Eating from the Tree of Ignorance
Part 2

Jan van Eijck, CWI Amsterdam and Uil-OTS Utrecht
Rineke Verbrugge, Institute of AI, University of Groningen

ESSLLI 2009, Bordeaux, July 22, 2009

Overview of Part 2

- Protocols for Preserving Anonymity and Privacy
- Public Key Cryptography: How does it Work?
- Epistemic Analysis: Effects of Telling Secrets
- Protocol for Anonymity: the Dining Cryptographers
- Epistemic Analysis
- Social Networks, Common Knowledge, Coordinated Action
- Disclosing the Truth, Not Disclosing the Truth, Lying
- Individual Ignorance vs Common Ignorance
- Fruits of Ignorance

In Praise of Ignorance

Rineke There is a saying "The innocent have nothing to fear", suggesting that only those with criminal intentions should worry about personal information getting public.

Jan I don't know where you got that from, but I think it is very dangerous. The distinction between the public and the private sphere is fundamental in Western democracies. It is also in the Universal Declaration of Human Rights, in article 12. I looked it up.

> No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honour and reputation. Everyone has the right to the protection of the law against such interference or attacks.

Protocols for Preserving Anonymity and Privacy

- Sending emails without cc:s (possibly with encryption). Or: Sending letters in closed envelopes (maybe with a seal). Meant to keep the contents of the message private.

  Third parties should remain ignorant of the message contents.

- Organizing a secret society on a need-to-know basis: meant to prevent the membership list of the society from becoming common knowledge.

  Members should remain (partly) ignorant about who are their fellow members.

- Finding out if you share a secret with someone, without disclosing the secret if you are not.

  Other party should not find out the secret if she does'nt already.

- Reviewing process of scientific papers: meant to preserve anonymity of the reviewer.

  Author should remain ignorant of the identity of the reviewer.

- Casting an anonymous vote.

  Others should not be able to detect your vote.

- Casting an anonymous receipt-free vote.

  Others should not be able to detect your vote; moreover, you should not be able to prove your vote. The vote is kept private even when the voter wishes to reveal it. This property is required in a setting with vote-buyers or coercers, where the voter wants to reveal his vote.

Public Key Cryptography: How does it Work?

- Suppose I tell you that 40285327 is the product of two primes, and challenge you to produce these primes. How would you do it? You are allowed to use a pocket calculator ...

```
Hugs.Base> 40285327 / 7
5755046.71428571

Hugs.Base> 40285327 / 509
79146.025540275

Hugs.Base> 40285327 / 5333
7553.97093568348

Hugs.Base> 40285327 / 5347
7534.19244436132
```

- Suppose instead I tell you that 7879 and 5113 are primes, and I ask you to calculate their product. Very easy:

$$\begin{array}{rl} 7879 & \\ 5113 & \times \\ \hline 23637 & \\ 78790 & \\ 787900 & \\ 39395000 & + \\ \hline 40285327 & \end{array}$$

- Multiplication of two large prime numbers is easy, but finding the prime factors of a large number is very difficult.
- No known method for finding the prime factors of a number is substantially better than trial and error.

The RSA Algorithm for Public Key Encryption [RSA78]

RSA (Rivest, Shamir, Adleman) public/private key generation:

1. Choose two large random prime numbers p and q,
2. Compute $n = pq$.
3. Compute the totient $\varphi(n)$ of n. This is the number of positive integers i with $i \leq n$ and $\gcd(i, n) = 1$ (i co-prime to n).
   $\varphi(n) = (p-1)(q-1)$.
   Example: $\varphi(15) = 8$, for $1, 2, 4, 7, 8, 11, 13, 14$ are co-prime to 15.
4. Choose an integer e with $1 < e < \varphi(n)$ and e co-prime to $\varphi(n)$.
   Release e as the public key exponent.
5. Compute d to satisfy $de = 1 + k\varphi(n)$ for some integer k.
   I.e., $de = 1 \pmod{\varphi(n)}$.
   Keep d as the private key exponent.

Encrypting and Decrypting

Alice transmits her public key (n, e) to Bob and keeps the private key d secret. Bob then wishes to send message M to Alice. First he turns m into a number smaller than n. Next he computes cipher c given by

$$c = m^e \pmod{n}$$

and transmits c to Alice.

Alice can recover m from c by using her private key d, as follows:

$$m = c^d \pmod{n}$$

From m, Alice can recover the original message M.

Why Does This Work?

$$c^d \equiv (m^e)^d \equiv m^{ed} \pmod{n}.$$

(Please take on faith that)

$$m^{ed} \equiv m \pmod{p}$$

$$m^{ed} \equiv m \pmod{q}.$$

It follows that

$$m^{ed} \equiv m \pmod{pq}.$$

and therefore

$$c^d \equiv m \pmod{n}.$$

Using Public Key Cryptography

- Public Key Cryptography is asymmetric. Analogy of a padlock: anyone can lock it, only someone with the key can unlock it.
- Never a need to send a (private) key over an insecure channel.
- Sending public keys around can do no harm.
- Analogy: sending open padlocks over the mail, and inviting recipients to use them to lock something and send it back to you.
- Use for encryption. Encrypt with the intended recipient's public key. Only the intended recipient can decrypt.
- Use for authentication. Sign a message with your private key. Anyone who has the corresponding public key can check that the signature is yours. Secure digital signatures.

Epistemic Analysis: Effects of Telling Secrets

Suppose p is a secret: Alice knows p, but Bob and Carol (often called 'Eve' for 'eavesdropper') do not. They do not even suspect that Alice knows.

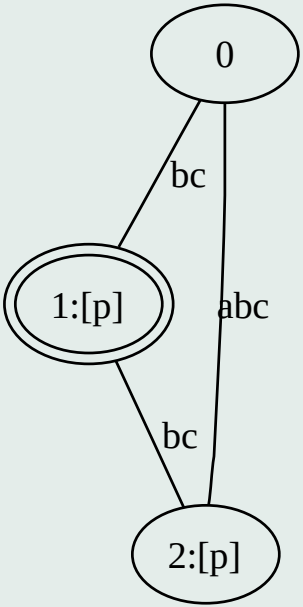

Now Alice tells Bob the secret:

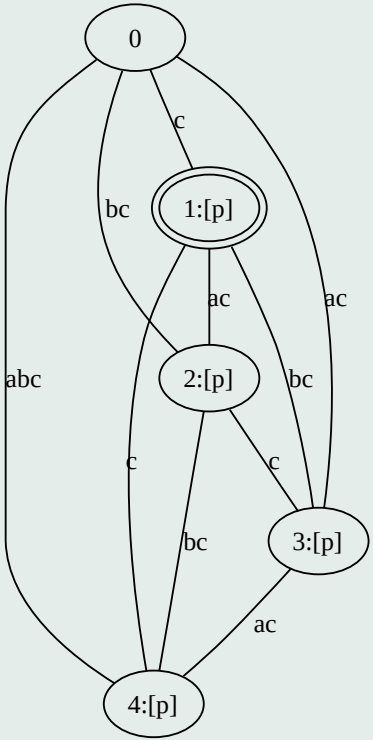

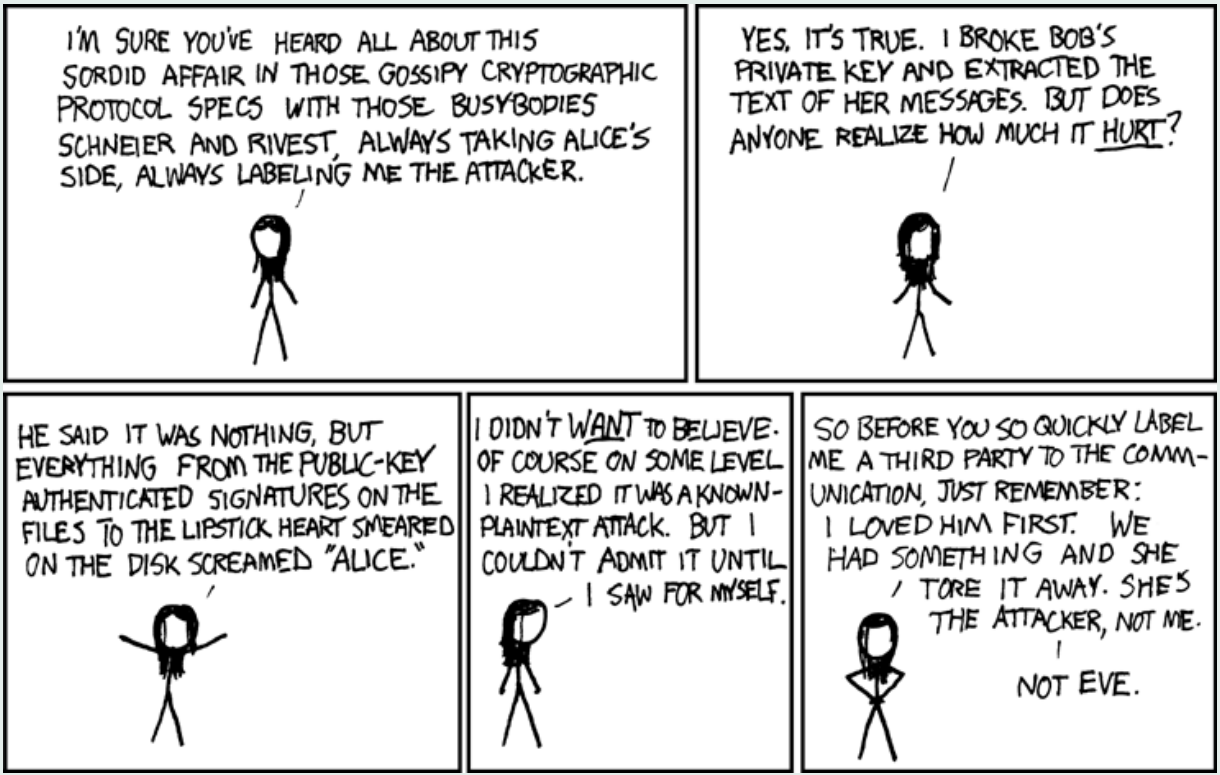

drawing by xkcd

Public Announcement

- Consider an epistemic model, with a set of worlds W.
- Then the effect of making a public announcement φ in that model is that all non-φ worlds disappear from the model.
- Public announcements can be used to create common knowledge. More on that tomorrow.

Example: the Dining Cryptographers

Chaum [Cha88]: three cryptographers are eating out. At the end of the dinner, they are informed that the bill has been paid, either by one of them, or by NSA (the National Security Agency).

They want to find out whether NSA paid or not.

They also want to respect each others rights to privacy: in case one of them has paid the bill, her identity should not be revealed to the two others.

Restrictions: **don't use a trusted outsider or ballot box**. **Assume that all conversations can be overheard.** In other words: the only communication that can be used is **public announcement**.

How can they do it?

Protocol

Each cryptographer tosses a coin with his righthand neighbour, with the result of the toss remaining hidden from the third person.

Each cryptographer then has a choice between two public announcements: that the coins that she has observed agree or that they disagree.

- If she has not paid the bill she will say that they agree if the coins are the same and that they disagree otherwise;
- if she has paid the bill she will say the opposite: she will say that they agree if in fact they are different and she will say that they disagree if in fact they are the same.

Why does this solve the problem?

Possible Situations

- All coins the same:
  - no liars: no disagreements
  - one liar: one disagreement
- One coin different:
  - no liars: two disagreements
  - one liar: one or three disagreements

Therefore...

- If there is an even number of disagreements, the NSA has paid
- If there is an odd number of disagreements, one of the cryptographers has paid.

Epistemic Analysis

- Start with a situation where noone knows anything, and this ignorance is common knowledge.
- Update with public announcement of 'at most one diner paid', so that this becomes common knowledge.
- Update with the information that every participant knows whether she has paid or not.
- Update with the results of the coin tosses.
- Update with appropriate group announcements of the results of the coin tosses.
- Update with appropriate public announcements about coin (dis)agreement.

Final Situation of Dining Cryptographer Scenario

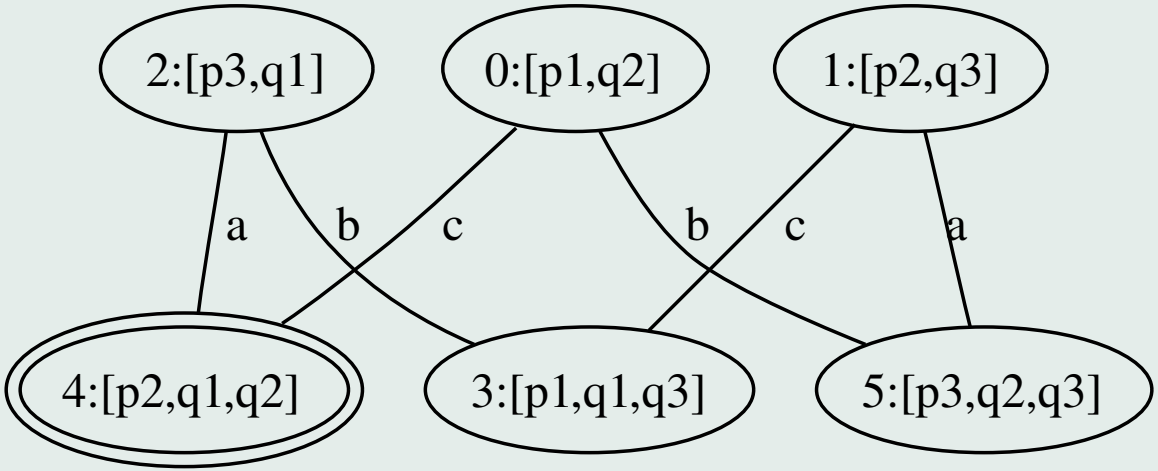

- Diner 2 paid, coins $1, 2$ show heads.
- $\neg K_a p_2, \neg K_c p_2$
- $C_{\{a,b,c\}}(\neg K_a p_2 \wedge \neg K_c p_2)$.
- $K_b(q_1 \wedge q_2 \wedge \neg q_3)$

More than three dining cryptographers

Question: find out how many out of N dining cryptographers have made contributions to the bill, without revealing their identities.

- Let someone start by whispering a number M larger than N into the ear of her lefthand neighbour.
- The neighbour then whispers a number to his lefthand neighbour, and so on.
- The ones who did not pay pass on the same number they heard.
- The ones who paid increase the number by one.

More than three dining cryptographers: how do they know?

- After one round, the initiator of the protocol hears the number K, and she knows that $K - M$ people contributed to the bill.
- In the second round, those who have contributed to the bill again increment the number they hear.
- In the course of the second round everyone finds out, by comparing the number they heard the first time with the number they heard the second time.

This procedure is due to Valentin Goranko.

The procedure is not quite as good as Chaum's original proposal for three cryptographers. Assuming that every conversation can be overheard this is an insecure procedure.

(Chaum also has a secure version for $N \geq 3$.)

Social Networks, Common Knowledge, Coordinated Action

Example from [Chw01]. Two social networks.

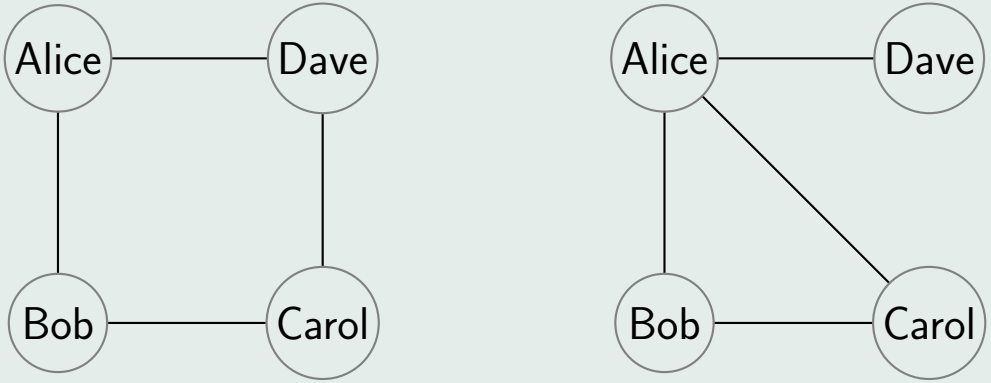

- Everyone thinks: "If I know for sure that at least two other people are going to take action, I will join in."
- Everyone communicates this intention to their neighbours.

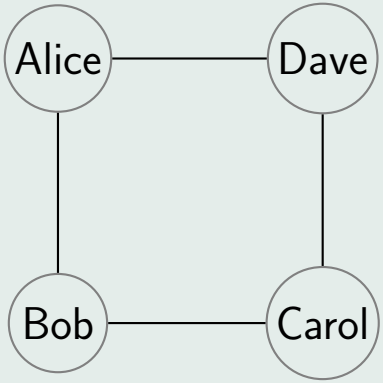

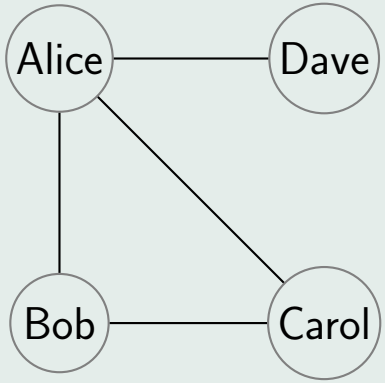

- It makes a difference whether I am in touch with the neighbours of my neighbours or not.
- Why? Because knowing that my neighbour will join if at least two other people join is not enough for me to be sure that he will join in. He can be sure about me. But how about his other neighbours?

Disclosing the Truth, Not Disclosing the Truth, Lying

Dr A. "By the way, were you one of the reviewers of my paper?"

Dr B. "I am sorry, but I think we should not discuss this matter. Maybe I was, maybe I was not. I am not going to tell you, for I believe in anonymity of reviewing."

Dr C, who actually was not a reviewer "Well, if I had been I would of course not been allowed to tell you. But in fact I was not."

Dr D, who actually was a reviewer "No, I did not review your paper."

Dr D to dr B "In fact I did review that crap, but of course I couldn't tell poor Dr A."

Now define the notion of an honest answer ...

Showing your ignorance may lead to knowledge on both sides

- There are situations where neither A nor B knows whether p. Then they meet; B asks whether p, and A truthfully answers, 'Yes, I know'.
- Suppose A has the additional piece of information that if p is not the case, then B knows that not p.
- The chair of the programme committee, Professor A, has been told by his secretary that all authors of rejected papers have been notified.
- When Doctor B meets Professor A at ESSLLI, then B's question 'Has my paper been accepted?' reveals to A that the answer must be 'Yes', for A reasons that otherwise B would have known.

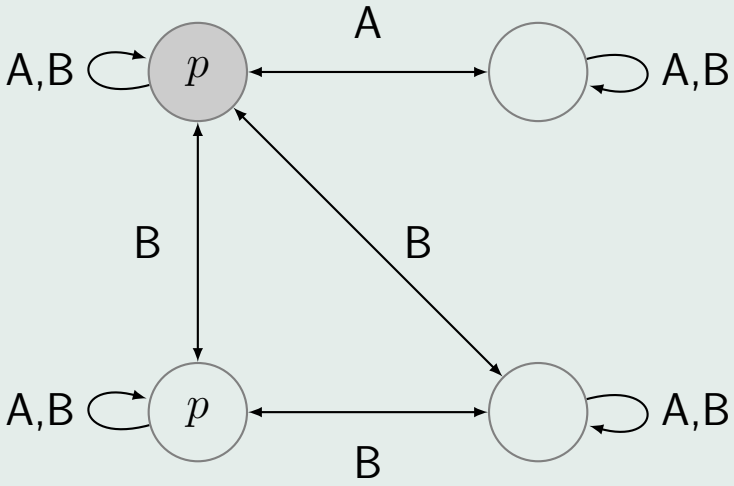
A
A,B
p
A,B
B
B
A,B
p
A,B
B

Individual Ignorance vs Common Ignorance

Individual Knowledge about φ:

$$K_a\varphi \vee K_a\neg\varphi.$$

Individual Ignorance:

$$\neg K_a\varphi \wedge \neg K_a\neg\varphi.$$

Common knowledge:

$$C\varphi \vee C\neg\varphi.$$

Common ignorance:

$$\neg C\varphi \wedge \neg C\neg\varphi.$$

Commonly known common ignorance:

$$C(\neg C\varphi \wedge \neg C\neg\varphi).$$

Fruits of Ignorance

- Common ignorance is compatible with individual knowledge for all individuals in the community.
- Even if everyone has individual knowledge that φ, public announcement of φ still has an epistemic effect. (Cf. Olmert's nuclear slip-up.)
- "φ-ambiguity" can only be maintained if there is common agreement that φ should remain ambiguous.
- Privacy protection is only possible through a communal effort. If φ concerns private information, then everyone should refuse to make public statements about φ.
- Challenge for Social Software Analysis: design a framework in which "public φ-hypocrisy" can be formally expressed and analyzed.

Social Software Design

One of the exercises of today:

A group of 100 prisoners, all together in the prison dining area, are told that they will be all put in isolation cells and then will be interrogated one by one in a room containing a light with an on/off switch. The prisoners may communicate with one another by toggling the light-switch (and in no other way). The light is initially switched off. There is no fixed order of interrogation. Every day one prisoner will get interrogated. At any stage every prisoner will be interrogated again sometime.

When interrogated, a prisoner can either do nothing, or toggle the light-switch, or announce that all prisoners have been interrogated. If that announcement is true, the prisoners will (all) be set free, but if it is false, they will all be executed. Can the prisoners agree on a protocol that will set them free?

References

[Cha88] D. Chaum. The dining cryptographers problem: unconditional sender and receiver untraceability. Journal of Cryptology, 1:65–75, 1988.

[Chw01] Michael Suk-Young Chwe. Rational Ritual. Princeton University Press, Princeton and Oxford, 2001.

[RSA78] R. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. Communications of the ACM, 21(2):120–126, February 1978.