a disk failure?). Zero-knowledge proofs, in contrast, can be immediately be checked, even if the secret is not public yet.

ZKP: An Illustrative Example

The idea behind zero-knowledge proofs is not very obvious and will surely take some time for you to digest. Therefore lets start with an illustrative example. The example will not be perfect, but hopefully explain the gist of the ideas. The example is called Alibaba's cave:

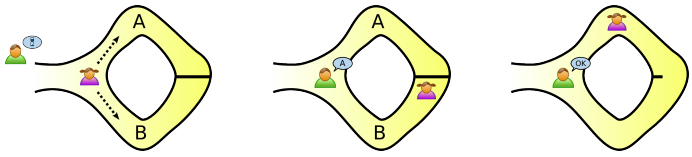

Step 1 Step 2 Step 3

Lets have a look at the picture in Step 1: The cave is a loop where at the end is a magic wand. The point of the magic wand is that Alice knows the secret word for how to open it. She wants to keep here word secret, but wants to convince Bob that she knows it.