Regret-freedom isn't free

Edmund L. Wong[†,*], Isaac Levy[†], Lorenzo Alvisi[†], Allen Clement[‡], and Mike Dahlin[†]

[†]Department of Computer Science, The University of Texas at Austin
[‡]Max Planck Institute for Software Systems

UTCS Technical Report, TR-11-20

Abstract

Cooperative, peer-to-peer (P2P) services—distributed systems consisting of participants from multiple administrative domains (MAD)—must deal with the threat of arbitrary (Byzantine) failures while incentivizing the cooperation of potentially selfish (rational) nodes that such services rely on to function. Although previous work has generally agreed that these types of participants need to be considered in any formal analysis, there have been differing viewpoints about what the conditions for rational cooperation, in the face of Byzantine failure, need to be. In this paper, we show that regret-freedom, a natural extension of traditional Byzantine fault tolerance that requires optimal choices regardless of how Byzantine failures occur, is unattainable in realistic cooperative services. We argue that protocols should instead aim to be regret-brave: take actions that are in a rational node's best interest based on some prior expectation of Byzantine failures. We demonstrate that by doing so we can provide strong guarantees without sacrificing real-world viability.

1 Introduction

Traditional fault-tolerant distributed computing relies on the assumption that nodes can be cleanly categorized as correct or faulty: the former can be counted on to run protocols that guarantee that systems will continue to provide desirable functionalities despite a limited number of the latter. The rise of cooperative, peer-to-peer (P2P) systems spanning multiple administrative domains (MAD) complicates this simple picture: much evidence suggests that a large number of peers in MAD services will free-ride (*e.g.*, [5, 29, 45]) or deviate from the assigned protocol if it is in their interest to do so (*e.g.*, [1, 45]). To maintain the service, it is essential to give these peers sufficient incentives to cooperate, and informal common-sense reasoning about incentives may still leave systems vulnerable to strategic attacks (*e.g.*, [33, 36, 42]). But what should be the basis for a rigorous treatment of MAD systems?

There is little controversy about the failure model. It is clear that one cannot simply assume that every peer will be rational, as in standard game theory: like other distributed systems, P2P services are susceptible to arbitrary failures. And, of course, some peers may simply be happy to run whatever protocol is assigned to them—similar to correct nodes in traditional distributed systems. P2P participants can therefore be regarded as being Byzantine, acquiescent,[1] or rational (*BAR*), and P2P services should be designed to function in such environments.

Where opinions differ is on what constitutes a best response for rational nodes, given the

[*]Author to whom correspondence should be addressed. Email: elwong@cs.utexas.edu; Address: 1 University Station, #C0500, Austin, TX 78712-0233; Telephone: 512.524.6526

[1]In previous work [6, 16, 34, 35] we called these nodes *altruistic*. We have since been made aware [3] of the risk of confusing such "altruistic" nodes (whose irrational generosity is only driven by obedience to the given protocol) with Byzantine nodes that are irrationally generous for arbitrary reasons. We believe that "acquiescent" better captures our original intentions.

arbitrary nature of Byzantine nodes. The challenge is to specify conditions, *i.e.*, a *solution concept* [22], under which we can assume that rational peers will cooperate and provide some functionality to the service.

A natural approach is to draw inspiration from traditional Byzantine fault tolerant (BFT) computing. In threshold-based BFT, as long as the number of Byzantine nodes does not exceed a threshold t, the system is guaranteed to provide its safety properties *independent of who the t Byzantine nodes are and how they behave*: looking back at an execution one can verify that, try as they may, at no point did the actions of Byzantine nodes undermine safety. Similarly, it is appealing to aim for a notion of equilibrium in which rational nodes maximize their utility *independent of who the t Byzantine nodes are and of how they behave*, *i.e.*, one can verify looking back at an execution (*i.e., ex-post*) that (coalitions of) rational nodes had, at no point, an incentive to deviate. This approach, elegantly formalized in the notion of (k, t)-robustness [2, 4], is in principle very attractive: it promises complete protection from the arbitrary nature of Byzantine failures with the guarantee of *no regrets*: under no circumstances, even after the identities and strategies of the Byzantine nodes become known, will rational players find themselves wishing to have chosen a different strategy.

The main result of this paper is that, despite its appeal, a solution concept that requires no regret from Byzantine actions is fundamentally unable to yield an equilibrium in many practical systems. A real-life example may help provide the intuition. Suppose you are just interested in maximizing your bottom line: should you spend some of your hard-earned income to insure your car against theft? If your car is never stolen, any money spent on insurance would be wasted. Yet, in a no-regret world, your decision should prove a best response *whether or not your car is ever stolen.* In fact, it should even prove a best response if you happen to meet an eccentric millionaire who will give $1,000,000 to any rugged individual who goes through life insurance-free. It is hard to conceive of a strategy that would prove a best response in all three cases.

Fortunately, rational agents that operate under uncertainty are often willing to cooperate without requiring absolute regret-freedom. For instance, when stock traders buy or sell shares, they are well aware of the possibility of regretting their actions. Nonetheless, they participate as long as they can maximize their utility with respect to their expectation of the worth of the traded asset, their comfort with risk, and what they believe will be the trends in the market. We argue that a similar *regret-braving* attitude can be used not just to respond to uncertainty in the environment, but also to develop a rigorous foundation for identifying Byzantine-tolerant equilibria.

To formalize this intuition and analyze rigorously the practical viability of different approaches to Byzantine-tolerant equilibria, we introduce the notion of a *communication game*, which captures the key characteristics of most fault-tolerant distributed systems. Specifically, our game models systems in which (a) to achieve some desired functionality, some nodes need to communicate; (b) bandwidth is not free; and (c) the desired functionality can be achieved despite t Byzantine failures. We ask: what kind of Byzantine-tolerant equilibria are possible in communication games? We find that communication games admit no non-trivial equilibria in solution concepts where rational players need to be completely regret-free with respect to who the t Byzantine nodes are *and* how they behave. Furthermore, even if the solution concept only requires some degree of regret-freedom (*i.e.*, only with respect to who the Byzantine nodes are *or* how they behave), equilibria can only be achieved in communication games under very limited circumstances. These results suggest that regret-free solution concepts such as (k, t)-robustness are unlikely to offer a viable theoretical framework for real-life deployments of BAR-tolerant systems.

Thankfully, we find that non-trivial equilibria in communication games *are* possible for compelling regret-braving solution concepts. Among the many possible, we consider two: in the first, rational nodes fear the worst from their Byzantine counterparts and will play a maximin strategy that guarantees the best worst-case outcome despite any possible Byzantine failure; in the other, rational nodes assign probabilities to various possible faulty behaviors and aim for a Bayesian equilibrium. We show that these regret-braving solution concepts admit simple, intuitive equilibria in a communication game even when the weakened versions of (k, t)-robustness could not.

The paper proceeds as follows. Section 2 formalizes how we model players and introduces the communication game that we use to compare solution concepts. Section 3 explores the land of the (regret-)free, showing how and why the ex-post optimality of (k, t)-robustness prevents equilibria. Section 4 describes instead the home of the (regret-)brave: we discuss two models of rational beliefs that admit useful equilibria in an example of the communication game. Section 5 discusses related work, and Section 6 concludes.

2 Model

A *communication game* is a general model that captures the basic requirements of many cooperative services.

Definition 2.1. *A communication game consists of some set of nodes $N = \{1, \ldots, n\}$ in which:*

- *Communication incurs some cost and does not generate immediate benefit to the sender;*
- *Communication incurs some cost to the receiver;*
- *Benefit is obtained from functionality which (a) can be achieved in the presence of up to $t < n$ Byzantine failures and (b) requires communication between some pair of nodes;*

For simplicity, we use the same communication cost γ for both sending and receiving.

Protocols are strategies played in the communication game, and strategies involve actions drawn from a non-empty, finite set. We refer to the service-assigned protocol as the *assigned strategy.* A *strategy profile* $\sigma = \{\sigma_x\}_{x \in N}$ is a complete set of strategies, one for each node, and Σ denotes the space of all possible strategy profiles σ that nodes may use. Every strategy profile σ results in some utility $U_x(\sigma)$ for every node x. We are only interested in *non-trivial* strategy profiles, in which some positive utility is expected for at least one node; this implies that some communication must occur. Finally, we refer to the actions that a node has performed or observed in the past as the node's *history.*

We are interested in systems that include Byzantine, rational, and (optionally) acquiescent nodes: in general, each node x belongs to a type *type* θ_x that falls into one of these groups. For simplicity, we assume that all rational nodes are of the same type R, and we ignore acquiescent nodes (who would anyway follow any strategy assigned to them). Because a Byzantine node may potentially play one of many different strategies, it is convenient to denote the node's type using the strategy it plays. Formally, if some Byzantine node z plays some strategy τ_z, then we say that $\theta_z = \tau_z$; the type space Θ then consists of $\Sigma \cup \{R\}$.

We focus on environments in which there is no trusted hardware nor trusted third-parties that monitor communication. Although a trusted mediator is useful [10, 28, 48], it is often impractical or even infeasible to provide one, and in practice few cooperative systems leverage trusted hardware to prove communication. We express this reality in the following assumption:

Assumption 2.2. *A node that sent a message m cannot unilaterally prove that it sent m.*

3 Byzantine regret-freedom in communication games

In BFT systems, safety properties hold regardless of how Byzantine failures occur. Ideally, one would like to achieve similar guarantees when it comes to ensuring rational cooperation. (k, t)-robustness [2, 4] is an elegant solution concept that captures this attractive intuition. A (k, t)-robust equilibrium is completely impervious to the actions of Byzantine nodes: rational nodes will never have to second-guess their decision even if the identities and strategies of the Byzantine nodes become known. Specifically, (k,t)-robustness offers two key properties. The first, *t-immunity* [2], captures the intuition that nodes following a strategy profile designed to handle up to t Byzantine failures should not be negatively affected by them.

DEFINITION 3.1. *A strategy profile σ is t-immune if, for all $T \subseteq N$ such that $|T| \leq t$, all strategy profiles τ, and (non-Byzantine) nodes $x \notin T$,*

$$U_x(\sigma_{-T}, \tau_T) \geq U_x(\sigma)$$

The second, *k-resilience* [2], addresses the possibility of collusion: a k-resilient strategy guarantees that a coalition of at most size k cannot deviate in a way that benefits every member.[2]

DEFINITION 3.2. *A strategy profile σ^* is k-resilient if, for all $K \subseteq N$ such that $|K| \leq k$, there exists no alternate strategy profile σ' such that for all $x \in K$,*

$$U_x(\sigma'_K, \sigma^*_{-K}) > U_x(\sigma^*)$$

The (k,t)-robustness solution concept is the combination of t-immunity, k-resilience, *and regret-freedom with respect to Byzantine failure*: regardless of how Byzantine failures occurs, (k,t)-robustness guarantees that no coalition of at most k nodes can ever do better than following the equilibrium strategy.

DEFINITION 3.3. *A strategy profile σ^* is a (k,t)-*robust equilibrium *if σ^* is t-immune and, for all (a) $K, T \subseteq N$ such that $K \cap T = \emptyset$, $|K| \leq k$, and $|T| \leq t$, and (b) strategy profiles τ, there does not exist an alternate strategy profile σ' such that for all $x \in K$,*

$$U_x(\sigma'_K, \tau_T, \sigma^*_{-\{K \cup T\}}) > U_x(\sigma^*_{-T}, \tau_T)$$

3.1 (k,t)-robustness is infeasible in communication games

We show that the very property—regret-freedom regardless of how Byzantine failures occur—that makes (k,t)-robustness so appealing, makes it infeasible in many real-world systems. The reason, fundamentally, is that if interacting with other nodes incurs a cost, then under some Byzantine strategies a rational node may realize in hindsight, independent of its chosen strategy, that it could have done better by avoiding all communication with nodes that are later found to be Byzantine.

THEOREM 3.4. *There exist no non-trivial (k,t)-robust equilibria in any communication game.*

Proof. Consider some non-trivial (k,t)-robust strategy σ^*. There must exist some node x which, with positive probability α under σ^*, sends a message to some other node z before receiving any other messages. Suppose that z is Byzantine. Since σ^* is (k,t)-robust, x must not be able to do better with some alternate strategy, regardless of who has failed and what a failed node will do. In particular, for all alternate strategies σ'_x for x and Byzantine strategies τ_z for z, it must be that

$$U_x(\sigma^*_{-x}, \tau_x) \geq U_x(\sigma^*_{-\{x,z\}}, \sigma'_x, \tau_z) \tag{1}$$

Suppose τ_z is the strategy in which z "crashes" immediately, *i.e.*, z never sends any messages. Let σ'_x be the strategy in which x does everything in σ^*_x, except x sends nothing to z. By Assumption

[2] Abraham et al. also define a strong version of collusion resilience in which there must not exist a deviation in which even *one* coalition member can do better [2, 4]. We focus on the weak version as Abraham et al. do in [4]. Since any strongly k-resilient equilibria is (weakly) k-resilient, our impossibility results hold in both versions.

2.2, x cannot prove that it communicated with z; it thus follows that $(\sigma^*_{-\{x,z\}}, \sigma'_x, \tau_z)$ has the same functionality as (σ^*_{-z}, τ_z) and is indistinguishable to any node in $N \setminus \{x, z\}$. Clearly, if z follows τ_z, x can do better by never communicating with z; x's outcome will not change (since z never communicates with anyone), and x's communication costs are lower. Formally,

$$U_x(\sigma^*_{-\{x,z\}}, \sigma'_x, \tau_z) = U_x(\sigma^*_{-z}, \tau_z) + \alpha\gamma > U_x(\sigma^*_{-z}, \tau_z)$$

which directly contradicts inequality (1). □

More broadly, Theorem 3.4 suggests that it may be hard to build (k, t)-robust equilibria for any game where a player's actions incur cost. Indeed, in all the games for which Abraham et al. derive (k, t)-robust equilibria [2, 4], a node's utility depends only on the game's outcome (e.g., in a secret sharing game based on Shamir's scheme, on whether a node can learn the secret) and is independent of how much communication is required to reach that outcome.

Discussion. (k, t)-robustness promises regret-freedom simultaneously along two axes: *who* the Byzantine nodes are and *how* they behave. Theorem 3.4 suggests that this may be too strong to require in practice. But what if we only require regret-freedom along only one axis? If we know exactly who the Byzantine nodes are, but not how they will behave, can we achieve regret-freedom in communication games? What if we do not know who is Byzantine, but we know their strategy?

3.2 What if we know who is Byzantine?

Let us assume that we know *exactly* who all the Byzantine players are before the game begins. This may already appear a strong assumption, but it is necessary, since if the identity of even one Byzantine node were unknown, Theorem 3.4 would still apply. We show that, even with this strong assumption, a solution concept that is regret-free with respect to the strategies of Byzantine nodes is possible only to the extent that it defines away the problem: the only possible equilibria are those in which rational nodes communicate only among themselves, completely excluding Byzantine nodes from the system. Furthermore, we show that many interesting communication games do not yield a regret-free equilibrium even if one takes the drastic step of excluding Byzantine nodes: specifically, communication games in which Byzantine nodes may take actions that can *improve* a rational node's utility by more than the cost of sending a *single* message have no regret-free equilibrium strategy, even if the identity of all Byzantine nodes are known a priori.

Definition 3.5. *A strategy σ^* is (k, T)-strategy-robust with respect to $T \subseteq N$ iff σ^* is $|T|$-immune and for all $K \subseteq N \setminus T$ such that $|K| \leq k$ and all Byzantine strategies τ, there does not exist some σ' such that for all $i \in K$,*

$$U_i(\sigma'_K, \tau_T, \sigma^*_{-(K \cup T)}) > U_i(\sigma^*_{-T}, \tau_T)$$

A (k, T)-*strategy*-robust equilibrium need only be a best response to the specified set T of Byzantine nodes. The following theorem shows that no (k, T)-*strategy*-robust equilibrium is possible unless rational nodes "blacklist" all nodes in T.

Theorem 3.6. *In a communication game, there does not exist any (k, T)-strategy-robust equilibrium σ^* where a rational node communicates with any node $z \in T$.*

Proof. This is similar to the proof of Theorem 3.4. Suppose the strategy τ_z that some Byzantine node z employs is the crash strategy: it never communicates.

Consider an alternate strategy strategy σ'_x in which some rational node x does everything in σ^*_x except communicate with z. Since σ^* is a (k, T)-*strategy*-robust equilibrium, $U_x(\sigma^*_{-z}, \tau_z) \geq U_x(\sigma'_x, \tau_z, \sigma^*_{-\{x,z\}})$. Yet, by Assumption 2.2, x cannot prove it communicated with z. It follows that $(\sigma^*_{-\{x,z\}}, \sigma'_x, \tau_z)$ has the same functionality as (σ^*_{-z}, τ_z) and is indistinguishable to any node in $N \setminus \{x,z\}$. Thus,

$$U_x(\sigma^*_{-\{x,z\}}, \sigma'_x, \tau_z) \geq U_x(\sigma^*_{-z}, \tau_z) + \alpha\gamma > U_x(\sigma^*_{-z}, \tau_z)$$

Contradiction. □

Although Theorem 3.6 does not rule out all (k,T)-*strategy*-robust equilibria, Theorem 3.7 proves that these equilibria, which must be regret-free for *any* Byzantine strategy, only exist in limited circumstances.

THEOREM 3.7. *No communication game can yield a (k,T)-*strategy*-robust equilibrium for any set $T \subseteq N$ of Byzantine nodes if (a) some node $x \in N \setminus T$ has at least one opportunity to send a message to a node in T and (b) there exists some Byzantine strategy in which x increases its payoff by more than γ after its first opportunity to communicate with some node in T.*

Proof. By contradiction. Let σ^* be some (k,T)-*strategy*-robust equilibrium. We know by Theorem 3.6 that if σ^* is (k,T)-*strategy*-robust, then any rational node x following σ^* never chooses to send to any member of T. We now construct a Byzantine strategy in which it may be in x's best interest to communicate with T. Let $\bar{\tau}_T$ be the Byzantine strategy which maximizes x's utility if all nodes in $N \setminus T$ play σ^*_{-T} after x's first action with respect to some $z \in T$; define every action prior to x's action arbitrarily. Similarly, let $\underline{\tau}_T$ be the Byzantine strategy which minimizes x's utility if all nodes in $N \setminus T$ play σ^*_{-T} after x's first action with respect to some $z \in T$; define every action prior to x's action the same as in $\bar{\tau}_T$.

By assumption, we know that T can affect x's utility by more than γ and thus $U_x(\sigma^*_{-T}, \bar{\tau}_T) - U_x(\sigma^*_{-T}, \underline{\tau}_T) > \gamma$. Let τ_T be the Byzantine strategy in which T plays $\bar{\tau}_T$ if x's first action with respect to any member of T is to communicate something; otherwise, T plays $\underline{\tau}_T$. Consider some alternate strategy σ'_x in which x plays as in σ^*_x, except x chooses to communicate in its first interaction with any member of T; x then plays the same strategy as in σ^*_x. Since σ^* and σ' are indistinguishable and equivalent to anyone outside of T and x, it follows that

$$U_x(\sigma'_x, \tau_T, \sigma^*_{-(\{x\}\cup T)}) = U_x(\sigma'_x, \bar{\tau}_T, \sigma^*_{-(\{x\}\cup T)}) > U_x(\sigma^*_{-T}, \bar{\tau}_T) - \gamma > U_x(\sigma^*_{-T}, \underline{\tau}_T) = U_x(\sigma^*_{-T}, \tau_T)$$

This contradicts the assumption that σ^* is a (k,T)-*strategy*-robust equilibrium. □

3.3 What if we know how Byzantine nodes behave?

Let us now consider a solution concept that assumes that the strategy played by every Byzantine node is known a priori and yields equilibria that are regret-free with respect to who the Byzantine nodes are.

DEFINITION 3.8. *The strategy profile σ^* is a (k,t,τ)-type-robust equilibrium iff σ^* is t-immune and for all $K, T \subseteq N$ such that $K \cap T = \emptyset$, $|K| \leq k$, and $|T| \leq t$, there does not exist some σ' such that for all $x \in K$,*

$$U_x(\sigma'_K, \tau_T, \sigma^*_{-(K\cup T)}) > U_x(\sigma^*_{-T}, \tau_T)$$

Despite the strong assumption on which they rely, (k,t,τ)-*type*-robust equilibria are impossible to achieve for many Byzantine behaviors. In particular, it follows immediately from Theorem 3.4 that no such equilibrium is possible if the known Byzantine strategy calls for any Byzantine node to crash at the very beginning of the game.

THEOREM 3.9. *There exist no non-trivial (k,t,τ)-type-robust equilibria in the communication game in which a Byzantine node z, following τ_z, crashes at the beginning of the game.*

Proof. Same as proof of Theorem 3.4. □

In general, it is possible to show that (k,t,τ)-*type*-robust equilibria are impossible whenever there is a point in the known Byzantine strategy after which a Byzantine node becomes "unresponsive": the node's behavior becomes independent of how the game has been played so far (*e.g.*, the node crashes or starts flooding all other nodes with messages). To simplify exposition, we assume that every node has some notion of time: at the very least, it must know, whenever it either has a chance to send or receive a message, whether this is before or after the point in which a Byzantine node becomes unresponsive.

THEOREM 3.10. *There exist no non-trivial (k,t,τ)-type-robust equilibria in the communication game in which z, following τ_z (a) before some commonly-known time r, plays as if it were playing σ^*_z, and (b) starting from time r, is unresponsive.*

Proof. Similar to proof of Theorem 3.4; by contradiction. Assume that σ^* is a non-trivial (k,t,τ)-type-robust equilibrium as described in the theorem statement. Consider the first rational node x who, with positive probability under σ^*, sends a message after time r to some peer z at some time $r_x \geq r$. Before time r, Byzantine and rational nodes play the same actions. Between time r and r_x, the same is still true: Byzantine nodes are unresponsive, and by assumption, x is the first rational node to send a message. As the equilibrium strategy cannot depend on who the Byzantine nodes are, x's choice to send a message must not depend on whether z is Byzantine. However, if z turns out to be unresponsive, x is clearly better off not sending to z (using the same argument as in the proof of Theorem 3.4).

This implies that under σ^*, no node sends any messages starting from time r. Now consider the last rational node y who, with positive probability α_y under σ^*, sends a message before time r. Let σ' be some alternate strategy in which x does not send anything at time r_x, and let V^* and V' be the expected utilities of playing σ^* and σ'. It follows that $V' = V^* + \alpha_y\gamma > V^*$, and thus x is better off playing σ'. Contradiction. □

4 Dealing with Byzantine failures through regret-bravery

The previous section suggests that solution concepts that require regret-freedom are unlikely to provide a viable theoretical framework for many BAR-tolerant systems. Yet, such systems do exist: previous work [6, 15, 34, 35, 39, 52] has circumvented these limitations by instead reasoning about rational adherence to a protocol given expectations regarding Byzantine failures.

In this section, we formalize this notion. We emphasize that *regret-braving* solution concepts do not restrict how Byzantine failures occur; rather, they explicitly specify how rational nodes calculate the expected utility of various strategies. After all, although Byzantine nodes can behave in any way, it does not follow that it is necessarily rational to aim to best respond to all possible Byzantine behaviors. Regret-braving solution concepts therefore provide a means to consider less demanding conditions under which rational nodes follow the assigned strategy.

There is a rich design space of regret-braving solution concepts; which one to use depends on the application. We provide two concrete examples. In the first, rational nodes best-respond to fearing the worst, *i.e.*, they follow a maximin strategy with respect to Byzantine failures.

DEFINITION 4.1. *The strategy profile σ^* is a k-resilient t-maximin equilibrium iff for any coalition $K \subseteq N$ such that $|K| \leq k$, there does not exist an alternate strategy profile σ' such that for all nodes $x \in K$,*

$$\min_{\substack{T\subseteq N\setminus K:\\ |T|\leq t}} \min_{\tau} U_x(\sigma'_K, \tau_T, \sigma^*_{-(K\cup T)}) \geq \min_{\substack{T\subseteq N\setminus K:\\ |T|\leq t}} \min_{\tau} U_x(\sigma^*_{-T}, \tau_T)$$

and for some $y \in K$, the inequality is strict.

In the second, rational nodes weigh the probabilities of various Byzantine failures; an equilibrium is thus these probabilities—known as *beliefs* in game theory parlance—and the strategy profile that is an expected best response given these beliefs. A set of beliefs $\mu = \{\mu_x\}_{x \in N}$ is, for each node, a probability distribution over sets of nodes and their types—whether they are rational, or Byzantine and playing a particular strategy. For notational simplicity, we define $\hat{\mu}_x(\tau_Z)$ to be x's belief that, given that x is of type R, all nodes $z \in Z$ are Byzantine and of type (*i.e.*, playing strategy) τ_z and, for all $y \notin Z$, y is of type R.

Definition 4.2. *The strategy profile/belief tuple (σ^*, μ^*) is a k-resilient Bayes equilibrium iff for all coalitions $K \subseteq N$ such that $|K| \leq k$, there does not exist an alternate strategy profile σ' such that for all $x \in K$,*

$$\sum_{T \subseteq N \setminus K} \sum_{\tau} \hat{\mu}^*_x(\tau_T) U_x(\sigma^*_{-T}, \tau_T) \geq \sum_{T \subseteq N \setminus K} \sum_{\tau} \hat{\mu}^*_x(\tau_T) U_x(\sigma^*_{-(K \cup T)}, \sigma'_K, \tau_T)$$

and for some $y \in K$, the inequality is strict.

In both definitions, we extend previous work that uses regret-brave solution concepts [6, 34, 35, 52] by explicitly considering collusion, which prior work avoided by either considering collusion a Byzantine failure or making informal arguments on the basis of experimental results. For simplicity, we use k-resilience (Definition 3.2); however, we could have used any notion of collusion resilience.

Finally, to show the viability of regret-brave solution concepts in a communication game, we consider a concrete communication game: a *quorum game*, which models protocols, such as secret-sharing (*e.g.*, [2, 14, 41, 43, 44, 48]), replicated state machines (*e.g.*,[13, 31, 47]), and terminating reliable broadcast (*e.g.*, [12, 32, 26, 49]), in which functionality is achieved iff some subset of nodes (a *quorum*) work together.

Definition 4.3. *A (synchronous) quorum game is an infinitely-repeated communication game where:*

- *There are at least 3 nodes ($n \geq 3$), of which at most t are Byzantine, where $t \leq n - 2$.*
- *The game repeats indefinitely. In every round, for each peer y, a rational node x decides whether to send a message ("contribute") or not ("snub") to y.*
- *At the end of the round, every node x simultaneously (1) observes who contributes to it and (2) receives its payoff.[3] x incurs a cost of γ for each node x contributes to and for each node that contributes to x; x incurs no cost for snubbing or being snubbed. x realizes a positive benefit of $b > 2n\gamma$ in any round where q other nodes (a quorum) contribute to x, where $t \leq q < n - t$.[4]*
- *The total payoff is the δ-discounted sum of each individual round's payoff, where $0 < \delta < 1$.*

Examples of t-maximin equilibria. We first show a 1-resilient t-maximin equilibrium in the quorum game.

Theorem 4.4. *Let the strategy profile σ^* be defined as follows: any node x following σ^*_x contributes to some node y iff y has never snubbed x in the past and x has been snubbed by at most t different*

[3]In game theory parlance, the game is a simultaneous game; in distributed systems, synchronous.

[4]Technically, the quorum size is $q + 1$: q other nodes and the node itself (we assume that it costs nothing for a node to contribute to itself). For simplicity, we will simply say that the quorum size is q.

nodes. Then σ^* is a 1-resilient t-maximin equilibrium in the quorum game if

$$\frac{b}{\gamma} \geq \frac{2(n-1) - (1-\delta^2)(q+t)}{\delta^2} \tag{2}$$

Proof. If all rational nodes follow σ^*, then each node will achieve quorum regardless of Byzantine behavior. x's utility is thus at least $\big(b - 2(n-1)\gamma\big)/(1-\delta)$.

Suppose that in round r, x snubs some node y that x was supposed to contribute to after history h_x^r. If h_x^r is not "expected" to occur (*i.e.*, given σ^* and the T and τ that minimize σ^*, h_x^r occurs with zero probability[5]) then the expected utility of deviating is at most that of σ^*, and the proof is trivially complete. Otherwise, suppose h_x^r is expected to occur. If y is rational and t Byzantine nodes, in addition to x, snub y in round r, y will snub all nodes from round $r+1$ at latest. If y and all t Byzantine nodes snub everyone in round $r+1$, then all other nodes snub everyone from round $r+2$ at latest. Therefore, x earns at most $b-(q+t)\gamma$ in rounds r and $r+1$ and 0 in subsequent rounds for a total of $(1+\delta)(b-(q+t)\gamma)$. Deviating is thus never worthwhile given inequality (2).

Now consider if x contributes to some node y that x, under σ^*, was supposed to snub after some history h_x^r. Again, if this is not expected to occur, the proof is trivially complete. Otherwise, since rational nodes never snub unless they were snubbed, x is contributing to a Byzantine node y, and it is obvious that this is never in x's best interest. □

We now prove a k-resilient t-maximin equilibrium for $k > 1$ in the quorum game. Although we argue that communication always has cost and the quorum game does not explicitly model communication that coalition members may perform to coordinate, our proof implicitly assumes that the coalition can coordinate its strategies. Thus, our results hold even if we augmented the game to allow coalition members to coordinate their strategies via cheap talk [17, 21].

Corollary 4.5. *The strategy profile σ^* as defined in Theorem 4.4 is a k-resilient t-maximin equilibrium if condition* (2) *holds, $k < q$, $q = n - t - 1$, and*

$$b \geq \left(\frac{1}{1-\delta} t + k + 1\right)\gamma \tag{3}$$

Proof. Since $q = n - t - 1$, a rational node needs the cooperation of all other rational nodes to achieve a quorum. As in the proof of Theorem 4.4, a rational node never snubs a node outside of the coalition. However, coalition members have an additional possible deviation: they may choose to help each other save on receiving extraneous contributions (stemming from the fault-tolerant nature of the quorum game, nodes typically send and receive contributions from more than q members) by "snubbing" one another without threat of punishment.

Consider some coalition K of size k that plays some alternate strategy σ'_K in which some $s \leq t$ rational nodes in K snub some $x \in K$. Then it is possible that $t - s + 1$ Byzantine nodes also snub x, causing x to lose the benefit it would have normally gained from playing σ^*, while saving at most $k\gamma$ from not having to contribute to coalition members and $(t+1)\gamma$ from not receiving contributions from $t+1$ other nodes. Furthermore, x now snubs $t-s+1$ nodes in the future, saving $(t-s+1)\gamma$ per round (in the worst case, they still continue to contribute to x). In order for σ'_K to be worthwhile, it must be the case that

$$\frac{\delta}{1-\delta}(t-s+1)\gamma - b + (k+t+1)\gamma > 0$$

This is never satisfied given inequality (3). □

[5] This is similar to being "off the equilibrium path" in traditional game theory.

Examples of Bayesian equilibria. One advantage of using the t-maximin solution concept is its simplicity: because we need only consider the worst possible case, t-maximin equilibria are simple to analyze. Unfortunately, although a rational node playing a t-maximin equilibrium may receive a safe, steady amount of utility, Byzantine failures are unlikely to always occur in the worst possible way, and a rational node willing to take a risk is likely to do better in expectation.

In the remainder of this section, we demonstrate that the Bayesian approach provides flexibility in how Byzantine nodes are modeled by rational nodes by showing equilibria given two different sets of beliefs. Our goal in these examples is to simply illustrate the existence of Bayesian equilibria, not to derive tight bounds for when these equilibria exist. Thus, for simplicity of exposition, we will be extremely optimistic about the utility earned by deviating and pessimistic about the utility earned by cooperating.

We first show that 1-resilient Bayes equilibria exist in a simple scenario that models the one used in the proof of Theorem 3.7: Byzantine nodes are likely to either crash or threaten to inflict communication costs unless rational nodes contribute.

THEOREM 4.6. *Define the strategy profile σ^* as follows. For any node, $x \in N$, let T_x^i be the set of nodes who have snubbed x in round i. x, following σ_x^*, does the following: (a) in round 0, x contributes to all nodes; (b) in round $r > 0$, if $|T_x^0| > t$ or there exists some round $i < r$ with $T_x^i \not\subseteq T_x^0$, then snub all nodes; otherwise, contribute to all nodes in $N \setminus T_x^0$.*

Let μ^* be some set of beliefs which place positive probability only on the following Byzantine strategies: (a) snub everyone (the crash strategy); and (b) snub everyone in the first round, and, in any subsequent round r, snub a node y iff y previously contributed to it.

Let ψ be the joint probability (based on μ^*) that the environment has exactly t Byzantine nodes and that a node, picking a peer at random, selects a rational one. Then (σ^*, μ^*) is a 1-resilient Bayes equilibrium in the quorum game if $\psi > 0$ and

$$\frac{b}{\gamma} \geq \frac{1}{\delta^2 \psi}(2(n-1) - q(1-\delta^2\psi)) \tag{4}$$

Proof. Consider some rational node x. If x follows σ^*, the worst x can do is achieve quorum in every round with everyone communicating with it, earning a total of

$$V^* = \frac{1}{1-\delta}(b - 2(n-1)\gamma)$$

x may deviate in the following ways:

- x snubs a node in round 0;
- x snubs a node that x is supposed to contribute to in round $r > 0$;
- x contributes to a node that x is supposed to snub in round $r > 0$ (because more than t nodes have snubbed x and/or that node and x have not been fully cooperative in the past).

We consider the first case: suppose x snubs some set of nodes $L \subset N$. We optimistically assume that (a) x is only hurt if there exists some rational node $y \in L$ that x snubbed and there are exactly t Byzantine nodes, which occurs with probability at least ψ, and (b) x earns the maximum round payoff that it can (*i.e.*, $b - q\gamma$ if it can achieve quorum and 0 if not) when deviating. Thus, if L contains only Byzantine nodes, then x earns $b - q\gamma$ in all rounds. Otherwise, x earns at most $b - q\gamma$ in rounds 0 and 1 and 0 in all subsequent rounds. x's total expected payoff from deviating is then

$$V' = \psi(1+\delta)(b - q\gamma) + (1-\psi)\frac{1}{1-\delta}(b - q\gamma)$$

It follows, from inequality (4), that $V^* \geq V'$, and thus x does not deviate.

Now consider any subsequent round $r > 0$. At this point, x knows exactly who the Byzantine nodes are. Thus, the second case, in which a rational node that deviates by snubbing some node y that has never snubbed it before, results in y snubbing everyone in round $r+1$, causing all nodes to snub x by round $r+2$. In this case, x earns at most $(1+\delta)(b-q\gamma)$ in rounds r and $r+1$ and 0 in all subsequent rounds. It is obvious that the argument used for round 0 holds here as well.

Finally, consider the third case: suppose x has followed σ^* until round r and contributes to a node z that x is supposed to snub in round r. Regardless of the reason, contributing to z does not affect z's strategy in any way and is thus never in x's best interest (see Theorem 3.10). $\square$

Finally, we prove a k-resilient Bayes equilibrium for $k > 1$ for a more realistic scenario.

THEOREM 4.7. *Define the strategy profile σ^* such that any node $x \in N$, following σ^*_x, contributes to any node y iff x and y have always contributed to each other in the past and x has been snubbed by at most t peers.*

Let τ be defined as the random crash strategy: in any given round, a node z playing τ_z has some positive probability ρ of crashing. Define the set of beliefs μ^* such that for any subset $K \subset N$ such that $|K| \leq k$, (a) $\hat{\mu}^*_K(\tau_T) = 0$ for any T such that $|T| \neq t$, and (b) $\hat{\mu}^*_K(\tau_{T_1}) = \hat{\mu}^*_K(\tau_{T_2}) > 0$ for any $T_1, T_2 \subset N$ such that $|T_1| = |T_2| = t$, $T_1 \cap K = \emptyset$, and $T_2 \cap K = \emptyset$.

Then (σ^*, μ^*) is a k-resilient Bayes equilibrium if

$$\frac{b}{\gamma} \geq \frac{n+t-1}{\rho^t \delta^2 (1-\delta)} \frac{n-k}{n-k-t} + n - t - 1 \tag{5}$$

Proof. Fix some rational node x and some coalition K, where $x \in K$ and $|K| \leq k$. We optimistically assume a rational node that deviates in round r only loses utility if t nodes crash on or before round r, which occurs with probability at least ρ^t.

As before, the minimum that x earns following σ^* is $(b - 2(n-1)\gamma)/(1-\delta)$. Suppose that x snubs some node $y \notin K$. Since the probability that a node is rational is uniform across all nodes, y is rational with probability at least $1 - t/(n-k)$, and with probability at least ρ^t, y will observe t other nodes snub it in round r. y then snubs everyone starting in round $r+1$, all non-coalition nodes snub everyone starting in round $r+2$, and x earns at most 0 in every round starting from round $r+2$. Otherwise, we assume x earns the maximum possible round payoff $b - q\gamma$. Thus, deviating is worthwhile only if

$$\rho^t \left(1 - \frac{t}{n-k}\right)(1+\delta)(b-q\gamma) + \left(1 - \rho^t\left(1 - \frac{t}{n-k}\right)\right)\frac{1}{1-\delta}(b-q\gamma) > \frac{1}{1-\delta}(b-2(n-1)\gamma)$$

This never holds given inequality (5).

Otherwise, suppose that $x \in K$ "snubs" its peer $y \in K$ to save on y's communication costs. Again, y, with probability at least ρ^t, will not achieve quorum if all t nodes crash on or before round r. However, unlike before, y only loses quorum for one round; we otherwise assume that it achieves the maximum round payoff $b - q\gamma$. Thus, deviating as a coalition is worthwhile only if

$$\rho^t \frac{\delta}{1-\delta}(b-q\gamma) + (1-\rho^t)\frac{1}{1-\delta}(b-q\gamma) > \frac{1}{1-\delta}(b-2(n-1)\gamma)$$

This never holds given inequality (5). $\square$

5 Related work

Outside of (k,t)-robustness [2, 4], Eliaz [20] also defined a solution concept which is effectively $(1,t)$-robustness. Gradwohl [24] explored regret-free equilibria with t arbitrary or colluding nodes in leader election and random sampling games. Our results still apply to the solution concepts used in these papers. Moscibroda et al. [39] use an approach similar to t-maximin to consider worst-case Byzantine behavior in the context of a computer virus propagation model.

Coalitions have been studied in depth in the game theory literature. Aumann [51] proposed a notion of collusion resilience which is the basis for k-resilience. Berheim et al. [11], Moreno et al. [38], Einy et al. [19], among others, have proposed weaker solution concepts which only consider deviations that are self-enforcing, in that there does not exist an even more profitable deviation for a sub-coalition of the coalition. All of these notions are complementary to regret-brave equilibria and can be used as a part of a regret-brave solution concept.

Our results are similar in spirit to previous work in mechanism design, where previous work [18, 23, 25, 30, 40, 46] has found that mechanisms that incentivize nodes to reveal their true preferences or types for every possible realization of their peers' types are often impossible or heavily restricted. Others [18, 40] found positive results by using Bayesian solution concepts instead of dominant ones. Mookherjee et al. [37] define when Bayesian incentive-compatible mechanisms can be replaced by equivalent dominant-strategy mechanisms.

Maximin strategies have been previously explored in conjunction with adversarial or possibly irrational agents. Alon et al. [7] quantify how, in a two-player zero-sum game, the payoff of playing a mixed maximin strategy is affected by an adversary who can choose its actions based on some information about its peer's realized strategy. Tennenholtz [50], extending the work of Aumann et al. [8, 9], explores how maximin strategies can approximate the payoff of a Nash equilibrium when a rational node may not want to rely on the rationality of its peers.

6 Conclusion

Distributed systems that span multiple administrative domains must tolerate the possibility that nodes may be Byzantine, rational, and (possibly) acquiescent. To formally reason about such services, we need a solution concept that rigorously guarantees rational cooperation without sacrificing real-world applicability. This paper argues that solution concepts based on regret-freedom, despite their intuitive correspondence to the traditional guarantees of fault-tolerant distributed computing, are unlikely to provide the basis for a viable theoretical framework for real-world systems. In particular, we believe that any practical solution concept should base a rational node's payoff not just on the game's outcome, but also on the cost of the actions required by a strategy. While our discussion here has focused on communication costs, other costs should be included, such as the computational costs discussed in the recent work of Halpern and Pass [27]. We believe that regret-brave solution concepts provide a framework for accounting for these costs that is both rigorous and realistic.

References

[1] Kazaa Lite. `http://en.wikipedia.org/wiki/Kazaa_Lite`.

[2] I. Abraham, D. Dolev, R. Gonen, and J. Halpern. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *PODC '06*, pages 53–62, July 2006.

[3] I. Abraham, D. Dolev, and J. Y. Halpern. Private communication.

[4] I. Abraham, D. Dolev, and J. Y. Halpern. Lower bounds on implementing robust and resilient mediators. In *Proceedings of the 5th Conference on Theory of Cryptography*, TCC'08, pages 302–319, Berlin, Heidelberg, 2008. Springer-Verlag.

[5] E. Adar and B. A. Huberman. Free riding on Gnutella. *First Monday*, 5(10):2–13, Oct. 2000.

[6] A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J.-P. Martin, and C. Porth. BAR fault tolerance for cooperative services. In *SOSP '05*, pages 45–58, Oct. 2005.

[7] N. Alon, Y. Emek, M. Feldman, and M. Tennenholtz. Adversarial leakage in games. In *Proceedings of Innovations in Computer Science (ICS 2010)*, pages 111–119.

[8] R. J. Aumann and M. Maschler. Some thoughts on the minimax principle. *Management Science*, 18(5):P54–P63, 1972.

[9] R. J. Aumman. On the non-transferable utility value: A comment on the Roth-Shaper examples. *Econometrica*, 53(3):667–677, 1985.

[10] E. Ben-Porath. Cheap talk in games with incomplete information. *Journal of Economic Theory*, 108:2003, 2003.

[11] B. D. Bernheim, B. Peleg, and M. D. Whinston. Coalition-proof nash equilibria i. concepts. *Journal of Economic Theory*, 42(1):1–12, 1987.

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

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

[14] B. Chor, S. Goldwasser, S. Micali, and B. Awerbuch. verifiable secret sharing and achieving simultaneity in the presence of faults. In *FOCS*, pages 383–395, 1985.

[15] A. Clement, H. Li, J. Napper, J.-P. Martin, L. Alvisi, and M. Dahlin. BAR primer. In *38th Annual IEEE/IFIP International Conference on Dependable Systms and Networks (DSN)*, June 2008.

[16] A. Clement, J. Napper, E. Wong, H. Li, J. Martin, L. Alvisi, and M. Dahlin. Theory of BAR games. Technical report, University of Texas Department of Computer Sciences Technical Report R-06-63, December 2005.

[17] V. P. Crawford and J. Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, Nov. 1982.

[18] C. d'Aspremont and L.-A. Gerard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11(1):25–45, Feb. 1979.

[19] E. Einy and B. Peleg. Coalition-proof communication equilibria. In W. A. Burnett, H. Moulin, M. Sales, and N. J. Schofield, editors, *Social Choice, Welfare, and Ethics: Proceedings of the Eighth International Symposium in Economic Theory and Econometrics*, pages 289–300, Cambridge, 1995. Cambridge University Press.

[20] K. Eliaz. Fault tolerant implementation. *Rev. of Econ. Studies*, 69:589–610, Aug 2002.

[21] J. Farrell and M. Rabin. Cheap talk. *The Journal of Economic Perspectives*, 10(3):103–118, Summer 1996.

[22] D. Fudenberg and J. Tirole. *Game Theory.* MIT Press, Aug. 1991.

[23] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, July 1973.

[24] R. Gradwohl. Rationality in the full-information model. In *Theory of Cryptography*, volume 5978 of *Lecture Notes in Computer Science*, pages 401–418. Springer Berlin / Heidelberg, 2010.

[25] J. Green and J.-J. Laffont. On coalition incentive compatibility. *The Review of Economic Studies*, 46(2):243–254, Apr. 1979.

[26] V. Hadzilacos and S. Toueg. Fault-tolerant broadcasts and related problems. *Distributed Systems (2nd Ed.)*, pages 97–145, 1993.

[27] J. Halpern and R. Pass. Game theory with costly computation. 2010.

[28] J. Halpern and V. Teague. Rational secret sharing and multiparty computation. In *Proc. 36th STOC*, pages 623–632, 2004.

[29] D. Hughes, G. Coulson, and J. Walkerdine. Free riding on Gnutella revisited: the bell tolls? *IEEE Distributed Systems Online*, 6(6), June 2005.

[30] P. Jehiel, M. Meyer-ter Vehn, B. Moldovanu, and W. R. Zame. The limits of ex post implementation. *Econometrica*, 74(3):585–610, 2006.

[31] L. Lamport. The part-time parliament. *ACM Transactions on Computer Systems*, 16(2):133–169, 1998.

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

[33] D. Levin, K. LaCurts, N. Spring, and B. Bhattacharjee. BitTorrent is an auction: analyzing and improving BitTorrent's incentives. *SIGCOMM Comput. Commun. Rev.*, 38(4):243–254, 2008.

[34] H. Li, A. Clement, M. Marchetti, M. Kapritsos, L. Robinson, L. Alvisi, and M. Dahlin. FlightPath: Obedience vs choice in cooperative services. In *OSDI '08*, pages 355–368, Dec 2008.

[35] H. C. Li, A. Clement, E. Wong, J. Napper, I. Roy, L. Alvisi, and M. Dahlin. BAR Gossip. In *OSDI '06*, pages 191–204, Nov. 2006.

[36] T. Locher, P. Moor, S. Schmid, and R. Wattenhofer. Free riding in bittorrent is cheap. In *HotNets '06*, Nov. 2006.

[37] D. Mookherjee and S. Reichelstein. Dominant strategy implementation of Bayesian incentive compatible allocation rules. *Journal of Economic Theory*, 56(2):378–399, 1992.

[38] D. Moreno and J. Wooders. Coalition-proof equilibrium. *Games and Economic Behavior*, 17:80–112, 1996.

[39] T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proc. 25th PODC*, pages 35–44, July 2006.

[40] R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, Apr. 1983.

[41] M. Naor and A. Wool. Access control and signatures via quorum secret sharing. In *Proc. 3rd CCS*, pages 157–168, 1996.

[42] M. Piatek, T. Isdal, T. Anderson, A. Krishnamurthy, and A. Venkataramani. Do incentives build robustness in BitTorrent? In *NSDI '07*, pages 1–14, Apr. 2007.

[43] M. Rabin. How to exchange secrets by oblivious transfer. Technical Report Tech. Memo TR-81, Aiken Computation Laboratory, Harvard U., 1981.

[44] T. Rabin and M. Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. In *Proc. 21st STOC*, pages 73–85. ACM Press, 1989.

[45] S. Saroiu, K. P. Gummadi, and S. D. Gribble. A Measurement Study of Peer-to-Peer File Sharing Systems. Jan. 2002.

[46] M. A. Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, Apr. 1975.

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

[48] A. Shamir. How to share a secret. *Comm. ACM*, 22(11):612–613, 1979.

[49] T. K. Srikanth and S. Toueg. Simulating authenticated broadcasts to derive simple fault-tolerant algorithms. *Distributed Computing*, 2(2):80–94, 1987.

[50] M. Tennenholtz. Competitive safety analysis: robust decision-making in multi-agent systems. *J. Artif. Int. Res.*, 17:363–378, November 2002.

[51] A. W. Tucker. Acceptable points in general cooperative n-person games. *Annals of Mathematics Study 40*, 4:287–324, 1959.

[52] E. L. Wong, J. B. Leners, and L. Alvisi. It's on me! The benefit of altruism in BAR environments. In *DISC 2010*.