Consolidated Review of

A Fistful of Bitcoins: Characterizing Payments Among Men with No Names

1. Strengths

Interesting paper (both in its "tutorial" aspects and analysis), with nice methods and results to assign addresses to real entities. The underlying problem, and the network analyzed in this paper are very interesting. The story and evolution of the heuristics was captivating, and the results can help better understand the Bitcoin network and its operations.

It answers an important question (what quality of anonymity does Bitcoin provide?) It presents an interesting new result (that Bitcoin transaction flows can be to a large extent "de-anonymized"). It observes and analyzes what the authors' call as 'peeling chains', and proposes interesting approaches towards tracking thefts and suspicious huge monetary transactions.

Paper performs a good job of collecting data, and their two heuristic seems reasonable (at least presently).

2. Weaknesses

Enhances previous methods and the results are somewhat incremental

Even though I enjoyed the story behind the evolution of Heuristic 2, the ad hoc nature of this refinement creates a big question mark for the results presented in this paper, and the conclusions reached.

1. Their first heuristic is well-known in the literature, while the weaker versions of the second heuristic has also already proposed and employed in a previous paper. This makes their efforts incremental.
2. Although their heuristics and corresponding reidentification attacks may work presently, they are not robust: Criminals knowing these attacks may try to increase false positive by various means. For example, by introducing some (gambling) game where multiple users may input to the same transaction. Or, as discussed by the authors, criminals may start employing Satoshi Dice as a mixing service to hide their peeling chains. Authors should discuss the robustness of their heuristics and their approach in general.
3. In absence of ground truth data, the authors used some approximating assumption, which again raises a question at least about the robustness of the approach. The authors says "if an address and transaction met the conditions of Definition 4.3 at one point in time (where time was measured by block height), and then at a later time the address was used again, we considered this a false positive." This will be a very small barrier for the fraudulent entities to cross in the future.
4. Conclusion of the analysis in Section 5 is not clear to me. In particular, what do we learn from or conclude in Section 5.1? Also, at the end of Section 1, I felt that deanonymization (or reidentification) is something desirable by the authors, however Section 5 ends on an opposite (or at least indifferent) note.

3. Comments

I like the problem studied in this paper, the story behind the methodology, and the results presented in this paper. Studying privacy in Bitcoin is a recent hot topic. This paper starts with a "tutorial"-like introduction, which I find as interesting and informative.

Even if the paper gets accepted, IMC may not be the best venue for this work. The IMC community will benefit from exposure to a new topic and an interesting result, but the authors will not receive rigorous technical feedback on their ideas.

As a non-specialist, given how the paper is written, I had trouble separating fundamental issues from the rest. For instance, in Section 3.1, it says that Eligius split the coin among the miners immediately, and thus the authors were unable to identify Eligius's addresses. So... why don't other mining pools do the same? Is there any cost associated with immediate coin splitting? Or is it simply that mining pools do not care if their addresses are identified?

More generally, it would be great if the paper answered (or provided some insight toward answering) this question: If the main Bitcoin players ***wanted*** to improve Bitcoin anonymity, would they be able to do it and how? How well would the proposed methodology work in the context of such an "adversarial" scenario? Differently said: What are the fundamental limits of the proposed methodology and how does its accuracy depend on specific behavior of the Bitcoin community?

Another relevant question: How (and how much) would one have to change Bitcoin design to improve anonymity? I realize this is a big question (that warrants multiple papers), but any discussion would be enlightening.

As mentioned before, I think there are major question marks especially with regard to the heuristics used, and the refinement methodology. Given the potential inaccuracies that are not caught by manual inspections explained in the paper, the high level results become less trustable. After all, the false positive rate that most of the refinement is based upon is an approximation as well. Having said all of that, these kinds of inaccuracies are expected to some extent from a work of this nature. The arguments conclusions reached in Section 5.2 are rather weak. Relying on historic volumes to argue that it is not possible cashing out at scale is not a strong argument. Making such statements more rigorous by providing details can make it more tangible and (somewhat) more reliable.

The second heuristic seems to add on the well known first heuristic, so although the results are somewhat incremental, I still find the contribution as interesting.

Although there are no break-through results in the paper, it extends our understanding of anonymity in the Bitcoin ecosystem.

4. Summary from PC Discussion

This paper seems to be liked by most of us, so I believe we will have an easy decision on this one.

Strengths:

- Tackles an interesting problem.
- Shows nice results regarding the ease of de-anonymizing transactions.

Weaknesses:

- Incremental contribution, first method is well known.
- The second heuristic, which is new, is not robust.
- There is no in depth discussion about potential "fixes" to the problem
- Our weakness - most of us are not experts in Bitcoin...

5. Authors' Response

The main concerns about our paper (weaknesses 2 and 3) seemed to focus on our second heuristic, in which we cluster addresses based on the mechanism by which change is made in the currently standard Bitcoin client. As noted, the heuristic lacks some robustness in the face of changing patterns in the Bitcoin network: if everyone stopped using one-time change addresses, then our second heuristic would essentially reduce to our first; worse yet, if an adversary were aware of our analysis and wished to create false positives to thwart us, it would be relatively easy to do so. These concerns seem accurate to us, so we now acknowledge the limitations of this heuristic when it is presented in Section 4. We also provide some argument in favor of its robustness against non-adversarial behavior in the conclusions, in which we briefly argue that evading the heuristic would require a certain sacrifice in usability.

The other main concern (weakness 4) questions what we wanted to demonstrate in Section 5. We now state the conclusions more clearly at the end of this section: in the introduction, we observe that many government agencies seem to be concerned about the ability of criminals to launder money using Bitcoin, and wonder whether or not these fears are grounded. While some people have suggested that criminals might be able to use Satoshi Dice due to its high volume, in Section 5.1 we explain more clearly why bets with and payouts from Satoshi Dice are completely transparent. In Section 5.2, we demonstrate our ability to follow the funds belonging to certain criminals within the Bitcoin economy, in many cases directly to known exchanges. We therefore conclude that, while in the future they could perhaps do more work to keep their flow of bitcoins anonymous, criminals currently provide ample opportunity to follow their illicitly-obtained bitcoins, and thus using Bitcoin to launder large quantities of money does not seem particularly feasible or attractive at present.

Finally, to address the concern that many people are not familiar with how Bitcoin works, we have fleshed out the protocol description in Section 2, and attempted to provide both more technical details and more discussion of the potential anonymity that Bitcoin currently provides via its use of pseudonyms.