COMS 6998-006 (Foundations of Blockchains): Homework #1

Due by 11:59 PM on Thursday, September 23, 2021

Instructions:

(1) Solutions are to be completed and submitted in pairs.

(2) We are using Gradescope for homework submissions. See the course home page for instructions, the late day policy, and the School of Engineering honor code.

(3) Please type your solutions if possible and we encourage you to use the LaTeX template provided on the Courseworks page.

(4) Write convincingly but not excessively. (We reserve the right to deduct points for egregiously bad or excessive writing.)

(5) Except where otherwise noted, you may refer to your lecture notes and the specific supplementary readings listed on the course Web page *only*.

(6) You are not permitted to look up solutions to these problems on the Web. You should cite any outside sources that you used. All words should be your own. Submissions that violate these guidelines will (at best) be given zero credit, and may be treated as honor code violations.

(7) You can discuss the problems verbally at a high level with other pairs. And of course, you are encouraged to contact the course staff (via the discussion forum or office hours) for additional help.

(8) If you discuss solution approaches with anyone outside of your pair, you must list their names on the front page of your write-up.

(9) Some of these problems are difficult, so your group may not solve them all to completion. In this case, you can write up what you've got (subject to (4), above): partial proofs, lemmas, high-level ideas, counterexamples, and so on.

Problem 1

(15 points) In Lecture 2 we studied the *Byzantine broadcast (BB)* problem, where a single node (the sender) has a private input and the goal is to broadcast that input to everyone else. For this problem, you can assume that the private input is either 0 or 1.

A closely related problem is *Byzantine agreement (BA)*. Here, each node $i \in \{1, 2, \ldots, n\}$ has its own private bit $b_i \in \{0, 1\}$. Up to f of the n nodes may be Byzantine, meaning they can behave in arbitrary ways. We are interested in deterministic BA protocols that satisfy the following two properties:

1. *Agreement:* the protocol always terminates with all honest nodes outputting the same bit.
2. *Validity:* if all honest nodes have the same private input $b \in \{0, 1\}$, then the protocol terminates with all such nodes outputting b.

We'll work in the same synchronous model used in Lectures 2 and 3.

(a) (5 points) Give one deterministic BA protocol that satisfies agreement (no matter what f is), and another that satisfies validity (no matter what f is). Prove that your protocols satisfy the property in question.

(b) (10 points) Suppose $f < n/2$. Prove that there exists a deterministic BB protocol satisfying validity and agreement (as defined for the BB problem) if and only if there exists a deterministic BA protocol satisfying validity and agreement (as defined for the BA problem).

Problem 2

(30 points) In this problem, we continue to study the Byzantine broadcast problem. (You can continue to assume that the sender's private input is either a 0 or a 1.) You can assume $n \geq 3$. Consider the following protocol (to be executed by all honest nodes):

$t = 0$: The sender sends its private bit (along with its signature) to all other nodes. The sender then outputs its own private bit and terminates.

$t = 1$: Every non-sender node i echoes what it heard from the sender to all the other non-sender nodes (with i's signature added).

$t = 2$: Every non-sender node collects all the votes it received (up to $n - 1$ votes, with at most one from the sender in round 0 and at most one from each non-sender node from roung 1) and chooses its output by majority vote. (If there's an equal number of votes for both 0 and 1, output 0.) If a node submits multiple votes for different values, all votes by that node are ignored.

For each of parts (a)–(f), provide either a convincing proof or an explicit counterexample.

(a) (5 points) Suppose $f = 1$ (i.e., at most one node is Byzantine, and all the rest are honest). Does this protocol satisfy agreement?

(b) (5 points) Suppose $f = 1$. Does this protocol satisfy validity?

(c) (5 points) Suppose $f = 2$. Does this protocol satisfy agreement?

(d) (5 points) Suppose $f = 2$. Does this protocol satisfy validity?

(e) (5 points) Recall the Dolev-Strong protocol from Lecture 2. Suppose we stop that protocol one time step earlier (with nodes committing to an output at time f rather than time $f + 1$), without changing anything else. Does the protocol continue to satisfy agreement?

(f) (5 points) Does the protocol in (e) continue to satisfy validity?

Problem 3

(15 points) Recall the impossibility result (for BB with $f \geq \frac{n}{3}$ and no PKI) from Lecture 3. In lecture we proved only the special case of $n = 3$ and $f = 1$. Extend this impossibility result to all n and f with $f \geq n/3$.

[Hint: try to avoid redoing the proof from scratch.]

Problem 4

(35 points) Consider again the impossibility result from Lecture 3. Recall that a protocol instructs a node what messages it should send in a given round, as a function of everything the node knows—its private input (if any) and the messages it has received thus far. In a *randomized* protocol, the outgoing messages in a given time step can additionally depend on coin flips performed locally at the node. Coin flips are private, meaning that the result of a coin flip is seen only by the node who flipped it.

(a) (15 points) Suppose $n = 3$ and $f = 1$. Prove that there is a randomized Byzantine broadcast (BB) protocol for which, no matter what the input and adversary strategy, the probability that at least one of validity or agreement is violated is at most 1/3. (The probability is over the random coin flips of all

the nodes. Note that an adversary strategy cannot depend directly on the outcomes of these (private) coin flips.)

[Don't worry about bit complexity issues. E.g., if you find it convenient to have a node sample a real number uniformly from $[0, 1]$, that's fine. You may also assume that a node can communicate a real number using a single message.]

[Hint: It may help to do (b) first, and then reverse engineer a protocol for which your argument is tight.]

(b) (13 points) Suppose $n = 3$ and $f = 1$. Prove that the failure probability of $1/3$ in part (a) cannot be improved upon by any randomized BB protocol.

[Hint: Study carefully the proof from lecture.]

(c) (7 points) Does your solution to Problem 3 apply also to randomized BB protocols, as studied in this problem? I.e., does your argument prove that the impossibility result in (b) holds more generally for all f and n with $f \geq n/3$? Explain why or why not.

Problem 5

(30 points) Consider the impossibility result from Problem 3 (for BB without PKI). (We're back to considering only deterministic protocols.)

(a) (10 points) Suppose we restrict our attention to protocols that can be implemented in polynomial time (meaning there is a fixed polynomial function p such that, in each round, the amount of computation that an honest node must perform is bounded above by $p(n)$, where n is the number of nodes). Suppose we restrict the adversary to run in time polynomial in n. Does the impossibility result still hold? Either way, prove your answer.

(b) (10 points) Suppose we give all n nodes access to a digital signature scheme, with its key generation, signing, and verification algorithms. A protocol can in particular instruct a node to invoke any of these algorithms at some point during its execution. In addition to the assumptions made in part (a), assume that a polynomially-bounded adversary cannot break cryptography (e.g., forge signatures). Does the impossibility result still hold? Either way, prove your answer.

(c) (10 points) In light of (a) and (b), why doesn't the Dolev-Strong protocol contradict the impossibility result from Lecture 3?