

Regret Freedom Isn't Free

Edmund L. Wong¹, Isaac Levy¹, Lorenzo Alvisi¹,
Allen Clement², and Mike Dahlin¹

¹ Department of Computer Science, The University of Texas at Austin

² Max Planck Institute for Software Systems

{elwong,isaac,lorenzo,dahlin}@cs.utexas.edu, aclement@mpi-sws.org

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. This paper investigates how to specify conditions (*i.e.*, a solution concept) for rational cooperation in an environment that also contains Byzantine and obedient peers. We find that *regret-free* approaches—which, inspired by traditional Byzantine fault tolerance, condition rational cooperation on identifying a strategy that proves a best response regardless of how Byzantine failures occur—are unattainable in many fault-tolerant distributed systems. We suggest an alternative *regret-braving* approach, in which rational nodes aim to best respond to their *expectations* regarding Byzantine failures: the chosen strategy guarantees no regret only to the extent that such expectations prove correct. While work on regret-braving solution concepts is just beginning, our preliminary results show that these solution concepts are not subject to the fundamental limitations inherent to regret freedom.

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.*, [5,24,36]) or deviate from the assigned protocol if it is in their interest to do so (*e.g.*, [1,36]). 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.*, [27,30,35]). 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.¹ 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 services should hence be designed to function in environments consisting of a mix of Byzantine, acquiescent,² and rational (or *BAR*) participants.

Building a *BAR*-tolerant system then involves two steps: 1) designing a Byzantine fault-tolerant protocol and 2) proving that rational peers will cooperate and follow the prescribed protocol. But how does one specify the conditions, *i.e.*, the *solution concept* [17], under which rational peers will be willing to cooperate?

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*. Similarly, it is appealing to aim for a notion of equilibrium in which rational nodes—either unilaterally or as a part of a coalition—cannot improve their utility by deviating *independent of who the t Byzantine nodes are and of how they behave*. This approach, elegantly formalized in the notion of (k, t) -robustness [3,4], is in principle very attractive: at equilibrium, peers will never have reason to regret their chosen strategy, which is guaranteed to prove a best response to any Byzantine strategy, independent of the identities of Byzantine nodes.

The main result of this paper is to show that, despite its appeal, a solution concept that guarantees *regret freedom* is fundamentally unable to yield non-trivial equilibria in games (which we name *communication games*) that capture three key characteristics of many practical fault-tolerant distributed systems: (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.

More, we find that weakening (k, t) -robustness, even considerably, seems unlikely to help. For example, suppose that, magically, all rational nodes in a communication game knew precisely the identity of all Byzantine nodes (but not their strategy); or, alternatively, that they knew their strategy (but not their identities). We find that in both cases a regret-free equilibrium can be achieved only under very limited circumstances.

These results are not interesting because of their proofs, which are straightforward, but because they show that in fault-tolerant distributed systems, conditioning rational cooperation on the expectation of regret freedom may be fundamentally too much to ask. Furthermore, the limitations of this approach

¹ Of course, arbitrarily faulty peers too can be modeled as rational peers who follow an unknown utility function. Unfortunately, doing so does not simplify the problem.

² We originally named these nodes *altruistic* [6] but have since been made aware [2] of the risk of confusing such peers (whose irrational generosity is only driven by obedience to the given protocol) with peers who are irrationally generous for arbitrary reasons. We believe that “acquiescent” better captures our original intentions.

appear hard to fix, since they are rooted in the universal quantifiers (*e.g.*, “for all strategies” or “for all sets of t Byzantine nodes”) that are at the very essence of regret freedom.

The second part of the paper points to a promising research agenda to overcome this impasse, an approach we call *regret braving*. Regret braving is motivated by the observation that rational agents that operate under uncertainty about the strategy of other players (as is the case when players are Byzantine) 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 follow a particular strategy as long as they cannot improve their utility with respect to their expectation about their environment—the worth of the traded asset, their comfort with risk, and what they believe will be the trends in the market—by deviating. Similarly, we consider solution concepts in which rational nodes aim to best respond to their *expectations* regarding Byzantine failures: the chosen strategy guarantees no regret only to the extent that such expectations prove correct.

We find that regret-braving solution concepts admit simple and intuitive equilibria for communication games where even the weakened versions of (k, t) -robustness could not. In particular, we consider two solution concepts: in the first, rational nodes 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 do *not* suggest that these solution concepts are the “right” ones or that they can be directly applied to every BAR-tolerant system; in fact, we believe that an exciting research opportunity lies in identifying increasingly realistic models for Byzantine failure expectations. What these preliminary results do show, however, is that regret-braving solution concepts are not subject to the fundamental limitations inherent to regret freedom.

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 why equilibria that base rational cooperation on regret freedom are fundamentally hard to achieve. Section 4 describes instead the home of the (regret) brave: we discuss two models of rational beliefs that admit useful equilibria in an instantiation of the communication game. Section 5 discusses related work, and Section 6 concludes.

2 Model

A *communication game* models any fault-tolerant system in which communication is not free and at least some nodes need to communicate in order to achieve the desired functionality.

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

- *Communication incurs some cost and does not generate immediate benefit to the sender,*

- *Communication incurs some cost to the receiver, and*
- *Benefit is obtained from functionality that (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, and we assume that messages are never lost.

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* $\sigma = (\sigma_x)_{x \in N}$ assigns a strategy σ_x to each node x , 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 . Following common game theory notation, we use (σ'_x, σ_{-x}) to denote the strategy profile in which x plays σ'_x and everyone else plays their component in σ (we also do this for sets of players, *e.g.*, (σ'_K, σ_{-K})), and we drop redundant parentheses when using a strategy profile as a parameter to a utility function, *e.g.*, $U_x(\sigma'_x, \sigma_{-x})$ vs. $U_x((\sigma'_x, \sigma_{-x}))$. We primarily focus on *non-trivial* strategy profiles, in which some positive utility is expected for at least one node; this implies that some communication must occur.

We are interested in systems that include Byzantine, rational, and (optionally) acquiescent nodes; each node x belongs to a *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). These assumptions do not affect our impossibility results, and they simplify the analysis for the positive results in regret braving—which, as in any game-theoretic analysis, depend on the types of players and solution concept. 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 neither trusted hardware nor trusted third-parties are used to monitor communication. Although a trusted mediator is useful [10,23,38], 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. it 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 rational cooperation to be achieved under similarly strong guarantees. (k, t) -robustness [3,4] 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* [3], captures the intuition that nodes following a strategy profile should not be adversely affected by up to t Byzantine failures.

Definition 3. *A strategy profile σ is t -immune if, for all $T \subseteq N$ such that $|T| \leq t$, all strategy profiles τ , and $x \notin T$,*

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

Note that t -immunity is *not* equivalent to Byzantine fault-tolerance, as t -immunity does not specify that a strategy profile σ must provide any sort of desirable safety or liveness properties despite t faults. In fact, any σ , fault-tolerant or not, is t -immune if it specifies actions so bad that Byzantine nodes, playing anything other than σ , cannot hurt a player's utility.

The second, *k-resilience* [3], addresses the possibility of collusion: a k -resilient strategy guarantees that a coalition of size at most k cannot deviate in a way that benefits every member.³

Definition 4. *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 5. *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 that makes (k, t) -robustness so appealing—regret freedom regardless of how Byzantine failures occur—makes it infeasible in many real-world systems. The reason, fundamentally, is that *communication always incurs cost but could potentially yield no benefit if one is communicating with a Byzantine node*. In other words, a rational node may realize in hindsight

³ 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 [3,4]. 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.

that it could have reduced its costs without affecting its benefits by avoiding all communication with Byzantine nodes, thus improving its utility. As any node can be Byzantine, this implies that the only possible (k, t) -robust equilibrium is one in which no node communicates.

Theorem 6. *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

$$U_x(\sigma_{-z}^*, \tau_z) \geq U_x(\sigma'_x, \tau_z, \sigma_{-\{x,z\}}^*) \quad (1)$$

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, x cannot prove that it communicated with z ; it thus follows that $(\sigma'_x, \tau_z, \sigma_{-\{x,z\}}^*)$ 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,

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

which directly contradicts inequality (1). □

More broadly, Theorem 6 suggests that it may be hard to build non-trivial (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 [3,4], a node's utility depends only on the game's outcome (*e.g.*, in a secret-sharing game based on Shamir's scheme, utility depends 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 6 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 6 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 *affect* a rational node’s utility by more than the cost of sending a *single* message have no regret-free equilibrium, even if the identity of all Byzantine nodes are known a priori.

We first define the equivalent of t -immunity (Definition 3) and (k, t) -robustness (Definition 5) for a fixed set T of Byzantine nodes.

Definition 7. A strategy profile σ is T -strategy-immune if for all strategy profiles τ and $x \notin T$,

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

Definition 8. A strategy profile σ^* is (k, T) -strategy-robust with respect to $T \subseteq N$ iff σ^* is T -strategy-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 $x \in K$,

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

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

Proof. Similar to proof of Theorem 6 (see [42]). □

Although Theorem 9 does not rule out all (k, T) -strategy-robust equilibria, Theorem 10 proves that these equilibria, which must be regret-free for *any* Byzantine strategy, only exist in limited circumstances.

Theorem 10. No communication game can yield a (k, T) -strategy-robust equilibrium for any set $T \subseteq N$ of Byzantine nodes if for some $x \notin T$ and some $z \in T$, (a) x has at least one opportunity to send a message to z and (b) for any strategy profile σ , there exist two Byzantine strategies τ_z and τ'_z such that τ_z and τ'_z are the same until x ’s first opportunity to communicate with z and

$$U_x(\sigma_{-z}, \tau_z) - U_x(\sigma_{-z}, \tau'_z) > \gamma$$

We omit the straightforward proof for lack of space (see [42]): in essence, if there exists a Byzantine strategy in which a rational node may *gain* by interacting with Byzantine nodes, then ignoring Byzantine players may not prove, in hindsight, an optimal strategy.

Theorem 10—unlike Theorem 6—provides conditions under which no (k, t) -strategy-robust equilibria exist, whether trivial or not. Since (k, t) -strategy-robust equilibria are a superset of (k, t) -robust equilibria, it naturally follows from Theorem 10 that no (k, t) -robust equilibria exist under the same conditions.

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 11. *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 6 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 12. *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 6. □

In general, it is possible to show (see [42]) that non-trivial (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,” *i.e.*, 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).

4 Dealing with Byzantine Failures through Regret Bravery

Finding a single strategy that is a best response against all possible Byzantine strategies or all possible t -sized subsets of Byzantine nodes (or both) appears fundamentally hard: regret-free solution concepts, for which rational cooperation depends on finding such a strategy, seem unlikely to provide a viable theoretical framework for many BAR-tolerant systems.

Regret bravery, the alternative we explore in this section, explicitly forgoes seeking a “universal” best response. Instead, it makes rational cooperation dependent on identifying a strategy that is a best response to the Byzantine behavior that rational nodes *expect* to be exposed to. Before we proceed to look at examples of regret-braving equilibria, we answer some natural questions.

Is aiming for a best response towards only a subset of all possible Byzantine behaviors in effect abdicating the general claims (and benefits) of Byzantine fault tolerance? No. Any BAR-tolerant protocol, independent of the underlying solution concept, must be a strategy that guarantees Byzantine fault tolerance. The choice of a solution concept is not about fault tolerance; rather, it specifies under which conditions rational nodes will be willing to follow a given strategy, fault-tolerant or not. Regret-braving solution concepts are motivated by the observation that rational nodes may be willing to cooperate even without the guarantee that the considered strategy will, in *all* circumstances, prove to be a best response.

Do regret-braving solution concepts limit how Byzantine node can behave? No more than a threshold t on the number of Byzantine faults limits a system to experience, in reality, more than t faults. Regret braving asks rational nodes to build a model of expected Byzantine behavior, but of course Byzantine nodes are in no way bound to follow that model. If Byzantine behavior does not match the expectation of rational nodes, then a regret-braving equilibrium strategy may not, in hindsight, prove to be a best response.

What is the right set of expectations when it comes to Byzantine behavior? It all depends on the application being considered. We discuss below two concrete examples inspired by approaches (maximin and Bayes equilibria) that have been extensively studied in the economics literature, but we do not claim that these solution concepts model “realistic” expectations for all distributed systems. For example, the maximin approach produces a best response to the expectation that the system always includes exactly t Byzantine nodes, when it may instead often be reasonable to expect that the actual number of Byzantine faults will be lower.⁴ Indeed, we believe that the challenge of finding equilibrium strategies under more flexible solution concepts is an extremely exciting research opportunity.

Regret Braving the *Quorum* Communication Game. 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 [38], replicated state machines [26] and terminating reliable broadcast [21] in which functionality is achieved if and only if some subset of nodes (a *quorum*) work together.

Definition 13. *A (synchronous) quorum game is an infinitely-repeated communication game where*

⁴ A worst-case attitude is actually not uncommon when designing fault-tolerant systems, even for benign failures. For instance, non-early stopping protocols for synchronous terminating reliable broadcast always run for $t + 1$ rounds, even in executions that experience no failures.

- There are at least 3 nodes ($n \geq 3$).
- The game repeats indefinitely. In every round, for each $y \in N$, a node $x \neq y$ decides whether to send a message (“contribute”) or not (“snub”) to y .
- At the end of the round, every $x \in N$ simultaneously (1) observes who contributed to it and (2) receives its payoff.⁵ 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 .⁶
- The total payoff is the δ -discounted sum of each individual round’s payoff, where $0 < \delta < 1$.

δ -discounting is a commonly-accepted way of handling utility in infinite-horizon games [17]. This models the reality that earning benefit (incurring cost) now is better (worse) than doing so later.⁷

We consider two concrete regret-braving solution concepts for the quorum game. In the first, rational nodes best-respond to fearing the worst, *i.e.*, they follow a maximin strategy with respect to Byzantine failures.

Definition 14. *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 $x \in K$,*

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

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. We use $\mu_x((R_{-T}, \tau_T) | R_K)$ to denote a rational node x ’s belief that all nodes $z \in T$ are Byzantine and of type (*i.e.*, playing strategy) τ_z and all nodes $w \notin T$ are rational (*i.e.*, of type R), given that there is some K (the coalition) in which $x \in K$ and all $y \in K$ are rational.

Definition 15. *The strategy profile/belief tuple (σ^*, μ^*) is a k -resilient Bayes equilibrium iff for all $K \subseteq N$ such that $|K| \leq k$, there does not exist an alternate*

⁵ In game theory parlance, the game is a simultaneous game; in distributed systems, synchronous.

⁶ 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 .

⁷ For example, it is often preferable to have a dollar now rather than later, since money can be invested and can earn interest in the meantime.

strategy profile σ' such that for all $x \in K$,

$$\begin{aligned} & \sum_{T \subseteq N \setminus K} \sum_{\tau} \mu_x^*((R_{-T}, \tau_T) | R_K) U_x(\sigma'_K, \tau_T, \sigma_{-(K \cup T)}^*) \\ & \geq \sum_{T \subseteq N \setminus K} \sum_{\tau} \mu_x^*((R_{-T}, \tau_T) | R_K) U_x(\sigma_{-T}^*, \tau_T) \end{aligned}$$

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

In both definitions, we extend previous work that uses regret-brave solution concepts [6,28,29,41] 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 4); however, we could have used any notion of collusion resilience, as this choice is orthogonal to how rational participants view Byzantine peers.

An example of a t -maximin equilibrium. We prove a k -resilient t -maximin equilibrium 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 actions. Thus, our results hold even if we augmented the game to allow coalition members to coordinate via cheap talk [12,16].

Theorem 16. *Let the strategy profile σ^* be defined as follows: any $x \in N$ following σ_x^* contributes to some $y \neq x$ iff x and y have always contributed to each other in the past and x has been snubbed by at most t different nodes. σ^* is a k -resilient t -maximin equilibrium if $q = n - t - 1$, $k \leq q$, and*

$$\frac{b}{\gamma} \geq \max \left(\frac{1 + \delta^2}{\delta^2} (n - 1), \frac{1}{1 - \delta} (t + k) + 1 \right) \tag{2}$$

Proof. (Sketch)⁸ Since $q = n - t - 1$, a rational node needs the cooperation of all other rational nodes to achieve a quorum; as $k \leq q$, a coalition cannot achieve quorum by itself.⁹ Consider some coalition K of size at most k . It can be easily verified that, given the conditions above, a coalition member $x \in K$ never snubs a cooperative, non-coalition node $y \notin K$ following σ^* . Intuitively, suppose x snubs y in some round r and y is not Byzantine. If t Byzantine nodes snub every node at least once by round r , y , having observed $t + 1$ snubs, will then snub every node in round $r + 1$. This causes all non-coalition nodes to follow suit and snub in round $r + 2$. It follows that all members of K , including x , will only receive up to $k - 1 < q$ other contributions for the remainder of the game starting from round $r + 2$. As this is not enough to achieve quorum, such a deviation results in the loss of benefit for the remainder of the game and is thus not worthwhile for K given the above conditions.

⁸ See [42] for the full details.

⁹ Recall that a node needs q other nodes to contribute in order to achieve quorum.

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.

Suppose that nodes in K play such an alternate strategy σ'_K in which some nodes in K snub, for the first time, some $x \in K$ in round r . Then Byzantine nodes may also snub x in order to cause x to lose quorum in round r . Specifically, by deviating, x may

- lose the benefit b it would have normally gained from playing σ^* ,
- save at most $(t + 1)\gamma$ from not receiving contributions from $t + 1$ members (the reason why x did not achieve quorum and lost benefit), and
- save at most $k\gamma$ from not contributing to other coalition members.

Therefore, as compared to σ_K^* , σ'_K loses x at least $b - (t + k + 1)\gamma$ in utility. However, in all subsequent rounds, x could save on contributing to

- Byzantine nodes that snubbed x in round r , saving at most $t\gamma$ per round (in the worst case, the Byzantine nodes still continue to contribute to x), and
- coalition members, saving at most $k\gamma$ per round.

This implies that x saves at most $\delta/(1 - \delta)(t + k)\gamma$ in utility over all subsequent rounds.

Thus, in order for σ'_K to be worthwhile for x , it must be the case that

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

which is never satisfied given inequality (2). □

An example of a Bayesian equilibrium. 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 and deviate from the prescribed strategy may be able 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 demonstrating a simple example of a k -resilient Bayes equilibrium. Our goal 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 use simple beliefs, optimistic bounds about the utility earned by deviating, and pessimistic bounds about the utility earned by cooperating.

Theorem 17. Define the strategy profile σ^* such that any $x \in N$, following σ_x^* , contributes to any $y \neq x$ iff x and y have always contributed to each other in the past and x has been snubbed by at most t peers, where t is some constant.

Let τ be defined as the random t -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 all subsets $K \subseteq N$ such that $|K| \leq k$ and all $y \in K$,

- $\mu_y^*((R_{-T}, \tau_T)|R_K) = 0$ for any T such that $|T| \neq t$, and
- $\mu_y^*((R_{-T_1}, \tau_{T_1})|R_K) = \mu_y^*((R_{-T_2}, \tau_{T_2})|R_K) > 0$ for any $T_1, T_2 \subseteq N \setminus K$ such that $|T_1| = |T_2| = t$.

Then (σ^*, μ^*) is a k -resilient Bayes equilibrium if $n > t + k$, $k \leq q$, and

$$\frac{b}{\gamma} \geq \frac{n + t - 1}{\rho^t \delta^2 (1 - \delta)} \frac{n - k}{n - k - t} + n - t - 1 \tag{3}$$

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 .

It can be easily verified that by following σ^* , each member of K , including x , earns no less than

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

in utility, since x can achieve quorum even if every Byzantine node crashes, so the “worst” that happens is x achieves quorum in every round while incurring cost from communication from everyone.

Suppose that x snubs some node $y \notin K$. Since the probability that a node is rational is uniform across all (non-coalition) 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 by 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 round payoff $b - q\gamma$. Thus, deviating is worthwhile only if

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

exceeds expression (4). This never holds given inequality (3).

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

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

which never holds given inequality (3). □

5 Related Work

Outside of (k, t) -robustness [3,4], Eliaz [15] also defined a solution concept which is effectively $(1, t)$ -robustness. Gradwohl [19] 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. [33] 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 [40] proposed a notion of collusion resilience which is the basis for k -resilience. Berheim et al. [11], Moreno et al. [32], Einy et al. [14], among others, have proposed weaker solution concepts that only consider deviations that are self-enforcing, meaning that there does not exist an even more profitable deviation for a sub-coalition within 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 [13,18,20,25,34,37] where mechanisms that incentivize nodes to reveal their true preferences or types for every possible realization of types are found to be often impossible or heavily restricted. Others [13,34] found positive results by using Bayesian solution concepts instead of dominant ones. Mookherjee et al. [31] define conditions in which 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 [39], 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 provides rigorous guarantees for 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 be able to admit equilibria in games where a rational node's payoff is not based simply on the outcome but also on the cost of the actions required to achieve said outcome. 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 [22]. We believe that regret-brave solution concepts provide a rigorous and realistic framework for games that account for these costs.

References

1. Kazaa Lite, http://en.wikipedia.org/wiki/Kazaa_Lite
2. Abraham, I., Dolev, D., Halpern, J.Y.: Private communication
3. Abraham, I., Dolev, D., Gonen, R., Halpern, J.: Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In: PODC 2006 (2006)
4. Abraham, I., Dolev, D., Halpern, J.Y.: Lower Bounds on Implementing Robust and Resilient Mediators. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 302–319. Springer, Heidelberg (2008), <http://portal.acm.org/citation.cfm?id=1802614.1802638>
5. Adar, E., Huberman, B.A.: Free riding on Gnutella. *First Monday* 5(10), 2–13 (2000), http://www.firstmonday.org/issues/issue5_10/adar/index.html
6. Aiyer, A.S., Alvisi, L., Clement, A., Dahlin, M., Martin, J.P., Porth, C.: BAR fault tolerance for cooperative services. In: SOSP 2005 (2005)
7. Alon, N., Emek, Y., Feldman, M., Tennenholtz, M.: Adversarial leakage in games. In: ICS 2010 (2010)
8. Aumann, R.J., Maschler, M.: Some thoughts on the minimax principle. *Management Science* 18(5), P54–P63 (1972)
9. Aumann, R.J.: On the non-transferable utility value: A comment on the Roth-Shaper examples. *Econometrica* 53(3), 667–677 (1985)
10. Ben-Porath, E.: Cheap talk in games with incomplete information. *Journal of Economic Theory* 108 (2003)
11. Bernheim, B.D., Peleg, B., Whinston, M.D.: Coalition-proof nash equilibria i. concepts. *Journal of Economic Theory* 42(1), 1–12 (1987)
12. Crawford, V.P., Sobel, J.: Strategic information transmission. *Econometrica* 50(6), 1431–1451 (1982)
13. d’Aspremont, C., Gerard-Varet, L.A.: Incentives and incomplete information. *Journal of Public Economics* 11(1), 25–45 (1979)
14. Einy, E., Peleg, B.: Coalition-proof communication equilibria. In: EISET 1995 (1995)
15. Eliaz, K.: Fault tolerant implementation. *Rev. of Econ. Studies* 69, 589–610 (2002)
16. Farrell, J., Rabin, M.: Cheap talk. *The Journal of Economic Perspectives* 10(3), 103–118 (1996)
17. Fudenberg, D., Tirole, J.: *Game Theory*. MIT Press (August 1991)
18. Gibbard, A.: Manipulation of voting schemes: A general result. *Econometrica* 41(4), 587–601 (1973)
19. Gradwohl, R.: Rationality in the full-information model. In: *Theory of Cryptography* (2010)
20. Green, J., Laffont, J.J.: On coalition incentive compatibility. *The Review of Economic Studies* 46(2), 243–254 (1979)
21. Hadzilacos, V., Toueg, S.: Fault-tolerant broadcasts and related problems (1993)
22. Halpern, J., Pass, R.: Game theory with costly computation (2010)

23. Halpern, J., Teague, V.: Rational secret sharing and multiparty computation. In: STOC 2004 (2004)
24. Hughes, D., Coulson, G., Walkerdine, J.: Free riding on Gnutella revisited: the bell tolls? *IEEE Distributed Systems Online* 6(6) (June 2005)
25. Jehiel, P., Meyer-ter Vehn, M., Moldovanu, B., Zame, W.R.: The limits of ex post implementation. *Econometrica* 74(3), 585–610 (2006), <http://dx.doi.org/10.1111/j.1468-0262.2006.00675.x>
26. Lamport, L.: Time, clocks, and the ordering of events in a distributed system. *CACM* (July 1978)
27. Levin, D., LaCurts, K., Spring, N., Bhattacharjee, B.: BitTorrent is an auction: analyzing and improving BitTorrent's incentives. *SIGCOMM Comput. Commun. Rev.* 38(4), 243–254 (2008)
28. Li, H., Clement, A., Marchetti, M., Kapritsos, M., Robinson, L., Alvisi, L., Dahlin, M.: FlightPath: Obedience vs choice in cooperative services. In: *OSDI 2008* (2008)
29. Li, H.C., Clement, A., Wong, E., Napper, J., Roy, I., Alvisi, L., Dahlin, M.: BAR Gossip. In: *OSDI 2006* (2006)
30. Locher, T., Moor, P., Schmid, S., Wattenhofer, R.: Free riding in bittorrent is cheap. In: *HotNets 2006* (2006)
31. Mookherjee, D., Reichelstein, S.: Dominant strategy implementation of Bayesian incentive compatible allocation rules. *Journal of Economic Theory* 56(2), 378–399 (1992)
32. Moreno, D., Wooders, J.: Coalition-proof equilibrium. *Games and Economic Behavior* 17, 80–112 (1996)
33. Moscibroda, T., Schmid, S., Wattenhofer, R.: When selfish meets evil: Byzantine players in a virus inoculation game. In: *PODC 2006* (2006)
34. Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29(2), 265–281 (1983)
35. Piatek, M., Isdal, T., Anderson, T., Krishnamurthy, A., Venkataramani, A.: Do incentives build robustness in BitTorrent? In: *NSDI 2007*, pp. 1–14 (April 2007)
36. Saroiu, S., Gummadi, K.P., Gribble, S.D.: A Measurement Study of Peer-to-Peer File Sharing Systems (January 2002)
37. Satterthwaite, M.A.: 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 (1975)
38. Shamir, A.: How to share a secret. *Comm. ACM* 22(11), 612–613 (1979)
39. Tennenholtz, M.: Competitive safety analysis: robust decision-making in multi-agent systems. *J. Artif. Int. Res.* 17, 363–378 (2002), <http://portal.acm.org/citation.cfm?id=1622810.1622822>
40. Tucker, A.W.: Acceptable points in general cooperative n-person games. *Annals of Mathematics Study* 40(4), 287–324 (1959)
41. Wong, E.L., Leners, J.B., Alvisi, L.: It's on Me! the Benefit of Altruism in BAR Environments. In: Lynch, N.A., Shvartsman, A.A. (eds.) *DISC 2010*. LNCS, vol. 6343, pp. 406–420. Springer, Heidelberg (2010)
42. Wong, E.L., Levy, I., Alvisi, L., Clement, A., Dahlin, M.: Regret freedom isn't free, <http://www.cs.utexas.edu/~elwong/research/publications/bar-no-regret-tr.pdf>