DIMACS Working Group on Measuring Anonymity
Notes from Session 2: Recent Perspectives

Scribe: Matthew Wright

In this session, we had two 15-minute talks based on submitted abstracts, one 20-minute talk based explaining important prior work, and about 40 minutes of “panel” discussion with the three speakers as panelists. The focus of the session was on recent work that fit within the anonymity measurement paradigms already being explored in recent years, as well as capturing a range of perspectives on the working group’s broader goals: Tor-specific analysis, the design of metrics, and dealing with the complexity in these systems.

Talk #1: **Aaron Johnson**, Chris Wacek, Rob Jansen, Micah Sherr, and Paul Syverson. *Traffic Correlation on Tor: Models, Metrics, and Results*. (15 min.)

Adversary Models:

- Node Adversary: with c of n nodes, fraction f of bandwidth, g of guard BW and e exit BW fraction.
  - Improved: g_t prob. of picking a guard within time t
  - Next step: How likely is correlation over time?
- Link adversary: ISP, ASes, IXPs
  - Why worry about all ASes/IXPs simultaneously?
  - What if my adversary controls many ASes/IXPs?
  - Improved: Fix the adversary’s location, multiple ASes/IXPs

Correlation metrics:

- Probability of compromise (fractions of)

What are some good metrics for this?

- Probability of compromise (fractions of)
- Days to compromise (human timescales)
- Getting the best or worst case

Talk #2: **Rajiv Bagai**. *A Graphical Framework for Representing Anonymity Metrics.* (15 min.)

A metric should capture the profile of a distribution:

- Existing anonymity metrics employ local observations of the profiles of the distribution. (Shannon entropy, Euclidian distance)
- BUT Profiles can intersect (any number of times)
- Plus, local information doesn’t capture everything.

Area-based metric: captures whole curve information.

- And is a weighted harmonic mean: a truer measure in some cases (ratios: e.g. resistance, speed)

- Is anonymity one of those cases?
- Well, probabilities are ratios...

Talk #3: **George Danezis**: *Bayesian approaches for assessing anonymity* (20 min.) – an invited talk.

See Carmela Troncoso's PhD thesis for the best reference: http://homes.esat.kuleuven.be/~ctroncos/index_archivos/Page842.htm

Probability view of anonymity

- Pr[A -> R_i | Observations/Knowledge and Constraints]
- Metrics pre-'09: summarizing this probability distribution in some way
  - Entropy, normalized entropy, min entropy, probability of error
  - This is not the interesting problem.
- Interesting problem: how do we compute this?
  - Example: constraint of length = 2 gives us a lot more info.
  - So even toy examples blow up fast with some constraints

Using a Bayesian approach

- The likelihood and the prior, normalizing factor in the denom.
- Capture all known traffic analysis as the "hidden state" (mixing)
- If only we could do this over all possible configurations
  - [in some sense, does not handle long-term intersection attacks]
  - There are a lot of configurations
    - Sample from the space (using AI techniques)

Long-term attacks

- Profile users over time (find Alice's friends)
- While also guessing at who sends what right now
  - Combine for a power boost
- Now, we can use the Bayesian approach to find the probabilities.
- And leverage this (also in a Bayesian way) to find the profiles.

Crowds

- (pass the hot potato message passing)
- Corrupt insider attack: uses Markov path length extension to avoid it
- Can show that it's optimal (the Wisdom of Crowds)

The likelihood function is at the heart of metrics.

- We need to know this to measure anonymity.
  - Intrinsic (to the protocol) vs. incidental hidden state
  - Random permutation vs. network delay
- Problems
  - How do we (safely) approximate incidental hidden state?
  - Computational/empirical vs. closed-form theorems
  - Long-term inference

Panel Discussion (40 min.)

Note: Discussion participants are labeled 'A' to 'Z' for each question.

- [to Ragiv] Any intuitive sense of what your harmonic mean metric means?
    - A: No. It seems that the problem is we don't know what anonymity is. That's why we don't know how to measure it.
    - B: There is no best summary of the distribution.
    - A: a distribution is like a giant tuple. A summarizing formula would be useful. But unless we know what we want, we can't figure that out.
    - C: I agree w/ George -- given the distribution, picking a useful quantity is easy and uninteresting and mostly done. Getting the distribution is hard.
    - D: a lot of our metrics don't compose nicely. How does the anonymity compose with the privacy of e-cash, e-voting, etc.? The security designers cannot rely on the anonymity as a black box, making their job much harder.
    - C: Composability is nice, but not having it is not the end of the world.
    - E: We don't need composability if we have the distribution.
        - Ed. note: I believe the point is that if we have the distribution, then that is enough information to compose correctly. Not sure if that's true…

- How important is communicating to the user?
    - A: This is what Aaron's paper does to some extent (time to compromise).
    - B: How do you answer the user's question about her privacy? [this question is perhaps out of scope]
    - C: Don't forget the users' needs in the design and metrics, even if you're not communicating to them.
    - D: Communicating is hard; you have to explain the threat model to even get started. We see confusion despite really clear communication on the Tor website.
    - E: And there's not just one threat model. There is a diversity of users with a diversity of needs.

- We also need to ask: what are the alternatives to Tor-like designs? We should measure them, too.

- [to George] Noting the optimality of Crowds: Is there a relationship between the Bayesian approach and the Shannon entropy?
    - A: It would seem so. The memory-less property is the key one. We didn't aim for this in our theorem or proof.

- The metric should be based on the actual attack success.
    - A: While that would be good, confidentiality loss is hard to see. There may be zero behavioral evidence.
    - B: Difference in statistics is based on ground truth. We don't have ground truth.

- A: Most people using anonymity (as a black box) are not much more sophisticated than my grandmother. I don’t see a big distinction between groups of users.