EE374 Internet-scale Consensus **Stanford, Winter 2021**

Homework #2[*]

Due: Tues, 2-February-2021, 11:59pm – Gradescope entry code: `N8XV23`

Please upload your answers timely to Gradescope. Start a new page for every problem. We strongly suggest LaTeX to type your answers. For the programming/simulation questions you can use any reasonable programming language (please no assembly, brainfuck, etc. ☺). Comment your source code and include the code and a brief overall explanation with your answers. A tentative point distribution (in % of the total) is provided in brackets. For most problems there is more than one valid way of solving them!

1. (30%) We did a partial analysis of Nakamoto's private attack on the k-deep confirmation rule in lecture 3, and we will complete it in this problem. As in the lecture, the attack begins at time 0 and it is on first honest block b, at level 1. The adversary has a fraction $\beta < 1/2$ of the mining power, the honest and adversary mine at rate λ_h and λ_a blocks per second respectively, and the network delay is assumed to be 0.

    a) What is the relation between β, λ_h and λ_a?

    b) Let E_m be the event that m adversary blocks are mined before m honest blocks are mined. Using what you learnt in the lecture, present a good bound on the probability of E_m.

    c) Let F be the event that the private attack succeeds on reversing the confirmation of block b under the k-deep confirmation rule. Express this event in terms of the events E_m, $m = 1, 2, \ldots$.

    d) Using your answers from previous parts or otherwise, show that the probability of F decreases exponentially with k. (**Hint:** The union bound may be useful here.)

    e) Simulate the longest chain protocol under Nakamoto's private attack, and estimate the confirmation error probability for $k = 5, 10, 15, 20$ and for adversarial hash power fraction $\beta = 0.3, 0.45$. Compare your results with the analytical bound you obtained in the previous part. (**Hint:** Make sure to repeat the runs of the protocol sufficiently often to generate reliable estimates of the error probabilities.)

[*] Version: 1 – Last update: 24-Jan-2021

2. (40%) Nakamoto's private attack was discussed in the context of reversing a confirmed block at level 1. In this problem, we will consider a more powerful attack applicable to blocks at arbitrary levels. This is an attack which is Nakamoto's private attack combined with a *pre-mining phase*. The attack is focused on reverting a transaction TX included in the block of the public chain at the i-th level.

   - *Pre-mining phase*: Starting from the genesis block, the attacker starts mining blocks in private to build a private chain. When the first honest block h_1 is mined on the genesis block, the attacker does one of two things: i) If the private chain is longer than the public chain at that moment, then the adversary continues mining on the private chain; ii) if the private chain is equal or shorter then the public chain, the attacker abandons the private chain it has been mining on and starts a new private chain on h_1 instead. The attacker repeats this process with all honest blocks $h_2, h_3, \dots h_{i-1}$.
   - *Private attack phase*: After block h_{i-1} is mined, the attacker will start Nakamoto's private attack from the current private chain it is working on, whether it is off h_{i-1} or the one it has been working on before h_{i-1} depending on which is longer.

   Answer the following questions. You may assume the same setting as in Problem 1.

   a) Suppose $\beta = 0.3$. What is the probability that the attacker will switch to h_1 when it is mined? What is the expected level at which the attacker is mining when h_1 arrives?

   b) Suppose honest (h) and adversarial blocks (a) are mined in the order:

   $$a_1, a_2, h_1, a_3, h_2, h_3, a_4, a_5, h_4.$$

   Draw the evolution of the block tree, always including both honest and adversary blocks.

   c) Let L^a_{i-1} be the level at which the adversary is mining just before the $(i-1)$-th honest block arrives. Let $G_{i-1} := L^a_{i-1} - i + 1$ be the advantage the adversary has over the public chain just after that time. The distribution of G_{i-1} depends on i. What happens when $i \to \infty$?

   d) Simulate this attack for large i and estimate the confirmation error probability for $k = 5, 10, 15, 20$ and adversarial hash power fraction $\beta = 0.3, 0.45$. Compare these results with those in Problem 1d). Are there significant differences? Why?

3. (30%) Assume the Ethereum chain where difficulty is adjusted such that on average a new block is created every 15 seconds ($\lambda = \frac{1}{15}$, unit $\frac{1}{\mathrm{s}}$). Suppose it takes Δ seconds to communicate a newly found block of size 20 KBytes to the remaining miners, during which the remaining miners continue to try to mine a new block off of the previous (now old) block.

   a) Why do you think that the mining rate in Ethereum is so much higher than that in Bitcoin?

   b) By either analysis or computer simulation (choose one), estimate the longest chain growth rate as a function of Δ. Assume that honest miners follow the longest-chain rule and break ties by mining on top of the block with the oldest timestamp. You can also assume that there is an infinite number of honest miners each with infinitesimal mining power.

   c) What is the effect of forking on the security and the throughput of the blockchain?

   d) Using the data on node latency in [1, Table 7, *Avg.* for *Ethereum* as measured by the *Dir*ect connection method] or `https://ethstats.net/` (specify which data you have used), estimate the percentage loss in chain growth rate in Ethereum compared to that in the ideal zero-delay case, and the resulting loss in security, i.e., the maximum fraction of adversarial hash power tolerable.

References

[1] S. K. Kim, Z. Ma, S. Murali, J. Mason, A. Miller, and M. Bailey. Measuring ethereum network peers. In *Proceedings of the Internet Measurement Conference 2018, IMC 2018, Boston, MA, USA, October 31 - November 02, 2018*, pages 91–104. ACM, 2018.