BAR Games

Allen Clement Jeff Napper Lorenzo Alvisi Mike Dahlin
{aclement,jmn,lorenzo,dahlin}@cs.utexas.edu
Laboratory for Advanced Systems Research
Department of Computer Sciences
The University of Texas at Austin

Abstract

This paper describes a general methodology for simplifying the design and analysis of BAR protocols. BAR protocols allow the participation of Byzantine, Altruistic, and Rational players. Because BAR protocols tolerate both arbitrary behaviors by some nodes and selfish behavior by the rest, they are appropriate for service and applications spanning multiple administrative domains. We focus our attention on IC-BFT protocols that (a) guarantee a set of safety and liveness properties to all non-Byzantine nodes and (b) insure that rational nodes follow the protocol faithfully. We rely on existing techniques to show that safety and liveness are maintained when all non-Byzantine nodes follow the protocol. In order to show that rational nodes follow the protocol faithfully, we decomponse the BAR game corresponding to the problem specification into a combination of a n-player benefit game and a collection of n 2-player cost games. We provide a set of sufficient properties to show that the protocol is a CBE in the benefit game and that rational nodes will thus follow the protocol faithfully. We present the first synchronous IC-BFT TRB protocol, basing our protocol design on these properties.

University of Texas at Austin Computer Sciences Technical Report 06-25

1 Introduction

This paper addresses the development of distributed protocols that target systems spanning multiple administrative domain (MADs). In these systems, nodes collaborate to provide some service that benefits each node, without any central authority owning and controlling all nodes: examples of these systems include the growing number of cooperative services built above the peer-to-peer paradigm [1,9,11,18,24, 35, 34]. MAD systems are attractive because their diffused control structure may yield services that are potentially less costly and more democratic than their more centralized counterparts.

A key challenge in building MAD cooperative services is dependability. As in a traditional distributed systems, nodes in a MAD system can deviate from their specification because they are *broken*, on account of bugs, errors in software configuration, or even malicious attacks. But MAD systems add a new dimension: without a central administrator to ensure that all unbroken nodes follow faithfully their assigned protocol, nodes may deviate from their specification also because they are *selfish* and are intent on maximizing their own utility. In fact, since it is not obvious how to bound the number of nodes that will opt to deviate selfishly from the protocol (especially if all is required of them to do so is simply to download some software from a web site), dependable MAD services must be designed under the assumption that potentially *all* nodes may deviate from their specification.

To make things worse, MAD services do not fit nicely within any of the traditional models used in distributed systems. It is theoretically possible to model all deviant nodes as Byzantine [7, 8, 17, 19, 21, 32, 33, 36, 40], but not much of interest can be said about a distributed system in which one cannot reasonably assume a fixed (and low!) threshold on the number of Byzantine nodes. Traditional game theoretic approaches [30] also fall short: they can successfully model systems in which every node is selfish, but they are not equipped to deal with Byzantine nodes.

In [2], we propose a new model, called BAR, that acknowledges the likely coexistence of Byzantine and selfish behaviors in any MAD service. The model owes its name to the initials of the three classes of nodes (Byzantine, Altruistic, and Rational) that it explicitly considers. *Byzantine* nodes behave arbitrarily: they may deviate in any way for any reason, regardless of the local or global consequences. *Altruistic* nodes follow a given protocol faithfully without consideration of their self interest. *Rational* nodes behave selfishly and will deviate from a given protocol if doing so improves their own utility.

While the BAR model seems to capture fairly naturally the characteristics of MAD services, its value cannot be measured only in terms of its expressiveness: ultimately we are interested in building dependable MAD services, not just in talking about them. It is consequently crucial to determine how hard it is design BAR tolerant protocols, i.e., protocols that will provably maintain their safety and liveness property under the BAR model. In this sense, the coexistence of selfish and Byzantine nodes appears fundamentally challenging: the misdeeds of Byzantine nodes are typically tolerated by requiring extra work of non-Byzantine nodes, and selfish nodes are unlikely to step up as volunteers to do this extra work.

In this paper, we view the problem specification in the BAR model as a game—a *BAR game*—through which it is possible to analyze the interactions of a set of n players corresponding to the nodes in the system. Reasoning about BAR games is hard, since it requires considering an exponential number of n-player strategies. A BAR protocol provides only a suggested strategy for each player to follow. Rational players independently choose the strategy that gives them maximal individual net benefit, whether that strategy is suggested by the protocol or derived independently, and they participate in the game only

if they expect to "win", i.e., to receive a net benefit from their participation. For a protocol designer, winning consists instead of orchestrating a distributed protocol that provably meets the problem's specification in the BAR model. A compelling correctness criterion is Incentive-Compatible Byzantine Fault Tolerance (IC-BFT) [2]. IC-BFT protocols provably meet the properties of the problem specification by suggesting strategies that, informally, are seen by Rational players as the most beneficial to them.

The main contribution of this paper is a general methodology for simplifing both the design and the analysis of IC-BFT protocols. The idea is a simple application of the divide and conquer principle: we separate the problem of designing an incentive compatible protocol from the complementary task of showing that the resulting protocol is Byzantine fault tolerant. The literature provides several examples to help us with the latter problem—this paper focuses on how to separate the two concerns cleanly and on how to address the first.

We use Cooperative Byzantine Equilibria (CBE) to show that rational players will not deviate from their assigned protocol. A protocol is a CBE in a game if, informally, each rational player does not find benefit in deviating unilaterally a specified the protocol. Finding a CBE is equivalent to identifying a protocol that corresponds to a local maximum in the individual benefit of each player. While this is simpler that identifying a global maximum, it is still challenging to achieve CBE protocols in a model, like BAR, where rational nodes will be tempted to shirk any extra work required to prevent Byzantine nodes from causing harm.

To simplify the design of a CBE protocol in the BAR model, we cast a BAR game as the combination of a n-player *benefit game* and a collection of n two-player *cost games*.

For a given protocol, the strategy employed by the two nodes involved in a particular cost game is defined by the set of messages that the two nodes exchange. An *ideal cost game* is a cost game that has only two viable strategies, cooperate or defect, and therefore effectively encapsulates in a binary choice much of the complexity that arises from selfish behavior. The n cost games determine the pairwise strategies used by the n players in the benefit game: nodes that play an ideal cost game need only choose in the benefit game whether to cooperate or defect with other players, including those that are Byzantine. This simple choice greatly reduces the number of strategies available to a benefit game player, simplifying the task of proving that the protocol is a CBE.

We show that if it is in the interest of rational players to maintain the safety and liveness properties associated with a problem specification, a protocol that meets the following conditions is a CBE for the associated benefit game: i) the protocol implies that cost games are ideal and ii) the protocol is minimal, in the sense that the loss of a single message between any pair of non-Byzantine nodes may lead to the violation of safety or liveness. Furthermore, we provide a set of conditions that are sufficient to ensure that a protocol will imply ideal cost games.

Finally, we demonstrate our methodology by presenting the first incentive compatible protocol for synchronous Terminating Reliable Broadcast (TRB). We show how our guidelines for protocol design can lead to an incentive compatible version of the classic Dolev and Strong TRB protocol [12] with the same asymptotic message complexity as the original.

2 Related Work

Monderer [22] presented the first framework for converting mechanisms defined in Game Theory to protocols. Nianb introduced Algorithmic Mechanism Design in [29] and Feigenbaum extended these concepts to Distributed Algorithmic Mechanism Design as a framework for building distributed systems out of rational participants [16]. Since then a variety of authors have modeled and built distributed systems using the constructs of game

theory [3, 10, 14, 15, 26, 25, 28, 27, 37, 38, 39]. Unfortunately, these solutions are brittle in the face of Byzantine deviations.

Eliaz defined the first notion of equilibrium to address Byzantine players in a rational context [13]. He presents an *ex ante* equilibrium, i.e. one that in retrospect is used to show that the strategy profile is an equilibrium despite the actions of faulty nodes. In Byzantine Nash Equilibria no rational player has incentive to deviate from its chosen strategy in the presence of Byzantine players. Aiyer et al. [2] introduced Byzantine Nash Equilibrium (BNE) as a way to analyze the behavior of players in BAR games; Moscibroda et al. [23] formally defined BNE. BNE provide information about a specific strategy profile, but do nothing to address the question of how rational players coordinate on that strategy profile. In this paper we make explicit the fact that protocols provide an external assignment of strategies to players and rely on correlated equilibrium [4, 31], providing some framework for the coordination of strategies.

Aiyer et al. [2] introduced the BAR model and presented a p2p backup system. One component of their system is a replicated state machine based on three phase commit and the TRB safety properties. In this paper we provide the first complete analysis of TRB under the BAR model.

3 Model

We consider the problem of designing communication protocols for MAD services in the context of the BAR model. In general, a problem specification under the BAR model consists of i) a set of *desired properties* that motivate rational players to participate and ii) a set of *design properties* that specify the solution to the problem. Proposed solutions are analyzed in the context of a *BAR game* consisting of a set of BAR players, a collection of possible strategies, and a payoff function mapping strategy profiles to net benefit, or *utility*. A *strategy* is a complete plan of action for a player, while a *strategy profile* is a mapping of a strategy to each player in the game. A *payoff function* maps strategy profiles to utility for each player: the measure of benefit minus the measure of costs. Benefit is measured by the degree to which the outcome of a strategy profile meets the desired properties specified by the problem with respect to a player. Costs instead are incurred by a player through the actions that constitute its strategy. A BAR protocol is a distinguished strategy profile. We consider a BAR protocol to be correct if it is Incentive-Compatible Byzantine Fault-Tolerant (IC-BFT) with respect to the problem specification; that is, the strategy profile is a Correlated Byzantine Equilibrium (CBE) with respect to the BAR game, and the design properties are met when the strategy profile is followed by all non-Byzantine players. A strategy profile is a CBE with respect to a BAR game if (a) the strategy profile is a Byzantine Nash Equilibrium in the game and (b) the strategy profile is specified by an external source. A CBE provides a certain amount of coordination between players in determining which strategies should be played.

The problem of implementing typical communication primitives such as Consensus and Terminating Reliable Broadcast (TRB) in MAD environments implies that the design properties of the problem are the safety and liveness properties of the communication primitive. We assume that the desired property for each non-Byzantine player is that the outcome meets the design properties of the system, a condition we call the *safety benefit assumption*. A strategy in the resulting BAR game is a mapping of messages received from all players to messages sent to each player. To measure benefit, we assume that benefit accrues to players iff the desired properties are met. In the context of communication protocols, costs are measured by messages sent. While players perform computational actions such as generating and verifying signatures and storage actions such as queueing messages for later process-

ing, we observe that these actions correspond to communication and consequently roll the costs of these actions into the costs of sending the associated messages.

We assume synchronous networks with reliable links. Each pairwise link between players can support at most k messages per round. We introduce some terminology to describe message sequences exchanged as part of a communication protocol. A message sequence between two players is *valid* with respect to a strategy execution if the sequence would be sent if all players were Altruistic; otherwise the message sequence is *invalid*. All valid message sequences are called *acceptable*; otherwise they are *unacceptable*. A message sequence is *possible* if the sending node is capable of sending the message sequence given the inputs to the strategy. We assume that cryptographic signatures are secure and cannot be forged. For example, a message sequence containing a message signed by player is not possible for a different player if she has not previously received the signed message.

4 BAR Games

In this section we focus on the IC portion of showing that a BAR protocol is IC-BFT. Showing that the protocol is a CBE for the relevant BAR game is sufficient to show IC.

In Section 4.1 we examine the cost game in more detail. We define a set four *sufficient conditions* for a cost game to act as an *ideal cost game* (ICG) that is easy to embed in a benefit game. To simplify protocol construction, in that section we also define four *protocol properties* that are sufficient to meet these conditions. In Section 4.2 we provide an additional property that, in conjuction with these conditions on the underlying cost game and the safety benefit assumption, is sufficient for the benefit game to provide a CBE.

4.1 The Cost Game

The *cost game* is a two-player game that encapsulates the decision of how to communicate with a specific other player. The possible strategies in the cost game are sending any of the message sequences defined as acceptable by the protocol or an arbitrary unacceptable message sequence. Players expect benefit $b > 0$ when their partner cooperates by sending a non-empty acceptable sequence. The precise value of b is specified by the overarching benefit game and corresponds to the share of global benefit that is attributed to this player. Costs incurred during the cost game are based on the messages sent during the game, for simplicity of presentation, we assume that any non-empty sequence of messages sent during a game incurs cost c where $0 < c < b$.

We define four properties of an *ideal cost game* (ICG) that can be embedded in the benefit game.

CG-1 **Acceptable Implies Valid.** Sending the valid sequence dominates sending any other possible acceptable sequence.

A receiver cannot distinguish between a valid or acceptable-but-invalid message sequence. CG-1 ensures that the most beneficial acceptable sequence is always the valid sequence and that a rational node will consequently never send acceptable-but-invalid message sequences.

CG-2 **Cooperate-Defect Game.** The cost game is a two-player binary cooperate-defect game, and the def

In a binary cooperate-defect game, a node will choose either to *cooperate* by sending an acceptable message sequence or to detectably *defect* by sending an unacceptable message sequence. This property simplifies the strategy space of the benefit game.

CG-3 **Cooperative Equilibrium.** The cost game has a cooperative equilibrium and if the game is non-degenerate the defect strategy has zero cost.

Note that the ideal cost game might require repeated instances of an underlying cost game to reach the cooperative equilibrium—for example, when using punishment for previous defection to incentivize future cooperation. In this case, each assignment of benefit in the benefit game is associated with one play of the ideal cost game but with repeated plays of the underlying cost game corresponding to the protocol. Utility of an ICG is reported as the average utility over the collection of iterated cost games. The zero cost defect strategy implies that a defecting Rational player will always send the empty sequence in the non-degenerate cost game.

In the rest of this section we consider protocol properties that are sufficient to insure CG-1 through CG-3 in turn.

4.1.1 Acceptable Implies Valid

The acceptable implies valid property of the ideal cost game requires that any acceptable message sequence sent by a Rational player be valid. The following protocol property is sufficient to ensure CG-1:

P-1 **Cheap Validity.** The valid sequence specified by the protocol is the cheapest acceptable sequence that can be sent for a given input.

Theorem 1. *Suppose a protocol has property P-1. Then, CG-1 holds for the corresponding cost game.*

Proof. P-1 states that for any input, the protocol specifies the optimal cooperate strategy, that is, to send the cheapest acceptable sequence. □

Typically, nondeterminism and incomplete knowledge of global state prevent nodes from differentiating between received valid and acceptable-but-invalid message sequences. Consequently, a Rational player may choose to exploit her partner by sending acceptable-but-invalid sequences that are less expensive than the valid sequence. Such deviations may have disastrous effects for the design properties of the protocol. Protocols designed to have property P-1 prevent Rational players from deviating in this undetectable fashion through careful construction of their message sequences. For example, in our IC-BFT TRB protocol presented in Section 5, to prevent Rational players from hiding behind the excuse that they have not heard from the leader, we require nodes to send "junk" messages that artificially inflate the cost of professing ignorance.

4.1.2 Cooperate-Defect Game

The cooperate-defect game property of the ideal cost game requires that there be a single viable cooperate strategy and a single viable defect strategy available to Rational players.

Theorem 1 implies that sending the valid message sequence is a dominant cooperate strategy. We rely on the following property to identify a dominant defect strategy:

P-2 **Zero-Cost Defect.** There exists a valid non-empty sequence that player A sends to a player B that defects by sending no messages to A.

P-1 makes the cheapest acceptable message sequence the valid one and P-2 makes the most beneficial unacceptable message sequence the empty sequence. In particular, this condition states that a player can obtain a benefit outcome in a given instance by defecting at no cost, reducing the choice of defect strategies to this cheapest message sequence. This condition does not imply that the defection is undetectable, but rather that the complete acceptable sequence will be sent before the detection can be detected.

Theorem 2. *Suppose a protocol has properties P-1 and P-2. Then, the corresponding cost game is a binary cooperate-defect game as specified by CG-2.*

Proof. Given a single benefit, a strategy dominates by having a lower cost than all other strategies. P-1 implies that the valid messages

specified by the protocol have the lowest cost of all cooperate strategies that send acceptable messages. By P-2 the benefit received by rational player A can be obtained by sending no messages if the player is willing to detectably defect. Sending no messages minimizes costs, resulting in a dominant defect strategy. □

P-1 subtly impacts the acceptability of the empty sequence:

Lemma 1. *Suppose P-1 holds for a protocol. The empty sequence is a degenerate specification that is either the only valid sequence or is unacceptable.*

Proof. Suppose the empty sequence and sequence S with cost $c > 0$ are both acceptable. Since the empty sequence has cost 0, Rational players will never send S, even when it is valid, contradicting P-1.

□

Lemma 1 splits the set of cooperative strategies into two classes based on the acceptability of the empty sequence: in *degenerate strategies* the empty sequence is the dominant (and only) cooperate strategy and in *non-degenerate strategies* the empty sequence is unacceptable and thus considered defection. These two classes of strategies result in three exemplar cooperate-defect games: *degenerate games* in which both players have degenerate strategies, *semi-degenerate games* in which only one player has a degenerate strategy, and *non-degenerate games* in which both players have non-degenerate strategies. Figure 1 shows the base payoff matrices for these three game forms.

4.1.3 Cooperative Equilibria

In this section we show that degenerate and non-degenerate cooperate-defect games have cooperative equilibria and thus fulfill CG-3. We also provide a protocol property sufficient to ensure that there are no semi-degenerate cost games.

Degenerate cost games are the easiest case to consider as cooperation strictly dominates defection, and thus the cooperative equilibrium is trivial identify.

Lemma 2. *Degenerate cost games fulfill CG-3.*

Proof. In degenerate cost games the cooperative strategy and defect strategy do not have identical costs, but neither provides benefit by definition. Since the cooperate strategy specifies sending no messages, defection requires sending a message, incurring non-zero costs. Hence, degenerate cost games have a cooperative equilibrium. □

We eliminate semi-degenerate games from consideration by introducing the following protocol property guaranteeing that all cost games are either degenerate or non-degenerate:

P-3 **Mutual Communication.** If there is an acceptable non-empty message sequence sent from A to B, then there is an acceptable non-empty message sequence that contains at least one non-padding message sent from B to A.

A *non-padding message* is fundamentally related to the protocol design properties. A *padding message* does not impact design properties and exists for the purpose of regulating selfish behavior. In our TRB protocol in Section 5, the $\perp$ messages are padding messages.

Lemma 3. *If P-1 through P-3 hold, then there are no semi-degenerate cost games in the protocol.*

Proof. By Theorem 2 the cost game is a cooperate-defect game. Lemma 1 implies that strategies are either degenerate or non-degenerate. P-3 requires that if one strategy in the communication game is non-degenerate, then the other is also non-degenerate. Hence a cost game cannot be semi-degenerate. □

We now address non-degenerate cost games. As a side effect of P-2 and Theorem 2, the non-degenerate game is an instance of the well known Prisoner's Dilemma [5, 6]. Tit-for-tat

	Coop	Def
Coop	**0, 0**	$0, -c$
Def	$-c, 0$	$-c, -c$

(a) degenerate game

	Coop	Def
Coop	$b, -c$	**0, 0**
Def	$b-c, -c$	$-c, 0$

(b) semi-degenerate game

	Coop	Def
Coop	$b-c, b-c$	$-c, b$
Def	$b, -c$	**0, 0**

(c) non-degenerate game

Figure 1: (a) Degenerate game where both players have a degenerate strategy. (b) Semi-degenerate game where the row player has a degenerate strategy and the column player has a non-degenerate strategy. (c) Non-degenerate game in which both players have a non-degenerate strategy. The quadrants in bold are the equilibrium solutions for the games.

strategies may provide a cooperative equilibrium in the indefinitely repeated Prisoner's Dilemma [30]. The following property is true of protocols that implement tit-for-tat:

P-4 **Shunning.** If player B sends player A an unacceptable sequence, A will shun B by sending no messages to B during an indefinite number of future exchanges.

Note that a rational node will begin shunning in a game as soon as she detects a defecting partner in order to reduce costs. The ideal cost game consists of an indeterminate length sequence of cost games, and implements infinite tits-for-tat. The minimum number of tits-per-tat can be found efficiently [20], allowing a fixed-length shunning mechanism within the ideal cost game that avoids the so-called dismal valley, although the details are beyond the scope of this paper.

Lemma 4. *If P-1 through P-4 hold, then non-degenerate cost games fulfill CG-3.*

Proof. If the protocol corresponds to a non-degenerate cost game, Theorems 1 and 2 implied by P-1 through P-2 hold that the game is a two-player cooperate-defect game, and P-2 implies that this game has the form of the Prisoner's Dilemma. The Prisoner's Dilemma with an indefinite horizon has a cooperative equilibrium in the presence of the punishment mechanism specified by P-4 [30], guaranteeing the cooperative equilibrium. P-2 guarantees that the zero cost defect strategy exists and is dominant, fulfilling CG-3. □

Finally, we prove that a protocol meeting the previous set of conditions corresponds to a cost game with a cooperative equilibrium as specified by CG-3.

Theorem 3. *Suppose P-1 through P-3 hold for a protocol. A repeated cost game with an indefinite horizon corresponding to the protocol has a cooperative equilibrium.*

Proof. Lemma 1 implies the cost game has three forms. By Lemma 3 there are no semi-degenerate cost games, leaving two cases to consider. By Lemma 2, degenerate cost games fulfill CG-3 and by Lemma 4 non-degenerate cost games fulfill CG-3. □

4.2 The Benefit Game

The *benefit game* is an n-player game where the strategies require choosing the strategies to play in the n 2 player cost games. If the

To play the benefit game, each player associates a share of the expected global benefit with the cost game associated with each other participant, distributing the benefit across a collection of two-player games. Players assign a share equal to the average expected benefit for reaching the cooperative equilibrium where rational players expect both players to cooperate. If the share of expected benefit is sufficient, then the player cooperates in the cost game; otherwise the player defects.

Utility is assigned based on the total amount and configuration of cooperation by players during the game. The benefit game is analyzed in expectation: for a given strategy profile, the expected utility for a player is the utility that would be received by that player in the BAR game when all cooperating non-Byzantine players send valid message sequences. Within this construction, Rational players seek to cooperate with the minimum

number of other players to insure an optimal benefit, as follows:

BG-1 **Correlated Byzantine Equilibrium.** The strategies implied by the protocol are a correlated Byzantine equilibrium for the benefit game.

We now consider how the ideal cost game—a cost game that meets CG-2, CG-1, and CG-3—is used in the benefit game and provide guidelines for designing protocols such that BG-1 holds. We say that a non-Byzantine player *participates* in an ideal cost game by playing a the cooperative strategy. First, we show that rational players participate in ideal cost games if they choose to play the game:

Lemma 5. *In the benefit game, non-Byzantine players will behave as Altruistic players if they participate in a non-degenerative pairwise ideal cost game.*

Proof. Altruistic nodes trivially behave Altruistically. If a rational node participates in a pairwise ideal cost game, then it plays the cooperative equilibrium strategy that holds by CG-3. By CG-1 the cooperative strategy sends valid sequences, matching the behavior of Altruistic nodes. □

To show that BG-1 holds for the benefit game, it remains to be shown that rational players participate in enough ideal cost games to obtain the protocol-specific benefit and that any unilateral deviation results in less utility. We show both by restricting protocols with the following condition:

P-5 **Minimal Communication.** If non-Byzantine player A shuns non-Byzantine player B by sending an unacceptable empty sequence, then the design properties of the protocol are violated.

In the presence of the safety benefit assumption, P-5 ensures that the protocol is fragile with respect to faulty behavior. Recall that the ideal cost game forces non-Byzantine players to choose between behaving Altruistic by cooperating or defecting by sending an unacceptable sequence, resulting in indefinite shunning between the players. Since the defect strategy is dominant, there is only one rational non-cooperative strategy, implying the defect strategy is the sole form of deviation open to a rational player. The defect strategy sends an unacceptable empty sequence. When P-5 holds, the protocol is designed to tolerate defection by the f Byzantine players, but with the addition of a defection by a rational player, the cost game might have $f + 1$ faults in the system when the shunned node cannot communicate with his partner so that design properties of the protocol is violated. When design properties are violated, all nodes do not receive benefit, motivating rational behavior as follows:

Theorem 4. *Suppose the cost games are ideal and the protocol has the property P-5. The benefit game has a CBE as specified by BG-1.*

Proof. The protocol specifies the ideal cost games a Rational player should participate in. If a Rational player's cooperate strategy is degenerate for an ideal cost game, then defection results in increased cost without an increase in benefit. Rational players will consequently not defect in those games.

If a Rational player's cooperative strategy is non-degenerate in an ideal cost game, then defection through sending no messages results in decreased cost with no decrease in benefit. However, if the strategy is non-degenerate then by P-5 it may result in the design properties being violated. By the safety benefit assumption violation of the design properties prevents Rational players from receiving benefit, thus decreasing the utility of the rational player. □

5 TRB in the BAR model

The Terminating Reliable Broadcast (TRB) problem states, informally, that a message broadcast by a distinguished sender must eventually be delivered by all non-Byzantine nodes. If the sender is faulty, nodes may deliver

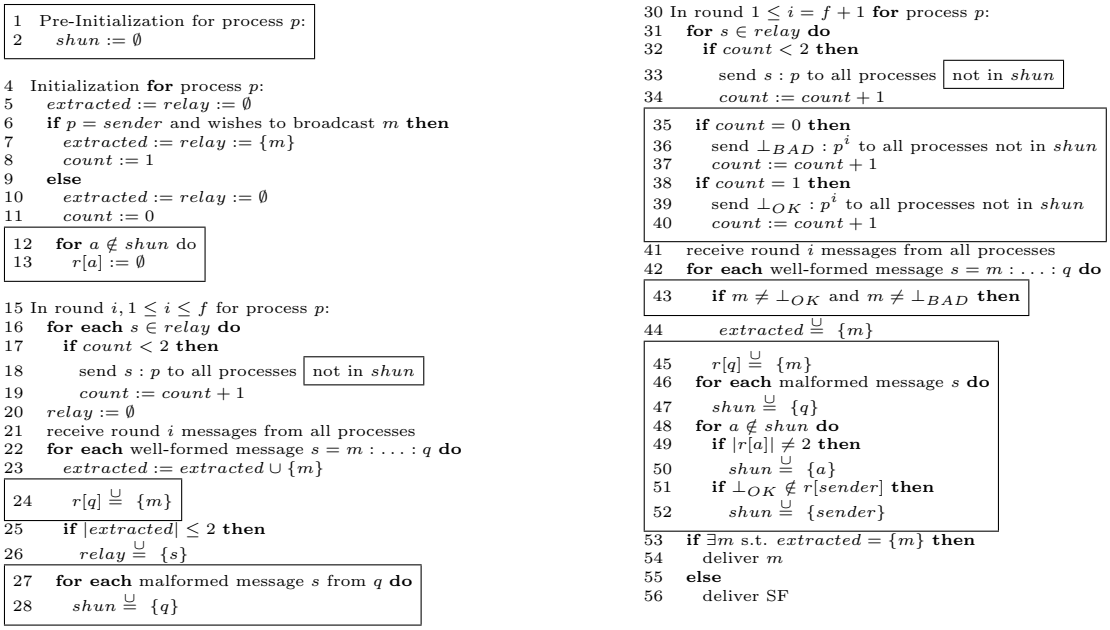

Figure 2: Incentive Compatible TRB with authenticated messages. The Pre-Initialization phase occurs once at the beginning of the benefit game. The rest of the protocol represents an instance of the cost game. The notation $\overset{\cup}{=}$ indicates assignment to the variable on the left-hand side, the variable's value (before assignment) unioned with the right-hand side, the notation $x : p^i$ indicates x signed i times by process p. The boxed portions are the additions to Dolev-Strong to meet P-1 through P-3.

sender-faulty (SF) to guarantee termination. TRB is a classic problem in distributed systems and we believe it represents an interesting case study for the construction of BAR protocols.

Dolev and Strong present a variation of Lamport's TRB protocol for Byzantine failures with message authentication [19] that is optimal with respect to the number of messages sent [12]. Figure 2 presents the Dolev-Strong algorithm—the boxes contain the additional code to create an IC-BFT version.

Players maintain two sets of messages: *relay*, containing messages that should be sent to other players and *extracted*, containing deliverable messages. To reach agreement, a message $m \in \mathcal{M}$ is delivered only when *extracted* $= \{m\}$; otherwise, SF is delivered. The $f + 1$ rounds of the protocol are designed to ensure all correct players have similar *extracted* sets. Byzantine players are thwarted if they attempt to inject new values by requiring each player to sign a received message (including previous signatures) before sending it to another player. A well-formed message then has enough distinct signatures to prove that it has passed through distinct players every round. The protocol lasts $f+1$ rounds to guarantee that every chain contains at least one correct player.

Initially, only the sender possesses the message m in the *relay* and *extracted* sets. Each player proceeds to execute $f+1$ rounds in which she *i)* removes, signs, and sends all values in *relay*, *ii)* receives messages from all players, and *iii)* adds well-formed, newly received messages to *extracted* and *relay*. Since SF will be delivered if *extracted* does not contain a single value, it is sufficient to send only two well-formed messages to every other non-Byzantine process, regardless of the number of distinct values received. The protocol obtains optimal message complexity taking steps (see lines 17, 19, 25, 32, and 34) to send only two distinct values. In the final $f + 1$ round, each player then decides a value based on *extracted* as described above.

In the rest of this section, we show that the modified protocol meets properties P-1 through

P-5. Theorem 4 then implies that the modified protocol is a Correlated Byzantine Equilibrium. Finally, we show that the modified protocol maintains the aka safety and liveness conditions of TRB as specified in [12].

Cheap Validity. In the original protocol, values are repeated only when received; a player might receive no values and correctly send nothing during the protocol. Thus, a rational player could always remain silent and rely on other players to do the work of disseminating the sender's message. To eliminate this acceptable silence, we first augment the protocol with $\perp$ messages that can only be sent in round $f + 1$. Adding these messages requires modifying the definition of well-formed to include $\perp_{OK}$ and $\perp_{BAD}$ received in round $f + 1$ and signed by the same node $f+1$ times. Once these modifications have been made, the acceptable sequences containing only well-formed messages are:

S-1 Two distinct non-$\perp$ values

S-2 $\perp_{OK}$ and a non-$\perp$ value

S-3 $\perp_{OK}$ and $\perp_{BAD}$

Since these are the only 3 acceptable sequences, rational players will choose to send the cheapest possible sequence during each cost game. We restrict the messages $m \in \mathcal{M}$ that may be sent by the sender to have size at most $\log |\mathcal{M}|$ and construct $\perp_{OK}$ and $\perp_{BAD}$ to have size $1+\log |\mathcal{M}|$. As a result of this construction, rational players will send S-1 in preference to S-2 and S-2 in preference to S-3, fulfilling P-1.

Zero-Cost Defect. S-3 allows for a zero-cost defect, fulfilling P-2. Since both $\perp_{OK}$ and $\perp_{BAD}$ are sent in round $f + 1$, a cooperating player will send an acceptable sequence before detecting the deviation.

Shunning. Unacceptable sequences in the modified TRB protocol do not contain two valid messages as described by S-1 through S-3. The modified protocol shuns other players that either send malformed messages or otherwise do not send two valid messages by round $f+1$. Players in the former set are added in any round (lines 27–28 and 46–47 in Fig. 2), while the latter are added at the end of round $f + 1$ (lines 48–50). Messages are never sent to players in the *shun* set, effecting the punishment specified by P-4.

Mutual Communication. The communication between players is symmetric in the modified TRB protocol—any sequence that can be sent by A to B can also be sent by B to A, implying P-3.

Minimal Communication. We show by counterexample that if a non-Byzantine player shuns another player, the agreement property of TRB is violated—that is, two non-Byzantine players may deliver different values. Suppose $f = 1$, $n = 3$, the sender p_B is Byzantine, and non-Byzantine player p_{nB} is shunning the other non-Byzantine player p_s. Further, p_B sends $v : p_B$ only to p_{nB} in the first round and then halts. In the second and final round, p_{nB} will not send any messages to the shunned player p_s. According to the modified protocol, player p_{nB} delivers v in round 2, while player p_s must deliver SF because he has not received any messages. This violation of the agreement property fulfills property P-5.

Safety and Liveness. Suppose all non-Byzantine players follow the modified TRB protocol, implying no player will send a defect message sequence or consequently, be shunned by a rational player. We argue that using the modified protocol, all non-Byzantine players will deliver the same value as if all non-Byzantine players were following the Dolev-Strong protocol [12]. The message delivered in both Dolev-Strong and the modified protocol depends upon the makeup of the *extracted* set. We note that the guard at line 43 is the sole line in the boxed portions of Figure 2—the modifications to Dolev-Strong—that affects the sets *relay* or *extracted* containing values sent or received. The additional guard prohibits $\perp$ messages in the *extracted* set representing values received, while the guard is unnecessary at line 23 because a well-formed message in rounds 1–f cannot consist of a $\perp$ message. Hence, the val-

ues sent and received in the modified protocol are identical to a correct execution of the original Dolev-Strong protocol, implying the delivered value will also be identical.

6 Conclusion

We believe that expressing a BAR game into a combination of an n player benefit game and n pairwise cost games can simplify significantly the difficult task of designing IC-BFT protocols. We have identified a set of sufficient conditions that can help designing a CBE protocol for the benefit game, ensuring that the protocol will be followed by all rational nodes. Finally, we have shown the practicality of our approach by deriving through it the first IC-BFT synchronous TRB protocol.

References

[1] E. Adar and B. Huberman. Free riding on gnutella. Technical report, Xerox PARC, Aug. 2000.

[2] A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J.-P. Martin, and C. Porth. BAR fault tolerance for cooperative services. In *Proc. 20th SOSP*, Oct. 2005.

[3] A. Akella, S. Seshan, R. Karp, S. Shenker, and C. Papadimitriou. Selfish behavior and stability of the internet: a game-theoretic analysis of tcp. In *Proc. SIGCOMM*, pages 117–130. ACM Press, 2002.

[4] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[5] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

[6] R. Axelrod. The evolution of strategies in the iterated prisoner's dilemma. In L. Davis, editor, *Genetic Algorithms and Simulated Annealing*, pages 32–41. Morgan Kaufman, 1987.

[7] G. Bracha and S. Toueg. Asynchronous consensus and broadcast protocols. *J. ACM*, 32(4):824–840, 1985.

[8] M. Castro and B. Liskov. Practical Byzantine fault tolerance and proactive recovery. *ACM Transactions on Computer Systems*, 20(4):398–461, 2002.

[9] B. Cohen. The BitTorrent home page. http://bittorrent.com.

[10] B. Cohen. Incentives build robustness in BitTorrent. In *Proc. 2nd IPTPS*, 2003.

[11] L. P. Cox and B. D. Noble. Samsara: honor among thieves in peer-to-peer storage. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles*, pages 120–132, 2003.

[12] D. Dolev and H. R. Strong. Authenticated algorithms for byzantine agreement. *Siam Journal Computing*, 12(4):656–666, Nov. 1983.

[13] K. Eliaz. Fault tolerant implementation. *Review of Economic Studies*, 69:589–610, Aug 2002.

[14] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *J. Comput. Syst. Sci.*, 63(1):21–41, 2001.

[15] J. Feigenbaum, R. Sami, and S. Shenker. Mechanism design for policy routing. In *Proc. 23rd PODC*, pages 11–20. ACM Press, 2004.

[16] J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *Proc. 6th DIALM*, pages 1–13. ACM Press, New York, 2002.

[17] J. Garay and Y. Moses. Fully Polynomial Byzantine Agreement for $n>3t$ Processors in $t + 1$ Rounds. *SIAM J. of Computing*, 27(1), 1998.

[18] A. Habib and J. Chuang. Incentive mechanism for peer-to-peer media streaming. In *12th IEEE International Workshop on Quality of Service.*, 2004.

[19] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.

[20] M. L. Littman and P. Stone. A polynomial-time Nash equilibrium algorithm for repeated games. *Decision Support Systems*, 39:55–66, 2005.

[21] D. Malkhi and M. Reiter. Byzantine quorum systems. *Distributed Computing*, 11(4):203–213, 1998.

[22] D. Monderer and M. Tennenholtz. Distributed games: from mechanisms to protocols. In *Proceedings of the sixteenth national conference on Artificial intelligence and the eleventh Innovative applications of artificial intelligence conference innovative applications of artificial intelligence*, pages 32–37. American Association for Artificial Intelligence, 1999.

[23] T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proc. 25th PODC*, 2006.

[24] A. Muthitacharoen, R. Morris, T. M. Gil, and B. Chen. Ivy: A read/write peer-to-peer file system. In *Proceedings of 5th Symposium on Operating Systems Design and Implementation*, 2002.

[25] T.-W. Ngan, A. Nandi, and A. Singh. Fair bandwidth and storage sharing in peer-to-peer networks. In *First IRIS Student Workshop*, Cambridge, Massachusetts, Aug. 2003.

[26] T. W. Ngan, D. Wallach, and P. Druschel. Enforcing fair sharing of peer-to-peer resources. In *Proc. 2nd IPTPS*, 2003.

[27] T.-W. Ngan, D. S. Wallach, and P. Druschel. Incentives-compatible peer-to-peer multicast. In *2nd Workshop on Economics of Peer-to-Peer Systems*, 2004.

[28] T.-W. J. Ngan, A. Nandi, A. Singh, D. Wallach, and P. Druschel. On designing incentives-compatible peer-to-peer systems.

[29] N. Nisanb and A. Ronenc. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, April 2001.

[30] M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[31] C. H. Papadimitriou and T. Roughgarden. Computing equilibria in multi-player games. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 82–91, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[32] M. Reiter. The Rampart toolkit for building high-integrity services. In *Dagstuhl Seminar on Dist. Sys.*, pages 99–110, 1994.

[33] R. Rodrigues, M. Castro, and B. Liskov. BASE: using abstraction to improve fault tolerance. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles*, pages 15–28. ACM Press, Oct. 2001.

[34] A. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. *Lecture Notes in Computer Science*, 2218:329–350, 2001.

[35] A. Rowstron and P. Druschel. Storage management and caching in past, a large-scale, persistent peer-to-peer storage utility. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles*, pages 188–201. ACM Press, 2001.

[36] F. B. Schneider. Implementing fault-tolerant services using the state machine approach: a tutorial. *ACM Computing Surveys*, 22(4):299–319, Sept. 1990.

[37] J. Shneidman and D. Parkes. Rationality and self-interest in peer to peer networks. In *Proc. 2nd IPTPS*, 2003.

[38] J. Shneidman, D. C. Parkes, and L. Massoulie. Faithfulness in internet algorithms. In *Proc. PINS*, Portland, USA, 2004.

[39] E. Tardos. Network games. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 341–342. ACM Press, 2004.

[40] J. Yin, J.-P. Martin, A. Venkataramani, L. Alvisi, and M. Dahlin. Separating agreement from execution for Byzantine fault tolerant services. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles*, pages 253–267. ACM Press, Oct. 2003.