COMS 6998-006 (Foundations of Blockchains): Homework #2

Due by 11:59 PM on Thursday, October 7, 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

(20 points) The FLP impossibility result (Lectures 4 and 5) concerns the Byzantine agreement (BA) problem. But what about the Byzantine broadcast (BB) problem discussed in Lecture 2?

(a) (5 points) Show that, in the asynchronous model with $f < n/3$, BA reduces to BB. (That is, given a BB protocol that satisfies validity and agreement, build a BA protocol that satisfies validity and agreement.) In light of the FLP result, what does this imply about the BB problem in the asynchronous model?

(b) (5 points) On Homework #1 you showed that BB reduces to BA in the synchronous model (no matter what f is). Does your reduction work also in the asynchronous model? If not, why not?

(c) (5 points) Give a simple and direct proof that, in the asynchronous model, no deterministic BB protocol always terminates while satisfying both validity and agreement (even for $f = 1$).

[Hint: Consider first the case of an adversarial sender who never sends any messages. What must the other nodes do?]

(d) (5 points) Show that the FLP impossibility result implies that there is no protocol for SMR that satisfies both consistency and liveness in the asynchronous model, even for $f = 1$.

[Hint: like in (a), use a suitable reduction.]

Problem 2

(35 points) Recall again the FLP impossibility result for BA in the asynchronous model.

(a) (5 points) The proof given in lecture is somewhat nonconstructive. Give a direct proof for the $n = 2$ (and $f = 1$) case, with an explicit description of the relevant adversary strategies.

(b) (5 points) A very restricted type of adversary is one who only uses *crash faults*. This means that the only misbehavior allowed by an adversarial node is: at some point in the protocol's execution (at a time decided by the adversary), the node never sends a message again (no matter how long the protocol runs). Explain why the proof from lecture does not immediately apply to crash-fault adversaries.

(c) (15 points) Modify the proof from lecture so that it holds also for crash-fault adversaries.

[If you must, you can rewrite the whole proof. Better would be to explain exactly which steps of the proof need to be changed, and what changes need to be made in each case.]

(d) (10 points) Suppose we assume that each node has its own public key-private key pair, and that all nodes' public keys are common knowledge prior to the start of the protocol. (I.e., the same PKI trusted setup assumption as for the Dolev-Strong protocol.) You can assume that the amount of computation performed by each adversarial node is polynomial (in the number n of nodes); message delivery continues to be completely arbitrary. Does the proof of the FLP impossibility result continue to hold? Give a compelling justification for your answer.

Problem 3

(30 points) The point of this problem is to see how randomization can help mitigate the FLP impossibility result for BA in the asynchronous model. Throughout, we assume a known upper bound f on the number of faulty (Byzantine) nodes, and assume that $f < n/5$. (Recall that the FLP result applies for deterministic protocols even when $f = 1$.)

Consider a protocol in which each node has a local notion of "rounds." Each round will have two phases (with messages sent and received in each phase), followed by a local update step. By "phase 17(b)," for example, we mean the second phase of the 17th round. Every message sent by an honest node i is annotated with the phase (from i's local perspective) to which it belongs.

Each node i maintains a bit x_i, initially set to its private input.

First phase. When node i reaches the (local) phase $r(a)$ (i.e., the first phase of round r), it sends its current bit x_i to all nodes (including itself), annotated as usual with the current phase $r(a)$. Node i then idles in phase $r(a)$ until it receives $r(a)$-phase messages from $n - f$ distinct nodes. (If it receives more than one $r(a)$-phase message from the same (dishonest) node, it ignores all but the first of these.)[1]

Second phase. If more than $(n+f)/2$ of the phase-$r(a)$ messages received agree on a common bit v, then send v to all nodes. Otherwise, send a null message $\perp$ to all nodes. Node i then idles in phase $r(b)$ until it receives $r(b)$-phase messages from $n - f$ distinct nodes.

Local update step.

- If more than $(n+f)/2$ of the phase-$r(b)$ messages received agree on a common bit v, then commit to v as the final output.
- Otherwise, if at least $f + 1$ of the received phase-$r(b)$ messages agree on a common bit v, set x_i to v.
- Otherwise, set x_i to 0 or 1, with 50/50 probability.

[1] The node discards any messages it receives that are associated with earlier rounds that the node has already moved on from. The node remembers any messages it receives that are associated with later rounds, to be taken into account once the node catches up and reaches the relevant round.

After the round-r local update step, a node proceeds to the first phase of round $r+1$.

(a) (7 points) Prove that the protocol is well defined. Specifically, prove that (no matter what the Byzantine nodes and the adversary controlling message delivery do), for every round r, no node will ever receive at least $f+1$ phase-$r(b)$ messages for two different values v.

(b) (7 points) Prove that, whenever the protocol terminates, it satisfies validity.

(c) (9 points) Prove that, whenever the protocol terminates, it satisfies agreement.

(d) (7 points) Prove that, with probability 1 (over the protocol's random coin flips), the protocol terminates in a finite number of steps.

Problem 4

(10 points) Recall the partially synchronous model from lecture: for an a priori known bound Δ and an unknown (adversarially chosen) time GST, a message sent at time t is guaranteed to arrive by time $\max\{t, GST\} + \Delta$. (The description of a protocol can depend on Δ but not GST.) In effect, the network is asynchronous until time GST, after which it is synchronous (with maximum message delay Δ).

One intuition we mentioned in lecture is that an ideal consensus protocol should adapt automatically to message delays, operating at close to the network speed. One formalization of this idea is: we'd like a protocol that works (i.e., satisfies safety and liveness) whenever it is run in the synchronous model with some adversarially chosen (but finite) maximum message delay Δ. (In this "unknown Δ" variant of partial synchrony, the protocol description cannot depend on Δ.)

(a) (5 points) Prove that if there is a Byzantine agreement protocol that guarantees agreement, validity, and (eventual, post-GST) termination in the original "GST version" of the partially synchronous model with f (Byzantine) faulty nodes, then there is also such a protocol in the "unknown Δ version" (i.e., one that satisfies agreement, validity, and termination, no matter what Δ is).

(b) (5 points) Prove that if there is a Byzantine agreement protocol that guarantees agreement, validity, and termination in the "unknown Δ version" of the partially synchronous model with f (Byzantine) faulty nodes (no matter what Δ is), then there is also such a protocol in the original "GST version" (i.e., one that satisfies agreement, validity, and eventual termination).

Problem 5

(35 points) This problem continues our study of Byzantine agreement protocols in the partially synchronous model (the original GST version).

(a) (19 points) Prove that with $n = 3$ nodes and at most $f = 1$ Byzantine node, no deterministic BA protocol satisfies agreement, validity, and eventual termination. Your proof should apply even assuming the availability of PKI.

[Hint: formalize the intuition given in lecture.]

(b) (6 points) Extend the impossibility result in (a) to all $n \geq 3$ and $f \geq n/3$.

(c) (5 points) Compare this impossibility result to the one from Lecture 3. In what way is it stronger?

(d) (5 points) Prove that this impossibility result extends to the SMR problem, in the following sense: if $f \geq n/3$, then no deterministic protocol for the SMR problem satisfies consistency (at all times) and eventual liveness (required only after GST).

Problem 6

(15 points) This problem is the same as the previous one, except with a crash-fault adversary (as defined in Problem 2).

(a) (10 points) Prove that with $n = 2$ nodes and at most $f = 1$ faulty node (with a crash-fault adversary), no deterministic BA protocol satisfies agreement, validity, and eventual termination.

(b) (5 points) Extend the impossibility result in (a) to all $n \geq 2$ and $f \geq n/2$.[2]

[2]Fact: there is also a matching positive result, meaning a deterministic BA protocol that, in the partially synchronous model and with $f < n/2$ crash-fault nodes, satisfies agreement, validity, and eventual termination. (You're well-positioned to prove this result, but this homework is already long enough. . .)