

A Denial of Service Attack against Fair Computations using Bitcoin Deposits

Jethro Beekman

July 2014

Abstract

Bitcoin supports complex transactions where the recipient of a transaction can be programmatically determined. Using these transactions, multi-party computation protocols that aim to ensure fairness among participants have been designed. We present a Denial of Service attack against these protocols that results in a net loss for some or all of the honest parties involved, violating those fairness goals.

1 Introduction

Several recent works by Andrychowicz et al. [1, 2] (Protocol “ADMM”) and Bentov and Kumaresan [3] (Protocol “BK”) describe multi-party computation schemes in which Bitcoin deposits are used to ensure fairness. The general idea is that parties in the computation make a deposit at the beginning of the computation, which honest parties will get back in the end. This incentivizes parties to share their result of the computation with the other parties.

In this work, we introduce a Denial of Service (DoS) attack that results in a net loss for honest parties, destroying the incentive for honest parties to participate. In our attack, dishonest parties will turn a profit at the cost of the honest parties, which incentivizes participants to cheat. This undermines the incentive structure of the underlying protocols. In particular, we note that the security models of ADMM and BK did not consider the possibility of network-level DoS. We show how a dishonest party can use network-level DoS against honest parties.

2 Background

ADMM and BK are protocols for secure multi-party computation that are intended to be fair. Traditional multi-party computation has the problem that one or more dishonest parties might be able to learn the result of the distributed computation and then walk away, so the honest parties never learn the result of the computation. A perfectly fair protocol is one where this cannot happen:

intuitively, either everyone learns the outcome of the computation, or no one does. In ADMM and BK, fairness is encouraged monetarily, but not guaranteed. Fairness is accomplished by having all parties initially pay a deposit. Dishonest parties who walk away forfeit their deposit and it is split among the honest parties as compensation, while honest parties receive their deposit back after the computation is finished. This is roughly how fairness and security are defined for ADMM¹ and BK.²

ADMM and BK use Bitcoin to define complex transactions like

$$P_3 \xleftarrow{\tau} \frac{P_1}{q} \xrightarrow{c} P_2$$

which means P_1 posts a Bitcoin transaction depositing q ; then, P_2 can post a Bitcoin transaction satisfying condition C and claiming q before time τ ; otherwise, after time τ , P_3 can post a Bitcoin transaction claiming q . The protocols use a sequence of these transactions to provide fairness.

For example, the 2-party BK protocol [3, §3.1] is defined as

$$P_1 \xleftarrow{\tau_2} \frac{P_1}{q} \xrightarrow{c_1 \wedge c_2} P_2$$

$$P_2 \xleftarrow{\tau_1} \frac{P_2}{q} \xrightarrow{c_1} P_1$$

with $\tau_1 < \tau_2$. If both parties are honest, P_1 satisfies C_1 ³ before τ_1 to receive q from P_2 , after which P_2 can then satisfy $C_1 \wedge C_2$ before τ_2 to receive q from P_1 . This means that no one loses their deposits and everyone learns the result of the computation. If P_1 is dishonest and does not satisfy C_1 in time, P_2 gets q back at τ_1 and later P_1 gets q back at τ_2 . In this case, no one loses their deposits and no one learns the result of the computation. If P_2 is dishonest and does not satisfy $C_1 \wedge C_2$ in time, P_1 has already gotten q at τ_1 and later P_1 gets q back at τ_2 . Here, P_2 learns the result of the computation while P_1 does not, but P_2 has a net loss of q and P_1 has a net gain of q .

ADMM is similar, but uses transactions of the form $P_2 \xleftarrow{\tau} \frac{P_1}{q} \xrightarrow{c} P_1$.⁴

3 Threat model and attack

Our threat model considers an adversary A that participates in a multi-party computation protocol using Bitcoin and is able to perform a network-level Denial

¹“Formally, we say that the protocol is *secure* if for any strategy of the adversary, that controls the network and corrupts the other parties, (1) the execution of the protocol terminates in [some time], and (2) the expected payoff of each honest party is at least negligible.” [2, §IV]

²“Loosely speaking, our notion of fair secure computation guarantees: an honest party never has to pay any penalty; [and] if a party aborts after learning the output and does not deliver output to honest parties, then every honest party is compensated.” [3, §2.1]

³The condition is satisfied by revealing a witness to a commitment [3].

⁴The difference is in the target parties: BK uses transactions of the form $P_1 \xleftarrow{\tau} \frac{P_1}{q} \xrightarrow{c} P_2$.

of Service attack against another party B in the same computation for extended periods of time. This inclusion of control over the network is consistent with the security definition for ADMM.¹

We show how an adversary A that pretends to be honest can turn another honest party B into a dishonest party in the eyes of the protocol by performing a DoS attack on B at the appropriate time.

Let's reconsider the 2-party BK protocol with P_1 being malicious. First, P_1 pretends to be honest and satisfies \mathcal{C}_1 before τ_1 to claim q . Then, P_1 immediately performs a DoS attack on P_2 , lasting at least until τ_2 . P_2 will be unable to post a transaction satisfying $\mathcal{C}_1 \wedge \mathcal{C}_2$ during this time, and at τ_2 P_1 will be able to claim its original deposit q back. P_1 now has $2q$ while it deposited just q at the beginning, for a net gain of q . P_2 lost its deposit, for a net loss of q , even though it might have been intending and trying to satisfy $\mathcal{C}_1 \wedge \mathcal{C}_2$.

Similar attacks work against ADMM and the same protocols extended for more than 2 parties.

4 Discussion

Denial of Service attacks are notoriously hard to defend against. Two potential solutions that future work could focus on are using very large time scales, and decoupling satisfying the condition from claiming the deposit.

Using very large time scales could give a party under attack an opportunity to reroute the message satisfying the condition. This is not a cure-all, as a powerful enough adversary would still be able to maintain the DoS attack for this prolonged period of time. Also, using longer time scales would mean that an honest party would have to wait longer to reclaim the deposit from a dishonest party.

Decoupling satisfying the condition from claiming the deposit would entail entrusting a third party or producing some kind of program. That third party or program would then satisfy a party's condition at the appropriate time, even if it is under attack.

5 Conclusion

In this work, we have demonstrated a Denial of Service attack against two recent fair multi-party computation protocols. This attack both defies the fairness aimed to be provided by these protocols and violates security definitions of those protocols. We highlight several avenues future research could explore.

References

- [1] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Łukasz Mazurek. Fair two-party computations via bitcoin deposits. In *1st Workshop on Bitcoin Research*, March 2014.

- [2] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on bitcoin. In *Security and Privacy (SP), 2014 IEEE Symposium on*, May 2014.
- [3] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In *Advances in Cryptology – CRYPTO 2014*, Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2014. (to appear).