

Algorithmic Mechanisms for Reliable Master-Worker Internet-based Computing

Evgenia Christoforou, Antonio Fernández Anta, Chryssis Georgiou, Miguel A. Mosteiro

Abstract—We consider Internet-based master-worker computations, where a master processor assigns, across the Internet, a computational task to a set of untrusted worker processors, and collects their responses. Examples of such computations are the “@home” projects such as SETI. In this work various worker behaviors are considered. *Altruistic* workers always return the correct result of the task, *malicious* workers always return an incorrect result, and *rational* workers act based on their self interest. In a massive computation platform, such as the Internet, it is expected that all three type of workers coexist. Therefore, in this work we study Internet-based master-worker computations in the presence of malicious, altruistic, and rational workers. A stochastic distribution of the workers over the three types is assumed. In addition, we consider the possibility that the communication between the master and the workers is not reliable, and that workers could be unavailable. Considering all the three types of workers renders a combination of game-theoretic and classical distributed computing approaches to the design of mechanisms for reliable Internet-based computing. Indeed, in this work we design and analyze two algorithmic mechanisms to provide appropriate incentives to rational workers to act correctly, despite the malicious workers’ actions and the unreliability of the communication. Only when necessary, the incentives are used to force the rational players to a certain equilibrium (which forces the workers to be truthful) that overcomes the attempt of the malicious workers to deceive the master. Finally, the mechanisms are analyzed in two realistic Internet-based master-worker settings, a SETI-like one and a contractor-based one, such as Amazon’s Mechanical Turk. We also present plots that illustrate the trade-offs between reliability and cost, under different system parameters.

Index Terms—Algorithmic Mechanism design, Internet-based Computing, Reliability and Fault-tolerance, Untrusted workers, Unreliable communication.



1 INTRODUCTION

1.1 Motivation and Prior Work

As an alternative to expensive supercomputing parallel machines, the Internet has recently become feasible as a computational platform for processing complex computational jobs. Several Internet-oriented systems and protocols have been designed to operate on top of this global computation infrastructure; examples include Grid systems [17], [57], the “@home” projects [6], such as SETI [38], Amazon’s Mechanical Turk [5], and peer-to-peer computing–P2PC [23], [60]. Although the potential is great, the use of Internet-based computing is limited by the untrustworthy nature of the platform’s components [6], [26], [31]. Let us take SETI as an example. In SETI, data is distributed for processing to millions of voluntary machines around the world. At a conceptual level, in SETI there is a machine, call it the *master*, that sends jobs, across the Internet, to these computers, call them the *workers*. These workers execute and report back the result of the task computation. However, these workers are not trustworthy, and hence might report incorrect results. In

SETI, the master attempts to minimize the impact of these bogus results by assigning the same task to several workers and comparing their outcomes (that is, redundant task allocation is employed [6]), but there are also other methods [14], [36], [59].

This problem has recently been studied under two different views: from a “classical” distributed computing view [20], [37], [40], [54] and from a game-theoretic view [21], [60]. Under the first view, the workers are classified as either *malicious* (Byzantine) or *altruistic*, based on a predefined behavior. The malicious workers have a “bad” behavior which results in reporting an incorrect result to the master. This behavior is, for example, due to a hardware or a software error or due to an ill-state of the worker such as being a wrongdoer intentionally. Altruistic workers exhibit a “good” behavior, that is, they compute and return the correct task result (called simply “result” throughout the paper). From the perspective of the master, the altruistic workers are the “correct” ones. Under this view, “classical” distributed computing models are defined (e.g., a fixed bound on the probability of a worker being malicious is assumed) and typical malicious-tolerant voting protocols or distributed verification mechanisms are designed.

Under the game-theoretic view, workers act on their own *self-interest* and they do not have an a priori established behavior, that is, they are assumed to be *rational* [2], [26], [55]. In other words, the workers decide on whether they will be *honest* and report the correct result, or *cheat* and report a bogus result, depending on which strategy increases their benefit or *utility*. Under this view, Algorithmic Mechanisms [2], [12], [49] are employed, where games are designed to provide

- University of Cyprus, evgenia.christoforou@gmail.com
- Inst. IMDEA Networks and URJC, antonio.fernandez@imdea.org
- University of Cyprus, chryssis@cs.ucy.ac.cy
- Kean University and URJC, mosteiro@cs.rutgers.edu

This work is supported in part by the Cyprus Research Promotion Foundation grant ΤΠΕ/ΠΛΗΡΟ/0609(BE)/05, Comunidad de Madrid grant S2009TIC-1692, Spanish MICINN grant TIN2008–06735-C02-01, and the National Science Foundation (CCF 1114930, CCF 0937829). Parts of this work reporting preliminary results have appeared in the Proceedings of IPDPS 2010 and NCA 2011.

the necessary incentives so that processors' interests are best served by acting "correctly". In particular, the master provides some reward (resp. penalty) should a worker be honest (resp. cheat). The design objective is for the master to force a desired unique *Nash equilibrium* (NE) [48], i.e., a strategy choice by each worker such that none of them has incentive to change it. That Nash equilibrium is the one in which the master achieves a desired probability of obtaining the correct result.

The above views could complement one another, if a certain computation includes only malicious and altruistic workers, or only rational workers. However, the pragmatic situation on the Internet is different: all three types of workers might co-exist in a given computation. One could assume that all workers are rational but, for example, what if a software bug occurs that makes a worker deviate from its protocol, and hence compute and return an incorrect result? This worker is no longer exhibiting a rational behavior, but rather an erroneous or irrational one. From the master's point of view such behavior can be seen as malicious.

In this paper we consider the possibility that all three types of workers co-exist. Furthermore, we consider the possibility that the communication between the master and workers is not reliable. This communication uncertainty can either be due to communication-related failures or due to workers being slow in processing messages (or even crashing while doing so). For instance, Heien et al. [31] have found that in BOINC only around 5% of the workers are available more than 80% of the time, and that half of the workers are available less than 40% of the time. This fact, combined with the length of the computation [35], justifies the interest of considering in the Internet-based master-worker framework the possibility of workers not replying. We introduce the unreliability of communication in our model assuming that a worker's reply is received by the master with some probability smaller than 1.

Since it is possible that a worker's reply does not reach the master, we also allow workers to abstain from the computation. Imagine the situation where a rational worker returns the correct result but its reply is not received by the master. As we explain in Section 2, in this case the master does not reward the worker, but the worker has incurred the cost of performing the task. Hence, it is natural to allow the workers to abstain from replying, specially when the communication reliability is low. This strategy choice makes the task of the master even more challenging, as it needs to provide the necessary incentives to encourage rational workers to reply and to do so truthfully, even in the presence of low communication reliability.

1.2 Contributions

We study Internet-based master-worker computations under the assumption that each worker's behavior is either malicious, altruistic or rational. Furthermore, we also assume that a worker's output may never be received. The presence of all three types of workers, naturally renders a combination of game-theoretic and classical approaches to the design of algorithmic mechanisms for distributed computing. Our model captures the hardest shortcomings of an Internet-based

platform, yielding mechanisms that are resilient to undesired worker behavior and uncertainty of reply. In particular our contributions are as follows:

leftmargin=5mm

- We identify a collection of realistic payoff parameters and reward models and we formulate the Internet-based master-worker computation problem as a *Bayesian game* [30] (Section 2). We assume a probability distribution of workers among the worker types. The master and the workers do not know the type of other workers, only the probability distribution. The rational workers play a game looking for a Nash Equilibrium, choosing to be honest, cheat or abstain while the malicious and altruistic workers have a predefined strategy, malicious cheat and altruistic are honest. The master does not participate in the game, it only designs the game to be played. The communication reliability is modeled by a parametric probability.
- We develop and analyze two algorithms (a time-based algorithm and a reply-based one) that provide incentives to the rational workers to return the correct result, despite the malicious workers' actions and the communication unreliability (Section 3). The algorithms are parametrized in terms of a probability of auditing p_A (defined in Section 2), and a parametric probability d modeling the communication reliability. Each of the algorithms implements an instance of the Bayesian game. Under a general worker-type probability distribution, we analyze the master's utility and probability of success (probability of obtaining the correct result) and identify the conditions under which the game has a unique NE. The reason to enforce a *unique* NE is to achieve correctness taking advantage of the presence of rational players. As we show in the proof of Theorem 6, if multiple NE were allowed, choosing deterministically to cheat would be also an equilibrium strategy. Thus, for the purpose of a worst-case analysis with respect to the probability of success, it would have to be assumed that rational players choose to cheat, yielding the presence of rationals irrelevant. The reason to aim for a NE at all is that, although it is known that equilibria do not always yield optimal solutions, it is a "safe" way for the rational players to obtain high utility satisfaction [50, Chapter 1]. More importantly, a NE is *stable*, that is, once proposed, it is against the interest of the players to individually deviate.
- Under specific worker-type probability distributions, we design a protocol in which the master chooses the values of p_A to guarantee a parametrized bound on the probability of success (Section 3). Once this is achieved, the master also attempts to maximize its utility. This protocol together with each of the above-mentioned algorithms comprise a mechanism. Note that the mechanisms designed (and their analyses) are general in that reward models can either be fixed exogenously or be chosen by the master. It is also shown that our mechanisms are the only feasible approaches for the master to achieve a given

bound on the probability of success.

- Under the constraint of the bounded probability of success, we show how to maximize the master utility in two real-world scenarios (Section 4). The first scenario abstracts a system of volunteering computing like SETI [38]. The second scenario abstracts a contractor-based application where a company buys computational power from Internet users and sells it to computation-hungry consumers. One such application is Amazon’s Mechanical Turk [5] where the master and the workers can be in fact humans that contribute time for solving problems for profit.
- Finally, to provide a better insight on the usability of our mechanisms, and to illustrate the trade-offs between reliability and cost, we have characterized the utility of the master for the above-mentioned scenarios via plots by choosing system parameters as derived by empirical evaluations of master-worker Internet-based systems in [16] and [19].

1.3 Related work

Prior examples of game theory in distributed computing include work on Internet routing [25], [39], [45], [52], resource/facility location and sharing [24], [27], containment of viruses spreading [47], secret sharing [2], [29], P2P services [3], [42], [43] and task computations [21], [60]. For more discussion on the connection between game theory and distributed computing we refer the reader to the surveys by Halpern [28] and by Abraham, Alvisi and Halpern [1], and the book by Nisan et al [50].

Eliasz [18] seems to be the first to formally study the co-existence of Byzantine (malicious) and rational players. He introduces the notion of *k-fault-tolerant Nash Equilibrium* as a state in which no player benefits from unilaterally deviating despite up to k players acting maliciously. He demonstrates this concept by designing simple mechanisms that implement the constrained Walrasian function and a choice rule for the efficient allocation of an indivisible good (e.g., in auctions). Abraham et al [2] extend Eliasz’s concept to accommodate colluding rational players. In particular they design a secret sharing protocol and prove that it is (k, t) -robust, that is, it is correct despite up to k colluding rational players and t Byzantine ones.

Aiyer et al. [3] introduce the BAR model to reason about systems with Byzantine (malicious), Altruistic, and Rational participants. They also introduce the notion of a protocol being BAR-tolerant, that is, the protocol is resilient to both Byzantine faults and rational manipulation. (In this respect, one might say that our algorithmic mechanisms designed in this work is BAR-tolerant.) As an application, they designed a cooperative backup service for P2P systems, based on a BAR-tolerant replicated state machine. Li et al [43] also considered the BAR model to design a P2P live streaming application based on a BAR-tolerant gossip protocol. Both works employ incentive-based game theoretic techniques (to remove the selfish behavior), but the emphasis is on building a reasonably practical system (hence, formal analysis is traded

for practicality). Recently, Li et al [42] developed a P2P streaming application, called FlightPath, that provides a highly reliable data stream to a dynamic set of peers. FlightPath, as opposed to the above-mentioned BAR-based works, is based on mechanisms for *approximate equilibria* [11], rather than strict equilibria. In particular, ϵ -Nash equilibria are considered, in which rational players deviate if and only if they expect to benefit by more than a factor of ϵ . As the authors claim, the less restrictive nature of these equilibria enables the design of incentives to limit selfish behavior rigorously, while it provides sufficient flexibility to build practical systems.

Gairing [25] introduced and studied *malicious Bayesian congestion games*. These games extend congestion games [53] by allowing players to act in a malicious way. In particular, each player can either be rational or, with a certain probability, be malicious (with the sole goal of disturbing the other players). As in our work, players are not aware of each other’s type, and this uncertainty is described by a probability distribution. Among other results, Gairing shows that, unlike congestion games, these games do not in general possess a Nash Equilibrium in pure strategies. Also he studies the impact of malicious types on the social cost (the overall performance of the system) by measuring the so-called *Price of Malice*. This measure was first introduced by Moscibroda et al [47] to measure the influence of malicious behavior for a virus inoculation game involving both rational (selfish) and malicious nodes. Alon et al. [4] studied the implications of Bayesian ignorance.

Besides investigating the co-existence of malicious and rational players, also the co-existence of altruistic and rational players has been considered. Hoefler and Skopalik [32] study congestion games with altruist players, assuming a level of altruism β_i for each player i : $\beta_i = 0$ being a pure selfish and $\beta_i = 1$ being a pure altruist player. The work of Kuznetsov and Schmid [41] describes arbitrary social relationships between players through a social range matrix. Their work considers the existence of different degrees of rationality or altruism, and the existence of malicious players as well. Their definition of maliciousness and altruism is with respect to the whole set of players. (I.e., malicious players aim at reducing the utility of the rest of players, while altruistic players aim at increasing these utilities.) Instead, in the context of master-worker task computing, we assume that maliciousness and altruism is with respect to the master. At the end of Section 2.1 we discuss the possibility of considering different levels (or types) of rational players in our work, as they do.

Monderer and Tennenholtz [46] consider computations where an interested party wishes to influence the behavior of rational agents in a game, without having control of the game. In our work the master designs the game (and the protocol) to be played by the workers. Hence, the work in [46], in some sense, complements our work, as it is applicable to situations where it is not possible for the central authority to modify the game rules.

In our work we have the master process auditing, when needed, and verifying the workers’ answers. The work in [40] proposes a distributed verification mechanism in which the rational workers verify each other’s tasks. This can potentially

disburden the master, especially in multi-round computations. The focus in our work is to design a mechanism in such a way, that auditing is applied only when necessary, rather than designing a verification mechanism. However, it would be interesting to investigate whether our mechanism could be combined with a distributed verification mechanism as then one presented in [40].

Distributed computation in presence of selfishness was studied within the scope of combinatorial agencies in Economics [7]–[9], [15]. The basic model considered is a combinatorial variant of the classical principal-agent problem [44]: A master (principal) must motivate a collection of workers (agents) to exert costly effort on the master’s behalf, but the workers’ actions are hidden from the master. Instead of focusing on each worker’s actions, the focus is on complex combinations of the efforts of the workers that influence the outcome. Based on these complex dependencies between the workers’ actions, a *technology function* is defined by the master. In [7], where the problem was first introduced, the goal was to study how the utility of the master is affected if the equilibria space is limited to pure strategies. To that extent, the computation of a few Boolean functions is evaluated. In [9] mixed strategies were considered: if the parameters of the problem yield multiple mixed equilibrium points, it is assumed that workers accept one suggested by the master. The work in [15] investigates the effect of auditing by allowing the master to audit some workers (by random sampling) and verify their work. In our work, we do not use any technology, instead we implement our own algorithm. Furthermore, the master decides probabilistically whether to audit all workers or none, and the master assumes no dependencies between the workers. Another important difference is that in our framework, the worker’s actions are not *hidden*. The master receives a response by each worker and it is aware that either the worker has truthfully performed the task or not. The outcome is affected by each worker’s action in the case that no auditing is performed, but via auditing the master can determine the exact strategy used by each worker and apply a specific reward/punishment scheme. In the framework considered in combinatorial agency, the master witnesses the outcome of the computation, but it has no knowledge of the possible actions that the worker might take. For this purpose, the master needs to devise contracts for each worker based on the observed outcome of the computation and not on each worker’s possible action (as in our framework). Finally, our scheme considers worker punishment, as opposed to the schemes in combinatorial agency where workers cannot be fined (limited liability constraint); this is possible in our framework as worker’s actions are contractible and verifiable (either it performs a task or not).

2 MODEL AND DEFINITIONS

2.1 Master-workers Framework and Worker Types

We consider a distributed system consisting of a master processor that assigns, over the Internet, a computational task to a set of n workers to compute and return the result. The master, based on the received replies, must decide on the

value it believes is the correct outcome of the task. The tasks considered in this work are assumed to have a unique solution; although such limitation reduces the scope of application of the presented mechanisms [56], there are plenty of computations where the correct solution is unique: e.g., any mathematical function.

Each of the n workers has one of the following types, *rational*, *malicious*, or *altruistic*. The type of any worker w is known only by w . That is, neither the master nor the other workers know the type of worker w . Furthermore, the number of workers of each type is unknown to everyone. With respect to the worker types, the only knowledge available is a probability distribution over those types. Specifically, it is known that each worker is independently of one of the three types with probabilities p_ρ, p_μ, p_α , respectively, where $p_\rho + p_\mu + p_\alpha = 1$. The knowledge of the distribution over types could be obtained, for example, statistically from existing master-worker applications. If such information is inaccurate it is enough to overestimate p_μ and underestimate p_α (as we do in Section 4.3.1 from SETI-like systems [16], [19]) to achieve correctness, although at a bigger expense. Malicious workers always cheat and altruistic workers are always honest, independently of how such behavior impacts their utilities. In the context of this paper, being honest means to compute and return the correct result, and cheating means returning some incorrect value. On the other hand, rational workers are assumed to be selfish in a game-theoretic sense, that is, their aim is to maximize their benefit (utility) under the assumption that other workers do the same. So, a rational worker decides to be honest, cheat or not reply to the master (workers may choose not to reply) depending on which strategy maximizes its utility. As a result, each rational worker cheats with probability p_C , it is honest with probability p_H , and does not reply with probability p_N , such that $p_C + p_H + p_N = 1$. It is assumed that if a worker decides not to reply, then it does not perform the task.

The above model implies that all rational workers share the same probability distribution over the possible strategies (cheat, be honest, abstain), i.e., all rational workers are of the same type. Otherwise, in order to model the individuality of the non-monetary part of each rational worker’s benefit/penalty, the distribution over types could be generalized to different types of rational workers instead of one. More precisely, define a probability distribution over each possible combination of payoffs in \mathbb{R}^4 , restricting signs appropriately, so that each rational worker draws independently its strategic normal form from this distribution. However, the analysis presented here would be the same but using expected payoffs, the expectation taken over such distribution. Thus, for the sake of clarity and without loss of generality, we assume that the strategic normal form is unique for all players.

2.2 Communication Unreliability

The communication network is considered to be unreliable, and workers could be unavailable. These are very realistic assumptions for Internet-based master-worker computations as suggested, for example, by the work of Heien et al. [31]. We

model this shortcoming by assuming that the communication with each worker fails stochastically and independently of other workers.

Furthermore, we assume two settings, one where the probability of communication failure depends on time (the more the master waits for replies the larger the probability of obtaining more replies), and a second one where the probability of communication failure is fixed (hence, the more workers the master hires the larger the number of replies). As we will see in Section 3, the first setting leads to a *time-based* mechanism and the second one to a *reply-based* mechanism.

In our analysis, we let $d_1 > 0$ be the probability of any worker being available and receiving the task assignment message by the master, $d_2 > 0$ be the probability of the master receiving the worker’s response (has the worker chosen to reply), and $d = d_1 \cdot d_2$ be the probability of a round trip, that is, the probability that the master receives the reply from a given worker; that is, d represents the communication reliability. Hence, d_2 is the probability value that the master achieves by waiting T time (for the time-based mechanism) or hiring n workers (for the reply-based mechanism).

2.3 Master’s Objectives, Auditing, Payoffs and Reward Models

The objective of the master is twofold. First, the master has to guarantee that the decided value is correct with probability at least $1 - \epsilon$, for a known constant $0 \leq \epsilon < 1$. Then, having achieved this, the master wants to maximize its own benefit (utility). As, for example, in [54], [20] and [21], while it is assumed that workers make their decision individually and with no coordination, it is assumed that all the (malicious and rational) workers that cheat return the same incorrect value. This yields a worst-case scenario (and hence analysis) for the master with respect to its probability of obtaining the correct result; it subsumes models where cheaters do not necessarily return the same answer. (In some sense, this can be seen as a cost-free, weak form of collusion.)

To achieve its objectives, the master employs, if necessary, *auditing* and *reward/penalizing* schemes. The master might decide to audit the response of the workers (at a cost). In this work, auditing means that the master computes the task by itself, and checks which workers have been truthful or not. We denote by p_A the probability of the master auditing the responses of the workers.

Furthermore, the master can reward and punish workers, which can be used (possibly combined with auditing) to encourage rational workers to be honest (altruistic workers need no encouragement, and malicious workers do not care about their utility). When the master audits, it can accurately reward and punish workers. When the master does not audit, it decides on the majority of the received replies, and may apply different reward/penalizing schemes. In this work we consider three reward models shown in Table 1. Each reward model is essentially different from the others and can be used depending on the specifics of the application considered.

Auditing or not, the master neither rewards nor punishes a worker from whom it did not receive its response. Due to the

\mathcal{R}_m	the master rewards the majority only
\mathcal{R}_a	the master rewards all workers
\mathcal{R}_\emptyset	the master does not reward any worker

TABLE 1: Reward models

unreliability of the communication, when the master does not receive a reply from a worker it can not distinguish whether the worker decided to abstain, or there was a communication failure in the round trip (it could be the case that the worker did not even receive the task assignment message). Hence, it would be unfair to punish a worker for not getting its response; imagine the case where the worker received the request, performed the task and replied to the master, but this last message got lost!

The payoff parameters considered in this work are detailed in Table 2. Note that the first letter of the parameter’s name identifies whose parameter it is. M stands for master and W for worker. Then, the second letter gives the type of parameter. P stands for punishment, C for cost, and B for benefit.

WP_C	worker’s punishment for being caught cheating
WC_T	worker’s cost for computing the task
WB_Y	worker’s benefit from master’s acceptance
MP_W	master’s punishment for accepting a wrong answer
MC_Y	master’s cost for accepting the worker’s answer
MC_A	master’s cost for auditing worker’s answers
MC_S	master’s cost for not getting a “sufficient” replies
MB_R	master’s benefit from accepting the right answer

TABLE 2: Payoffs. All these parameters are non-negative.

Observe that there are different parameters for the reward WB_Y to a worker and the cost MC_Y of this reward to the master. This models the fact that the cost to the master might be different from the benefit for a worker. In fact, in some applications they may be completely unrelated. For example, in scenarios such as SETI, workers carry out the computation for free. Nevertheless, the master may still incur in some costs for processing the replies, posting a list of participants, etc. Although workers are not penalized for not replying, our model allows the possibility for the master to be penalized for not getting enough replies (parameter MC_S); the actual number of “enough” replies is quantified in Section 3. This provides an incentive for the master to choose (when it can) more workers to assign the task (especially if d is small) or to increase their incentives for replying; if convenient, MC_S could be set to zero. As usual in algorithmic mechanism design, we include a punishment in addition to the incentive. This is an implementation of a “carrot and stick” incentive-based mechanism when dealing with rational workers. Such mechanism is possible when the workers’ actions are contractible and verifiable as in our model (unlike the case of combinatorial agencies). Nevertheless, observe that, if needed, the punishment may be disabled setting $WP_C = 0$ (as some instances here). Among the parameters involved, we assume that the master has the freedom of choosing WB_Y and WP_C ; by tuning these parameters and choosing n , the master can achieve the desired trade-offs between correctness and cost. All other parameters can either be fixed because they are

system parameters or may also be chosen by the master.

2.4 Game Theory Concepts and Problem Formulation

We study the problem under the assumption that the rational workers, or *players*, will play a game looking for an equilibrium (recall that malicious and altruistic workers have a predefined strategy: malicious cheat, and altruistic are honest). The master does not play the game, it only defines the protocol and the parameters to be followed (i.e., it designs the game or mechanism). The master and the workers do not know the type of other workers, only the probability distribution. Hence, the game played is a so-called game with imperfect information or *Bayesian game* [30]. The action space is the set of pure strategies $\{\mathcal{C}, \mathcal{H}, \mathcal{N}\}$, and the belief of a player is the probability distribution over types.

More formally, the Internet-based Master-Worker computation considered in this work is formulated as the Bayesian game $\mathcal{G}(W, \varepsilon, \mathcal{D}, A, p_A, d_1, d_2, \mathcal{R}, pfs)$, where W is the set of n workers, $1 - \varepsilon \in [0, 1]$ is the desired success probability of the master obtaining the correct result, \mathcal{D} is the type probability distribution $(p_\rho, p_\mu, p_\alpha)$, $A = \{\mathcal{C}, \mathcal{H}, \mathcal{N}\}$ is the workers' actions space, p_A is the probability of the master auditing the workers' responses, d_1 and d_2 are the probabilities characterizing the reliability of the communication ($d = d_1 \cdot d_2$), \mathcal{R} is one of the reward models given in Table 1, and pfs are the payoffs as described in Table 2. Each player knows in advance the distribution over types \mathcal{D} , the total number of workers (n), the probability characterizing the communication reliability (d_1, d_2) and its normal strategic form, which is assumed to be unique.

The core of the mechanisms we develop is the computation of p_A . Based on the type distribution, the master must choose a value of p_A that would yield a *Nash Equilibrium* that best serves its purposes. Recall from [51], that for any finite game, a mixed strategy profile σ is a *mixed-strategy Nash equilibrium* (MSNE) if, and only if, for each player i ,

$$\begin{aligned} U_i(s_i, \sigma_{-i}) &= U_i(s'_i, \sigma_{-i}), \forall s_i, s'_i \in \text{supp}(\sigma_i), \\ U_i(s_i, \sigma_{-i}) &\geq U_i(s'_i, \sigma_{-i}), \\ &\forall s_i, s'_i : s_i \in \text{supp}(\sigma_i), s'_i \notin \text{supp}(\sigma_i), \end{aligned}$$

where s_i is the strategy used by player i in the strategy profile s , σ_i is the probability distribution over pure strategies used by player i in σ , σ_{-i} is the probability distribution over pure strategies used by each player but i in σ , $U_i(s_i, \sigma_{-i})$ is the expected utility of player i when using strategy s_i with mixed strategy profile σ , and $\text{supp}(\sigma_i)$ is the set of strategies in σ with positive probability.

The above definition applies to our setting as follows. First notice that there is no NE where some players choose a pure strategy and others do not, because the game is symmetric for all rational players. (Should many types of rational players be considered, then we would have to consider such a NE.) Assume first that there is a NE with mixed-strategies (that is, a NE where no strategy is chosen with probability 1). Then, the expected utility of a worker is the same for each pure strategy that such worker can choose with positive probability,

and it is not less than the expected utility of a pure strategy with probability zero of being chosen (if there is any). On the other hand, if only pure strategies are included in a NE (that is, there is only one strategy that can be chosen), that means that the expected utility of a worker is not less than the expected utility on the remaining pure strategies. Let us illustrate with an example. If $p_C = 1/2, p_H = 1/2, p_N = 0$ is a NE, that means that the expected utility of a worker is the same if it cheats or is honest, and it is not less than the utility if it does not reply. On the other hand, if $p_C = 0, p_H = 1, p_N = 0$ is a NE, then the expected utility of an honest worker is not less than the expected utility of cheating or not replying.

We denote by $\Delta U_{S_1 S_2}$ the difference on the expected utilities of a rational worker when choosing strategy S_1 over strategy S_2 . Then, for the purposes of the game we consider, in order to find conditions for equilibria, we want to study for each player i

$$\begin{cases} \Delta U_{\mathcal{H}\mathcal{C}} = \pi_{\mathcal{H}} \cdot w_{\mathcal{H}} - \pi_{\mathcal{C}} \cdot w_{\mathcal{C}} \\ \Delta U_{\mathcal{H}\mathcal{N}} = \pi_{\mathcal{H}} \cdot w_{\mathcal{H}} - \pi_{\mathcal{N}} \cdot w_{\mathcal{N}} \end{cases} \quad (1)$$

The expression $\pi_{\bullet} \cdot w_{\bullet}$ denotes the utility of the worker when choosing strategy \bullet ; we present the components of the expression in detail in Section 3.

The following notation will be used throughout:

$$\mathbf{P}_q^{(n)}(a, b) \triangleq \sum_{i=a}^b \binom{n}{i} q^i (1-q)^{n-i}.$$

The notation used throughout the paper is summarized in Table 3.

3 ALGORITHMIC MECHANISMS

In this section we present the mechanisms we design and analyze them. In particular, we show two different algorithms that the master runs in order to obtain the result of the task. Each of these algorithms is essentially an instance of the game we defined in the previous section. Before running one of the algorithms, the master must choose an appropriate value of p_A ; it does so by running a protocol we also present in this section. This protocol, together with each of the algorithms the master runs to obtain the tasks, comprises a mechanism.

3.1 Algorithms

As discussed in Section 2.2, we consider two different settings for modeling communication unreliability, which yield two different algorithms.

Figure 1 presents the *time-based* algorithm. Based on how the probability of communication failure depends on time, the master fixes a time T , it sends the specification of the task to be computed to n workers, and waits for replies. Once time T is reached, the master gathers all received replies, and chooses to audit the answers with probability p_A . If the answers were not audited, it accepts the result of the majority (ties are broken at random). Then, it applies the corresponding reward model.

Figure 2 presents the *reply-based* algorithm. Here the master, by appropriately choosing n , fixes k , an estimate of the minimum number of replies that wants to receive with high

$W = \{1, 2, \dots, n\}$	set of n workers
M	master processor
d_1	probability of a worker being available and receiving the task assignment message by the master
d_2	probability of the master receiving the worker's response (has the worker chosen to reply)
d	$d = d_1 \cdot d_2$, probability that the master receives a reply from a given worker
p_ρ	probability of a worker to be of rational type
p_μ	probability of a worker to be of malicious type
p_a	probability of a worker to be of altruistic type
p_A	probability that the master audits (computes task and checks worker answers)
P_{succ}	probability that the master obtains correct answer
ε	known constant $\varepsilon \in [0, 1]$, $1 - \varepsilon$ desired bound on the probability of success
$\{\mathcal{C}, \mathcal{H}, \mathcal{N}\}$	action space of a worker
p_C	probability of a worker to cheat
$p_{\mathcal{H}}$	probability of a worker to be honest
$p_{\mathcal{N}}$	probability of a worker not replying
s	strategy profile (a mapping from players to pure strategies)
s_i	strategy used by player i in the strategy profile s
s_{-i}	strategy used by each player but i in the strategy profile s
σ	mixed strategy profile (mapping from players to prob. distrib. over pure strat.)
σ_i	probability distribution over pure strategies used by player i in σ
σ_{-i}	probability distribution over pure strategies used by each player but i in σ
$U_i(s_i, \sigma_{-i})$	expected utility of player i with mixed strategy profile σ
$supp(\sigma_i)$	set of strategies of player i with probability > 0 in σ
$\Delta U_{S_1 S_2}$	difference on the expected utilities of a rational worker when choosing strategy S_1 over strategy S_2
$\mathbf{P}_q^{(n)}(a, b)$	$\sum_{i=a}^b \binom{n}{i} q^i (1-q)^{n-i}$

TABLE 3: Summary of Symbols

probability. (We discuss in the next subsection how k is computed and what is the probability of not receiving at least that many answers). The master sends the task specification to the n workers and gets replies. If at least k replies are received, then the master chooses to audit the answers with probability p_A and proceeds as the other protocol. In case that less than k replies are received, then the master does nothing and it incurs penalty MC_S .

Notice that both algorithms are one-shot, in the sense that they terminate after one round of communication between the master and the workers. This enables fast termination and avoids using complex cheater detection and worker reputation mechanisms. The benefit of one-round protocols is also partially supported by the work of Kondo et al. [35] that have demonstrated experimentally that there are common tasks that may take much more than one day of CPU time to complete. Having said that we do note that a multi-round computation could be analyzed by computing the expectations and probabilities in our analysis along all rounds. We leave this as subject of future work.

Each of the above algorithms basically implements an instance of the game we presented in Section 2.4. The master designs the game and the rational workers play looking for a Nash Equilibrium (NE) in an effort to maximize their benefit. Therefore, based on the type distribution, the master must choose the value of p_A that would yield a *unique* NE that best serves its purposes. The reason for uniqueness is to force all workers to the same strategy; this is similar to *strong implementation* in Mechanism Design, cf., [7], [49]. Multiple equilibria could be considered that could perhaps favor the utility of the master. However, in this work, correctness is the priority which, as shown later, our mechanisms guarantee.

In order to make the computation feasible to the workers, the master sends together with the task and the chosen value of p_A a *certificate* pointing out the only possible equilibrium. The certificate includes the strategy that the workers must play to achieve the unique NE together with the appropriate data to demonstrate this fact. These data include the system parameters/payoff values and the reward model; together with the value of p_A is enough to verify uniqueness (see the analysis in Section 3.2.3).

Recall that the main objective of the master is to achieve probability of accepting the correct result of at least $1 - \varepsilon$. Once this is achieved, then it seeks to maximize its utility as well. Based on the type distribution, it could be the case that the master may achieve this without relying on actions of the rational workers (e.g., the vast majority of workers are altruistic). Such cases fall into what we call the *free rationals scenario*. The cases in which the master needs to enforce the behavior of rational workers ($p_{\mathcal{H}}$) fall into what we call the *guided rationals scenario*. In this scenario, the master must choose p_A so that the benefit of the rational workers is maximized when $p_C = p_{\mathcal{N}} = 0$; in other words, rational workers choose to be honest ($p_{\mathcal{H}} = 1$) and hence they compute and truthfully return the correct result. The protocol ran by the master for choosing p_A is presented in Figure 3. Together with each of the algorithms in Figures 1 and 2 comprise our *mechanisms*. The analysis of the mechanisms and the lemmas referenced in Figure 3 are given in the next subsection.

Note that both designed mechanisms are useful and can be used depending on the setting. For example:

(a) As discussed in Section 2.2, the probability of the communication failure could depend on time, or be fixed. The master could have knowledge (e.g., based on statistics) of only one

```

1 send(task,  $p_A$ , certificate) to  $n$  workers
2 wait time  $T$  for replies
3 upon expire of time  $T$  do
4   audit the answers with probability  $p_A$ 
5   if the answers were not audited then
6     accept the majority
7   end if
8   apply the reward model

```

Fig. 1: Master's Time-based Algorithm

```

1 send(task,  $p_A$ , certificate) to  $n$  workers
2 if at least  $k$  replies are received then
3   audit the answers with probability  $p_A$ 
4   if the answers were not audited then
5     accept the majority
6   end if
7   apply the reward model
8 end if

```

Fig. 2: Master's Reply-based Algorithm

```

1 if  $\Pr[\text{majority honest} \mid \text{all rationals honest}] < 1 - \varepsilon$  then
2    $p_C \leftarrow 1$ ;  $p_N \leftarrow 0$ ;  $p_A \leftarrow 1 - \varepsilon / \sum_{i=k}^n r_i c_i$ ;
3 elseif  $\Pr[\text{majority honest} \mid \text{all rationals cheat}] \geq 1 - \varepsilon$  then
4    $p_C \leftarrow 1$ ;  $p_N \leftarrow 0$ ;  $p_A \leftarrow 0$ ;
5 elseif  $\Pr[\text{majority honest} \mid \text{all rationals honest}] \geq 1 - \varepsilon$  and
6    $\Delta U_{\mathcal{H}\mathcal{C}}(p_{\mathcal{H}} = 1, p_A = 0) \geq 0$  and  $\Delta U_{\mathcal{H}\mathcal{N}}(p_{\mathcal{H}} = 1, p_A = 0) \geq 0$  then
7    $p_C \leftarrow 0$ ;  $p_N \leftarrow 0$ ;  $p_A \leftarrow 0$ ;
8 else
9    $p_C \leftarrow 0$ ;  $p_N \leftarrow 0$ ; set  $p_A$  as in Lemma 4;
10 if  $U_M(p_A, p_N, p_C) < U_M(p_A = (1 - \varepsilon) / \sum_{i=k}^n r_i, p_N = 1, p_C = 0)$  then
11    $p_N \leftarrow 1$ ;  $p_A \leftarrow (1 - \varepsilon) / \sum_{i=k}^n r_i$ ;

```

Fig. 3: Master protocol to choose p_A . The expressions of k , r_i , and c_i are defined in Section 3.2

of the two settings. In such a case, it has no choice other than using the mechanism designed for that setting.

(b) It is not difficult to see that the time-based mechanism is more likely to use auditing than the other one, on the other hand, the reply-based mechanism runs the risk of not receiving enough replies. Hence, the time-based mechanism would be more preferable in case the cost of auditing is low, and the reply-based mechanism in case the cost of auditing is high and the value of parameter MC_S is small.

3.2 Equilibria Conditions and Analysis

We begin the analysis of our mechanisms by elucidating the following probabilities, expected utilities, and equilibria conditions. For succinctness, the analysis of both mechanisms is presented for a minimum number of replies k , where $k = 1$ for the time-based mechanism and $k \geq 1$ for the reply-based mechanism. For the latter, for a given worker type distribution, the choice of n workers, and d , even if all rational workers choose not to reply, the master will receive at least $E = nd(p_\alpha + p_\mu)$ replies in expectation. Thus, using Chernoff bounds, it can be shown that the master will receive at least $k = E - \sqrt{2E \ln(1/\zeta)}$ replies with probability at least $1 - \zeta$, for $0 < \zeta < 1$ and big enough n (e.g., $\zeta = 1/n$).

3.2.1 Probabilities and expected utilities.

Given the description of the mechanisms and the system parameters, it is not difficult to compute the following:

$$\Pr(\text{worker cheats} \mid \text{worker replies}): q = \frac{p_\mu + p_\rho p_C}{1 - p_\rho p_N}$$

$\Pr(\text{worker does not cheat} \mid \text{worker replies}):$

$$\hat{q} = \frac{p_\alpha + p_\rho p_{\mathcal{H}}}{1 - p_\rho p_N} = 1 - q$$

$\Pr(\text{reply received from worker}): r = d(1 - p_\rho p_N)$

$\Pr(\text{reply not received from worker}): \hat{r} = 1 - r$

Then, $r(q + \hat{q}) + \hat{r} = 1$.

$\Pr(i \text{ out of } n \text{ replies received}): r_i = \binom{n}{i} r^i \hat{r}^{n-i}$

$\Pr(\text{majority honest} \mid i \text{ replies received}):$

$$h_i = \sum_{j=0}^{\lfloor i/2 \rfloor - 1} \binom{i}{j} q^j \hat{q}^{i-j} + (1 + \lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \hat{q}^{\lceil i/2 \rceil}.$$

$\Pr(\text{majority cheats} \mid i \text{ replies received}):$

$$c_i = \sum_{j=\lceil i/2 \rceil + 1}^i \binom{i}{j} q^j \hat{q}^{i-j} + (1 + \lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lceil i/2 \rceil} q^{\lceil i/2 \rceil} \hat{q}^{\lfloor i/2 \rfloor}.$$

$\Pr(\text{master obtains correct answer}):$

$$P_{succ} = \sum_{i=k}^n r_i (p_A + (1 - p_A) h_i) \quad (2)$$

$E(\text{utility of master}):$

$$U_M = - \sum_{i=0}^{k-1} r_i \cdot MC_S + \sum_{i=k}^n r_i (p_A \alpha_i + (1 - p_A) \beta_i) \quad (3)$$

where,

$$\alpha_i = MB_{\mathcal{R}} - MC_{\mathcal{A}} - nd(p_\alpha + p_\rho p_{\mathcal{H}}) MC_{\mathcal{Y}}$$

$$\beta_i = MB_{\mathcal{R}} h_i - MP_{\mathcal{W}} c_i - MC_{\mathcal{Y}} \gamma_i$$

and where, $\gamma_i = 0$ for \mathcal{R}_\emptyset , $\gamma_i = i$ for \mathcal{R}_a , and for \mathcal{R}_m is,

$$\gamma_i = \sum_{j=\lceil i/2 \rceil + 1}^i \binom{i}{j} j (\hat{q}^j q^{i-j} + q^j \hat{q}^{i-j}) \\ + (1 + \lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lceil i/2 \rceil} [\lceil i/2 \rceil (\hat{q}^{\lceil i/2 \rceil} q^{\lfloor i/2 \rfloor} + q^{\lceil i/2 \rceil} \hat{q}^{\lfloor i/2 \rfloor})].$$

3.2.2 General Equilibria Conditions

Recall from Section 2.4 that Equation (1) states the conditions we want to study for each player i . In particular, as discussed there, we want $\Delta U_{\mathcal{H}C} \geq 0$ and $\Delta U_{\mathcal{H}N} \geq 0$.

The components of the vectors denoted by w_\bullet in (1) correspond to the different payoffs received by the given worker for each of the various events that may outcome from the game when the worker has chosen strategy \bullet , and the components of the vectors denoted by π_\bullet correspond to the probabilities that those events occur. Their detail values are given in Tables 4, 5, and 6; Table 7 lists the used notation. These conditions are defined so that a pure NE where $p_{\mathcal{H}} = 0$ is precluded.

3.2.3 Analysis Based on the Worker-type Distribution

Appropriate strategies to carry out the computation with the desired probability of success under the free rationals and guided rationals scenarios are considered in this section. It is important to stress again that, in order to obtain a mechanism that is useful for any of those scenarios we do not restrict ourselves to a particular instance of payoffs or reward models leaving those variables as parameters. Thus, we focus our study here on how to choose p_A to have the probability of success bounded by $1 - \varepsilon$ for each of the reward models assuming that the payoffs have already been chosen by the master or are fixed exogenously. For settings where payoffs and reward models are a choice of the master, its utility can be easily maximized choosing those parameters conveniently in Equation 3, as demonstrated in Section 4.

Although known, the worker-type distribution is assumed to be arbitrary. Likewise, the particular value of ε is arbitrary given that it is an input of the problem. Finally, although the priority is to obtain $P_{succ} \geq 1 - \varepsilon$, it is desirable to maximize the utility of the master under such restriction. Therefore, as it can be seen in Figure 3, the protocol the master runs for choosing p_A takes into account both the free rationals and guided rationals scenarios as discussed in Section 3.1.

We now proceed to analyze the different cases, first considering the free rationals scenario and then the guided rationals one.

3.2.3.1 Free Rationals.: Here we study the various cases where the behavior of rational workers does not need to be enforced. As mentioned before, the main goal is to carry out the computation obtaining the correct output with probability at least $1 - \varepsilon$. Provided that this goal is achieved, it is desirable to maximize the utility of the master. Hence if, for a given instance of the problem, the expected utility of the master utilizing the mechanism presented is smaller than the utility of just setting p_A big enough to guarantee the desired probability of correctness, independently of the outcome of the game, the

latter is used. We establish this observation in the following lemma.

Lemma 1. *In order to guarantee $P_{succ} \geq 1 - \varepsilon$, it is enough to set $p_A = (1 - \varepsilon) / \sum_{i=k}^n r_i$, making $p_N = 1$.*

Proof: Conditioning Equation 2 to be $\geq 1 - \varepsilon$, it is enough to make $p_A \geq \frac{1 - \varepsilon}{\sum_{i=k}^n r_i}$. Given that $\sum_{i=k}^n r_i$ is the probability that k or more replies are received, it is minimized when $p_N = 1$. Therefore, the claim follows. \square

We consider now pessimistic worker-type distributions, i.e., distributions where p_μ is so large that, even if all rationals choose to be honest, the probability of obtaining the correct answer is too small. Hence, the master has to audit with a probability big enough, perhaps bigger than the minimum needed to ensure that all rationals are honest. Nevertheless, for such p_A , rational workers still might use some NE where $p_{\mathcal{H}} < 1$. Thus, the worst case for P_{succ} has to be assumed. Formally,

Lemma 2. *In order to guarantee $P_{succ} \geq 1 - \varepsilon$, it is enough to set $p_A = 1 - \varepsilon / \sum_{i=k}^n r_i c_i$, making $p_C = 1$ and $p_N = 0$.*

Proof: Conditioning Equation 2 to be $\geq 1 - \varepsilon$, $p_A \geq 1 - \varepsilon / \sum_{i=k}^n r_i c_i$. Given that $\sum_{i=k}^n r_i c_i$ is the probability that k or more replies are received and the majority of them cheat, it is maximized when $p_C = 1$ (hence, $p_N = 0$). Therefore, the claim follows. \square

Now, we consider cases where no audit is needed to achieve the desired probability of correctness. I.e., we study conditions under the assumption that $p_A = 0$. The first case occurs when the type-distribution is such that, even if all rational workers cheat, the probability of having a majority of correct answers is at least $1 - \varepsilon$. A second case happens when the particular instance of the parameters of the game force a unique NE such that rationals are honest, even if they know that the result will not be audited. We establish those cases in the following lemma.

Lemma 3. *If any of the following holds:*

- $\sum_{i=k}^n r_i h_i \geq 1 - \varepsilon$ making $p_C = 1$ and $p_N = 0$; or
- $\sum_{i=k}^n r_i h_i \geq 1 - \varepsilon$ making $p_C = 0$ and $p_N = 0$ and there is a unique NE for $p_{\mathcal{H}} = 1$ and $p_A = 0$,

then, in order to guarantee $P_{succ} \geq 1 - \varepsilon$, it is enough to set $p_A = 0$.

Proof: Conditioning Equation 2 to be $\geq 1 - \varepsilon$ under the assumption that $p_A = 0$, it is enough

$$\sum_{i=k}^n r_i h_i \geq 1 - \varepsilon. \quad (4)$$

To find the condition for the case where even if all rationals cheat the probability of success is big enough, we replace $p_C = 1$ and $p_N = 0$ in Eq.(4). For the condition when the NE corresponds to some $p_C < 1$, we observe the following. Replacing in $\Delta U_{\mathcal{H}C}$ and $\Delta U_{\mathcal{H}N}$ for each reward model the value $p_A = 0$, it can be shown that $\Delta U_{\mathcal{H}C}(p_C, p_A = 0)$ is non-increasing in the interval $p_C \in [0, 1]$ for all three reward

		\mathcal{R}_m	\mathcal{R}_a	\mathcal{R}_\emptyset
w_C	w_C^{AR}	$-WP_C$	$-WP_C$	$-WP_C$
	w_C^{CR}	WB_y	WB_y	0
	$w_C^{\mathcal{H}R}$	0	WB_y	0
	$w_C^{\mathcal{X}\widehat{R}}$	0	0	0
$w_{\mathcal{H}}$	$w_{\mathcal{H}}^{AR}$	$WB_y - WC_{\mathcal{T}}$	$WB_y - WC_{\mathcal{T}}$	$WB_y - WC_{\mathcal{T}}$
	$w_{\mathcal{H}}^{CR}$	$-WC_{\mathcal{T}}$	$WB_y - WC_{\mathcal{T}}$	$-WC_{\mathcal{T}}$
	$w_{\mathcal{H}}^{\mathcal{H}R}$	$WB_y - WC_{\mathcal{T}}$	$WB_y - WC_{\mathcal{T}}$	$-WC_{\mathcal{T}}$
	$w_{\mathcal{H}}^{\mathcal{X}\widehat{R}}$	$-WC_{\mathcal{T}}$	$-WC_{\mathcal{T}}$	$-WC_{\mathcal{T}}$
$w_{\mathcal{N}}$	$w_{\mathcal{N}}^{\mathcal{X}\mathcal{X}}$	0	0	0

TABLE 4: Payoff vectors. Refer to Table 7 for notation.

π_C	π_C^{AR}	dp_A
	π_C^{CR}	$d(1-p_A) \sum_{i=0}^{n-1} \binom{n-1}{i} r^i \widehat{r}^{n-1-i} \left(\sum_{j=\lceil i/2 \rceil}^i \binom{i}{j} q^j \widehat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \widehat{q}^{\lfloor i/2 \rfloor} \right)$
	$\pi_C^{\mathcal{H}R}$	$d(1-p_A) \sum_{i=0}^{n-1} \binom{n-1}{i} r^i \widehat{r}^{n-1-i} \left(\sum_{j=0}^{\lfloor i/2 \rfloor} \binom{i}{j} q^j \widehat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \widehat{q}^{\lfloor i/2 \rfloor} \right)$
	$\pi_C^{\mathcal{X}\widehat{R}}$	$d_1(1-d_2)$
$\pi_{\mathcal{H}}$	$\pi_{\mathcal{H}}^{AR}$	dp_A
	$\pi_{\mathcal{H}}^{CR}$	$d(1-p_A) \sum_{i=0}^{n-1} \binom{n-1}{i} r^i \widehat{r}^{n-1-i} \left(\sum_{j=\lceil i/2 \rceil + 1}^i \binom{i}{j} q^j \widehat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \widehat{q}^{\lfloor i/2 \rfloor} \right)$
	$\pi_{\mathcal{H}}^{\mathcal{H}R}$	$d(1-p_A) \sum_{i=0}^{n-1} \binom{n-1}{i} r^i \widehat{r}^{n-1-i} \left(\sum_{j=0}^{\lfloor i/2 \rfloor} \binom{i}{j} q^j \widehat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \widehat{q}^{\lfloor i/2 \rfloor} \right)$
	$\pi_{\mathcal{H}}^{\mathcal{X}\widehat{R}}$	$d_1(1-d_2)$
$\pi_{\mathcal{N}}$	$\pi_{\mathcal{N}}^{\mathcal{X}\mathcal{X}}$	d_1

TABLE 5: Probability vectors for the time-based mechanism. Refer to Table 7 for notation.

$\pi_{\mathcal{C}}$	$\pi_{\mathcal{C}}^{AR}$	$dp_{\mathcal{A}} \sum_{i=k-1}^{n-1} \binom{n-1}{i} r^i \hat{r}^{n-1-i}$
	$\pi_{\mathcal{C}}^{\mathcal{C}\mathcal{R}}$	$d(1-p_{\mathcal{A}}) \sum_{i=k-1}^{n-1} \binom{n-1}{i} r^i \hat{r}^{n-1-i} \left(\sum_{j=\lceil i/2 \rceil}^i \binom{i}{j} q^j \hat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \hat{q}^{\lceil i/2 \rceil} \right)$
	$\pi_{\mathcal{C}}^{\mathcal{H}\mathcal{R}}$	$d(1-p_{\mathcal{A}}) \sum_{i=k-1}^{n-1} \binom{n-1}{i} r^i \hat{r}^{n-1-i} \left(\sum_{j=0}^{\lfloor i/2 \rfloor - 1} \binom{i}{j} q^j \hat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \hat{q}^{\lceil i/2 \rceil} \right)$
	$\pi_{\mathcal{C}}^{\mathcal{X}\hat{\mathcal{R}}}$	$d_1(1-d_2) + d \sum_{i=0}^{k-2} \binom{n-1}{i} r^i \hat{r}^{n-1-i}$
$\pi_{\mathcal{H}}$	$\pi_{\mathcal{H}}^{AR}$	$dp_{\mathcal{A}} \sum_{i=k-1}^{n-1} \binom{n-1}{i} r^i \hat{r}^{n-1-i}$
	$\pi_{\mathcal{H}}^{\mathcal{C}\mathcal{R}}$	$d(1-p_{\mathcal{A}}) \sum_{i=k-1}^{n-1} \binom{n-1}{i} r^i \hat{r}^{n-1-i} \left(\sum_{j=\lceil i/2 \rceil + 1}^i \binom{i}{j} q^j \hat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \hat{q}^{\lceil i/2 \rceil} \right)$
	$\pi_{\mathcal{H}}^{\mathcal{H}\mathcal{R}}$	$d(1-p_{\mathcal{A}}) \sum_{i=k-1}^{n-1} \binom{n-1}{i} r^i \hat{r}^{n-1-i} \left(\sum_{j=0}^{\lfloor i/2 \rfloor} \binom{i}{j} q^j \hat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \hat{q}^{\lceil i/2 \rceil} \right)$
	$\pi_{\mathcal{H}}^{\mathcal{X}\hat{\mathcal{R}}}$	$d_1(1-d_2) + d \sum_{i=0}^{k-2} \binom{n-1}{i} r^i \hat{r}^{n-1-i}$
$\pi_{\mathcal{N}}$	$\pi_{\mathcal{N}}^{\mathcal{X}\mathcal{X}}$	d_1

TABLE 6: Probability vectors for the reply-based mechanism. Refer to Table 7 for notation.

$w_{\bullet\bullet}$	payoff of event $\bullet \wedge \bullet \wedge \bullet$
$\pi_{\circ\bullet}$	probability of event $\bullet \wedge \bullet$, conditioned on the event \circ
$\ell_j^{\bullet\bullet}$	the worker has chosen strategy $j \in \{\mathcal{C}, \mathcal{H}, \mathcal{N}\}$
$\ell_{\bullet}^{\mathcal{A}\bullet}$	the master audits
$\ell_{\bullet}^{\mathcal{C}\bullet}$	the master does not audit and the majority cheats
$\ell_{\bullet}^{\mathcal{H}\bullet}$	the master does not audit and the majority does not cheat
$\ell_{\bullet}^{\mathcal{R}}$	the communication is successful and the master receives enough replies
$\ell_{\bullet}^{\hat{\mathcal{R}}}$	the communication fails or the master does not receive enough replies
\mathcal{X}	true (equivalent to “any value”)

TABLE 7: Notation for Tables 4, 5, and 6; $\ell \in \{w, \pi\}$.

models, and $\Delta U_{\mathcal{HN}}(p_{\mathcal{N}}, p_{\mathcal{A}} = 0)$ is non-increasing in the interval $p_{\mathcal{N}} \in [0, 1]$ for all three reward models as well. Thus, if $\Delta U_{\mathcal{HC}}(p_{\mathcal{C}} = 1, p_{\mathcal{A}} = 0) \geq 0$ and $\Delta U_{\mathcal{HN}=1}(p_{\mathcal{N}} = 1, p_{\mathcal{A}} = 0) \geq 0$, the rate of growth of $\Delta U_{\mathcal{HC}}$ and $\Delta U_{\mathcal{HN}}$ implies a single pure NE at $p_{\mathcal{H}} = 1$. Then, replacing $p_{\mathcal{C}} = 0$ and $p_{\mathcal{N}} = 0$ in Eq.(4) the claim follows. \square

3.2.3.2 Guided Rationals.: We now study worker-type distributions such that the master can take advantage of a specific NE to achieve the desired bound on the probability of success. Given that the scenario where all players cheat was considered in the free rationals scenario, here it is enough to study $\Delta U_{\mathcal{HC}}$ and $\Delta U_{\mathcal{HN}}$ for each reward model, conditioning $\Delta U_{\mathcal{HC}}(p_{\mathcal{C}} = 1) \geq 0$ and $\Delta U_{\mathcal{HN}}(p_{\mathcal{N}} = 1) \geq 0$ to obtain appropriate values for $p_{\mathcal{A}}$. As proved in the following lemma, the specific value $p_{\mathcal{A}}$ assigned depends on the reward model, and it is set so that a unique pure NE is forced at

$p_{\mathcal{H}} = 1$ (rendering the rationals truthful), and the correctness probability is achieved.

Lemma 4. *If $\sum_{i=k}^n r_i h_i < 1 - \varepsilon$ making $p_{\mathcal{C}} = 1$ and $p_{\mathcal{N}} = 0$, and $\sum_{i=k}^n r_i h_i \geq 1 - \varepsilon$ making $p_{\mathcal{C}} = 0$ and $p_{\mathcal{N}} = 0$ then, in order to guarantee $P_{succ} \geq 1 - \varepsilon$, it is enough to set $p_{\mathcal{A}}$ as follows.*

For \mathcal{R}_{\emptyset} ,

$$p_{\mathcal{A}} = \frac{WC_{\mathcal{T}}}{d_2 W_{By} \sum_{i=k-1}^{n-1} r'_i} \quad (5)$$

For \mathcal{R}_a ,

$$p_A = \frac{WC_T}{d_2(WB_y + WPC) \sum_{i=k-1}^{n-1} r'_i} \quad (6)$$

$$d_2 WB_y \sum_{i=k-1}^{n-1} r'_i \geq WC_T \quad (7)$$

For \mathcal{R}_m ,

$$p_A = \frac{WC_T/d_2 - WB_y \sum_{i=k-1}^{n-1} r'_i (h'_i - c'_i)}{(WB_y + WPC) \sum_{i=k-1}^{n-1} r'_i - WB_y \sum_{i=k-1}^{n-1} r'_i (h'_i - c'_i)} \quad (8)$$

$$p_A = \frac{WC_T/d_2 - WB_y \sum_{i=k-1}^{n-1} r'_i h'_i}{WB_y \sum_{i=k-1}^{n-1} r'_i - WB_y \sum_{i=k-1}^{n-1} r'_i h'_i} \quad (9)$$

Where

$$\begin{aligned} r'_i &= \binom{n-1}{i} r_i \hat{r}^{n-1-i}, \\ h'_i &= \sum_{j=0}^{\lfloor i/2 \rfloor} \binom{i}{j} q^j \hat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lceil i/2 \rceil} q^{\lceil i/2 \rceil} \hat{q}^{\lfloor i/2 \rfloor}, \\ c'_i &= \sum_{j=\lceil i/2 \rceil}^i \binom{i}{j} q^j \hat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor) \frac{1}{2} \binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \hat{q}^{\lceil i/2 \rceil}, \end{aligned}$$

for $p_C = 1$ in conditions (6) and (8), and for $p_N = 1$ in conditions (5), (7) and (9).

Proof: We compute the general conditions for each reward model from Equations (1). (Refer to Tables 4, 5, and 6 for details.) Recall that, for succinctness, the analysis of both mechanisms is presented for a number of replies k , where $k = 1$ for the time-based mechanism and $k = nd(p_\alpha + p_\mu) \left(1 - \sqrt{\frac{2 \ln(1/\zeta)}{nd(p_\alpha + p_\mu)}}\right)$ for the reply-based mechanism.

Conditions for reward model \mathcal{R}_\emptyset :

$$\Delta U_{\mathcal{H}C} = dp_A (WB_y + WPC) \sum_{i=k-1}^{n-1} r'_i - WC_T d_1 \geq 0$$

$$\Delta U_{\mathcal{H}N} = dp_A WB_y \sum_{i=k-1}^{n-1} r'_i - WC_T d_1 \geq 0$$

Thus, it is enough to use the latter condition only.

Conditions for the reward model \mathcal{R}_a :

$$\Delta U_{\mathcal{H}C} = dp_A (WB_y + WPC) \sum_{i=k-1}^{n-1} r'_i - WC_T d_1 \geq 0$$

$$\Delta U_{\mathcal{H}N} = d WB_y \sum_{i=k-1}^{n-1} r'_i - WC_T d_1 \geq 0$$

Conditions for the reward model \mathcal{R}_m :

$$\begin{aligned} \Delta U_{\mathcal{H}C} &= dp_A (WB_y + WPC) \sum_{i=k-1}^{n-1} r'_i - d_1 WC_T \\ &+ d(1 - p_A) WB_y \sum_{i=k-1}^{n-1} r'_i (h'_i - c'_i) \geq 0 \end{aligned} \quad (10)$$

$$\begin{aligned} \Delta U_{\mathcal{H}N} &= dp_A WB_y \sum_{i=k-1}^{n-1} r'_i - d_1 WC_T \\ &+ d(1 - p_A) WB_y \sum_{i=k-1}^{n-1} r'_i h'_i \geq 0 \end{aligned} \quad (11)$$

Notice that $\sum_{i=k-1}^{n-1} r'_i h'_i$ is the probability that at least $k-1$ other workers reply, and the majority of them is honest and $\sum_{i=k-1}^{n-1} r'_i c'_i$ is the probability that at least $k-1$ other workers reply, and the majority of them cheat. It can be seen that, when p_N is fixed, the equilibria condition 10 for this model is non-increasing on $p_C \in [0, 1 - p_N]$ as follows. Only $\sum_{i=k-1}^{n-1} r'_i (h'_i - c'_i)$ depends on p_C in this condition. When p_C increases and p_N is fixed, the probability that the majority of repliers is honest decreases. On the other hand, the probability that the majority cheats increases with p_C , but given that it is negated the slope is negative. Likewise, it can be seen that, when p_C is fixed, the equilibria condition 11 for this model is non-increasing on $p_N \in [0, 1 - p_C]$ as follows. Only $\sum_{i=k-1}^{n-1} r'_i h'_i$ depends on p_N in this condition. When p_N increases and p_C is fixed, the probability that the majority of repliers is honest decreases. Therefore, replacing in the above conditions for $\Delta U_{\mathcal{H}C}(p_C = 1) \geq 0$ and $\Delta U_{\mathcal{H}N}(p_N = 1) \geq 0$ the claim follows. \square

3.3 Correctness and Optimality

The following theorem proves the correctness of the mechanisms presented in Section 3.1. Its proof is the simple aggregation of the results presented in Section 3.2.

Theorem 5. *For any given system parameters, the values of p_A chosen after running the protocol depicted in Figure 3 satisfy that $P_{succ} \geq 1 - \varepsilon$.*

We now argue that only two approaches are feasible to bound the probability of accepting an incorrect value. In this respect, the strategy enforced by the mechanisms we designed is optimal.

Theorem 6. *In order to achieve $P_{succ} \geq 1 - \varepsilon$, the only feasible approaches are either to enforce a NE where $p_{\mathcal{H}} = 1$ or to use a p_A as shown in Lemma 2.*

Proof: It can be seen as in Lemma 4 that $\Delta U_{\mathcal{H}C}$ is non-increasing for $p_C \in [0, 1 - p_N]$ and $\Delta U_{\mathcal{H}N}$ is non-increasing for $p_N \in [0, 1 - p_C]$. Then, the only NE that can be made unique corresponds to $p_{\mathcal{H}} = 1$. Consider any other NE where $p_{\mathcal{H}} < 1$ (which is not unique). Then $p_C = 1$ and $p_N = 1$ are also both NE. In face of more than one equilibrium to choose from, different players might choose different ones. Thus, for the purpose of a worst-case analysis with respect to the probability of correctness, it has to be assumed the worst case, i.e. p_A has to be set as in Lemma 2. \square

3.4 Computational Issues

In Sections 3.1 and 3.2.3 we discussed a protocol for the master to choose appropriate values of p_A for different scenarios. A natural question is what is the computational cost of this protocol. In addition to simple arithmetical calculations, there

are two kinds of relevant computations required: binomial probabilities and verification of conditions for Nash equilibria. Both computations are n -th degree polynomial evaluations and can be carried out using any of the well-known numerical tools [34] with polynomial asymptotic cost. These numerical methods yield only approximations, but all these calculations are performed either to decide in which case the parameters fit in, or to assign a value to p_A , or to compare utilities. Given that these evaluations and assignments were obtained in the design as inequalities or restricted only to lower bounds, it is enough to choose the appropriate side of the approximation in each case.

Regarding the computational resources that rational workers require to carry out these calculations, notice that the choice of p_A in the mechanisms either yields a unique NE in $p_{\mathcal{H}} = 1$ or does not take advantage of the behavior of rational workers (Theorem 6). Furthermore, $p_C = 1$ was assumed as a worst case (with respect to probability of success). Notice from Tables 4–7 and the equilibrium conditions (eq. (1)) that setting $WP_C = WB_Y = 0$ for the cases where we do not use the behavior of the rational workers, $p_C = 1$ is a dominant strategy. (Recall that WB_Y and WP_C can be chosen by the master.) Thus, the mechanisms are enriched so that rational workers are enforced to use always a unique NE, either $p_C = 0$ or $p_C = 1$.

4 PUTTING THE MECHANISMS INTO ACTION

In this section two realistic scenarios in which the master-worker model considered could be naturally applicable are proposed. For these scenarios, we determine how to choose p_A and n in the case where the behavior of rational workers is enforced, i.e., under the conditions of Lemma 4. Again, for succinctness, the analysis of both mechanisms is presented for a number of replies k .

4.1 SETI-like Scenario

The first scenario considered is a volunteering computing system such as SETI@home, where users accept to donate part of their processors idle time to collaborate in the computation of large tasks. In this case, we assume that workers incur in no cost to perform the task, but they obtain a benefit by being recognized as having performed it (possibly in the form of prestige, e.g., by being included on SETI’s top contributors list). Hence, we assume that $WB_Y > WC_T = 0$. The master incurs in a (possibly small) cost MC_Y when rewarding a worker (e.g., by advertising its participation in the project). As assumed in the general model, in this model the master may audit the values returned by the workers, at a cost $MC_A > 0$. We also assume that the master obtains a benefit $MB_{\mathcal{R}} > MC_Y$ if it accepts the correct result of the task, and suffers a cost $MP_{\mathcal{W}} > MC_A$ if it accepts an incorrect value. Also it is assumed, as stressed before, that $d > 0$ (there is always a chance that the master will receive a reply from the worker).

Plugging $WC_T = 0$ in the lower bounds of Lemma 4 it can be seen that, for this scenario and conditions, in order to achieve the desired P_{succ} , it is enough to set p_A

arbitrarily close to 0 for all three models. So, we want to choose $\delta \leq p_A \leq 1$, with $\delta \rightarrow 0$, so that the utility of the master is maximized. Using calculus, it can be seen that U_M is monotonic in such range, but the growth of such function depends on the specific instance of the master-payoff parameters. Thus, it is enough to choose one of the extreme values of p_A . Replacing in Equation 3, we get

$$U_M \approx - \sum_{i=0}^{k-1} r_i MC_S + \sum_{i=k}^n r_i \max\{\alpha_i, \beta_i\}, \quad (12)$$

where $p_N = 0$ and α_i, β_i as in Equation (3). The approximation given in Equation (12) provides a mechanism to choose p_A and n so that U_M is maximized for $P_{succ} \geq 1 - \varepsilon$ for any given worker-type distribution, reward model, and set of payoff parameters in the SETI scenario.

4.2 Contractor Scenario

The second scenario considered is a company that buys computational power from Internet users and sells it to computation-hungry costumers, such as Amazon’s Mechanical Turk [5]. In this case the company pays the users an amount $S = WB_Y = MC_Y$ for using their computing capabilities, and charges the consumers another amount $MB_{\mathcal{R}} > MC_Y$ for the provided service. Since the users are not volunteers in this scenario, we assume that computing a task is not free for them (i.e., $WC_T > 0$), and that rational workers must have incentives to participate (i.e., $U > 0$). As in the previous case, we assume that the master verifies and has a cost for accepting a wrong value, such that $MP_{\mathcal{W}} > MC_A > 0$. Also as before we assume that $d > 0$ and $p_N = 0$.

As mentioned before, using calculus it can be seen that U_M is monotonic on p_A but the growth depends on the specific instance of master-payoff parameters. Thus, the maximum expected utility can be obtained for one of the extreme values. Trivially, 1 is an upper bound for p_A . For the lower bound, p_A must be appropriately bounded so that the utility of rational workers is positive and $P_{succ} \geq 1 - \varepsilon$. For example, for the \mathcal{R}_0 model, using Lemma 4 and conditioning $U > 0$, we get,

$$U_M = - \sum_{i=0}^{k-1} r_i MC_S + \sum_{i=k}^n r_i \max \left\{ \alpha_i, \beta_i + (\alpha_i - \beta_i) \frac{WC_T}{d_2 WB_Y \sum_{i=k-1}^{n-1} r'_i} \right\} \quad (13)$$

As in the previous section, the approximation given in Equation (13), and similar equations for the other reward models which are omitted for clarity, provide a mechanism to choose p_A and n so that U_M is maximized for $P_{succ} \geq 1 - \varepsilon$ for any given worker-type distribution, reward model, and set of payoff parameters in the contractor scenario.

4.3 Graphical Characterization of Master’s Utility

In this section, to provide a better insight of the usability of our mechanisms and to illustrate interesting trade-offs between

reliability and cost, we present a graphical characterization of the master’s utility, both in the SETI-like and the Contractor settings.

4.3.1 SETI-like Scenario

We begin by considering the timed-based mechanism, then the reply-based one, and then the special case of reliable communication. Recall that the only knowledge available on the workers type is a probability distribution. Such knowledge could be obtained statistically from existing master-worker applications such as [16], [19]. To err on the safe side, we overestimate p_μ and underestimate p_α with respect to those statistics.

4.3.1.1 Timed-based Mechanism.: For this mechanism, we consider $MC_A = 1$ as our normalizing parameter and we take $MP_W = 100$, $MC_S = 10$ and $MB_R = 4$ as realistically large enough values (with respect to $MC_A = 1$). Our experiments concluded that using other values for these parameters do not change qualitatively the results. We choose $p_\mu \in [0, 0.5]$ as we believe this is a reasonable interval. As it can be seen from the empirical evaluations of SETI-like systems reported in [16] and [19], p_μ is less than 0.1. So we took a larger range on p_μ to examine its general impact on the utility of the master. We choose $[0, 0.1]$ as the range of MC_Y , to reflect the small cost incurred by the master for maintaining a workers contribution list.

We consider three plot scenarios where we vary p_μ and MC_Y as discussed above:

(a) We fix $d = 0.9$ and $n = 75$ and compute the master’s utility for all three reward models. The results are depicted in Figure 4(a).

(b) We fix $n = 75$, we consider the \mathcal{R}_m model and compute the master’s utility over $d = 0.5, 0.9, 0.99$. See Figure 4(b).

(c) We fix $d = 0.9$, we consider the \mathcal{R}_m model and we compute the master’s utility over $n = 15, 55, 75$. The results are depicted in Figure 4(c).

In all plots we can notice a threshold where the behavior of the utility changes; this depicts the transition point in which the master’s strategy changes from non-auditing to auditing.

In Figure 4(a) we see that for all the reward models, the master does not audit until p_μ gets around 0.35. This behavior is reasonable, since in the presence of more malicious workers the master must audit to ensure correctness. Once auditing, the utility of the master becomes the same in all three reward model, as the same reward/penalize scheme is deployed. As expected, when the master does not audit, it gets its higher utility from \mathcal{R}_\emptyset and its lower utility from \mathcal{R}_a . The utility of the master for \mathcal{R}_m seems to balance nicely between the other two reward models. This may suggest that the \mathcal{R}_m model is the most stable among the three. A final observation is that as MC_Y gets bigger, for \mathcal{R}_m and \mathcal{R}_a , the utility of the master gets smaller; this is natural, since by increasing the payment to the workers the master decreases its own benefit.

In Fig 4(b) we observe that for smaller values of d we get a higher utility for the master. This is because the master receives fewer replies, and hence it rewards a smaller number of workers. As before, for any d , as MC_Y increases, U_M drops. An important observation is that for $d = \{0.9, 0.99\}$ and for

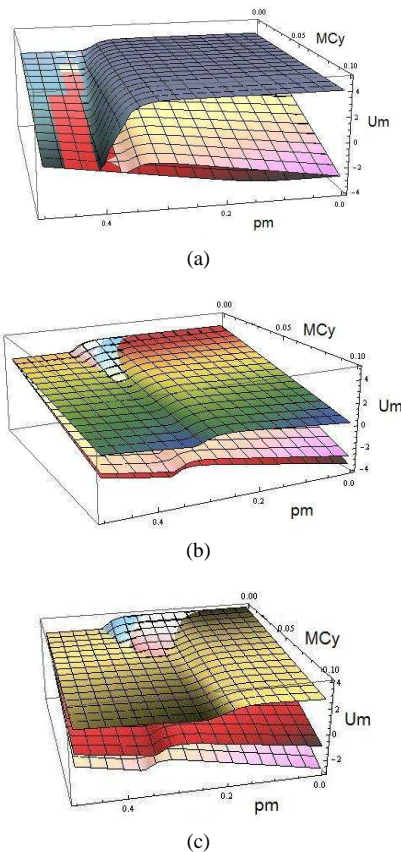


Fig. 4: Time-based Mechanism in the SETI-like scenario: Master’s utility for the three plot scenarios: (a) The upper plane corresponds to \mathcal{R}_\emptyset , the middle to \mathcal{R}_m , and the third to \mathcal{R}_a . (b) The upper plane corresponds to $d = 0.5$, the middle to $d = 0.9$, and the third to $d = 0.99$. (c) The upper plane corresponds to $n = 15$, the middle to $n = 55$, and the third to $n = 75$.

large values of MC_Y , the master’s utility is higher as it audits; the cost of rewarding the workers increases so much, that it is better for the master to audit.

In Figure 4(c) we notice that the utility of the master decreases as the number of workers increases; again, this is due to the reward it must provide to the workers. Observe that for $n = 15$, the master chooses to change its strategy to auditing for a smaller value of p_μ ; this is because as the master gets fewer replies, the probability of having a majority of incorrect replies gets bigger for smaller values of p_μ .

4.3.1.2 Reply-based Mechanism.: We now turn our attention to the reply-based mechanism. Our aim is to observe how the minimum number of replies (k) is affected by the number of workers chosen by the master (n) and the probability distribution of rational workers (p_ρ). Furthermore, we depict how k is affecting the utility of the master. As with the previous mechanism, we set $MC_A = 1$, $MP_W = 100$, $MC_S = 10$ and $MB_R = 4$. In the second scenario plotted, we choose a majority of rational workers to depict their effect on the master’s utility.

We consider two plot scenarios:

(a) We vary n from 65 to 95, p_ρ for 0 to 1, and we compute

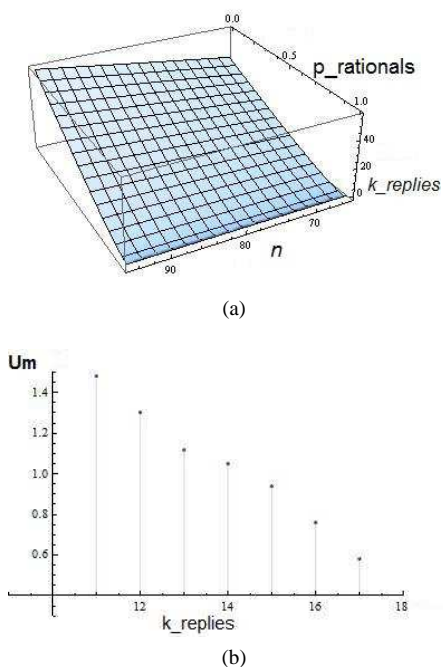


Fig. 5: Plots of the SETI-like Scenario for the Reply-based Mechanism.

the appropriate k that the master should choose for each n . The results are depicted in Figure 5(a).

(b) We use the \mathcal{R}_m , we fix $p_p = 0.6$, $d = 0.9$, $MC_y = 0.05$, we vary k and we compute the utility of the master. See Figure 5(b).

In Figure 5(a) we observe that as n increases, naturally, k increases as well. An interesting observation is that as p_p increases, k decreases. This is explained as follows: k is computed based on the number of malicious and altruistic workers that exist (since they always reply). Therefore, as these become fewer, k is naturally reduced. In Figure 5(b) we observe how the utility of the master is affected by k ; as k increases, the utility of the master decreases. This is because as the master gets more replies, it has to reward more workers.

4.3.1.3 Reliable Network.: We also provide the graphical characterization for the master's utility for the case that the communication is reliable, that is, $d = 1$. From this simpler case we can better study the trade-offs between reliability and cost without the complications of unreliable communication and workers not replying. Since the master receives all replies from the workers, the two mechanisms conceptually become the same (in other words, there is no sense to study two mechanisms under the knowledge that all messages are received). Notice that in this case MC_S is not applicable, hence its value is set to zero. As before, we set $MC_A = 1$ and $MP_W = 100$. We plot for values $p_\mu \in [0, 0.5]$ and $MC_y \in [0, 0.1]$. We consider three scenarios, applying the R_\emptyset model and varying p_μ and MC_y as discussed above. In particular:

(a) We fix $n=5$ and compute the utility of the master for $MB_{\mathcal{R}} = \{1, 4\}$; the results are depicted in Figure 6(a).

(b) We fix $n=15$ and compute the utility of the master for

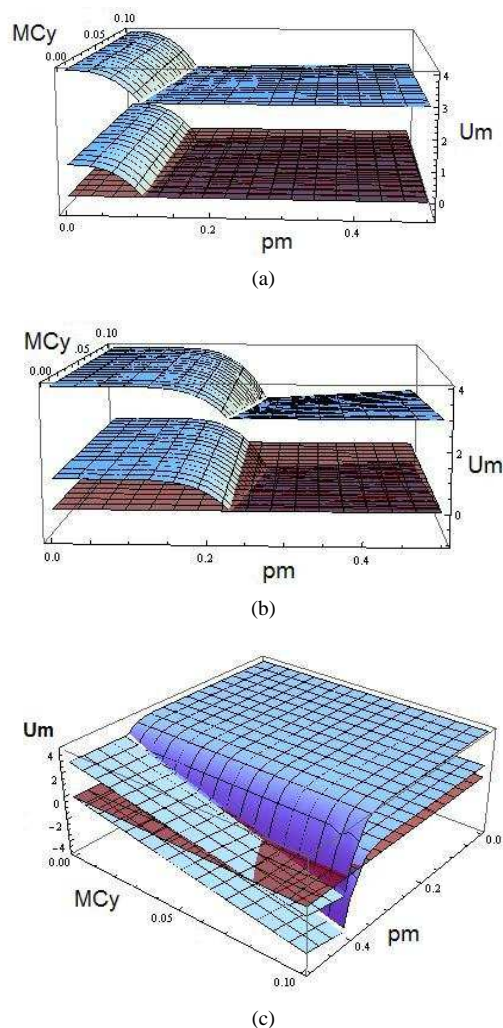


Fig. 6: Plots of the SETI-like scenario for $d = 1$. The upper plane corresponds to $MB_{\mathcal{R}} = 4$ the lower plane to $MB_{\mathcal{R}} = 1$ and the red flat plane to $U_M = 0$. (a) $n = 5$. (b) $n = 15$. (c) $n = 75$.

$MB_{\mathcal{R}} = \{1, 4\}$; the results are shown in Figure 6(b).

(c) We fix $n=75$ for both values of $MB_{\mathcal{R}}$ mentioned earlier; Figure 6(c) depicts the corresponding results.

All plots include a reference surface plane $U_M = 0$. Here we have only presented the R_\emptyset model because it is the simplest one. However, for the other reward models the plots depict more or less the same behavior, with the difference that before the threshold point (where the master does not audit) the utility of the master also depends on MC_y (e.g. Figure 4(c)).

In Figure 6 we observe, as expected, that the higher the value of $MB_{\mathcal{R}}$, the higher the utility of the master is, without this affecting the shape of the plot. In all plots we see a threshold where the behavior of the utility changes; this is the transition point where the master's strategy changes from non-auditing to auditing. For all three plots in Figure 6, we generally observe a smaller utility when the master audits than when it does not. Recall that we apply the R_\emptyset model when the master follows a non-auditing strategy; thus the master rewards the honest workers only when it audits and this decreases

its own utility proportionally to the value of payment to the workers (MC_Y). Another interesting observation is the sharp declining curve before the threshold (the master follows a non-auditing strategy). This curve reflects the fact that as p_μ increases the probability of the master getting an incorrect reply increases, and thus the utility of the master decreases accepting an incorrect reply. Notice that this declining curve is much sharper in Figure 6(c), since the larger the number of workers the more acute the impact of a high p_μ .

A significant difference between the number of chosen workers, is the threshold value of p_μ where the master changes its strategy to auditing. The larger the number of workers, the bigger the transition value (p_μ value) that the master starts to audit. This is due to the large reward it must provide when it audits, combined with the fact that having more workers increases the probability of getting the correct reply. We also notice that U_M increases slightly after the threshold, as p_μ increases. Although this behavior is not expected, we believe it is because the master has resolved to auditing in order to guarantee getting the correct value, and thus the fewer honest workers it has to reward, the greater its benefit.

4.3.2 Contractor Scenario

We now consider the contractor scenario (e.g., Amazon’s Mechanical Turk). Recall that in this setting $WC_T > 0$, and the workers are willing to participate only if their utility is positive (they are not volunteers as in the SETI-like setting). For this scenario we focus on the special case of reliable communication to illustrate how the cost for computing the task (WC_T) affects the trade-offs between reliability and cost (which we could not study in the SETI-like setting).

Figure 7 illustrates the utility of the master for the R_0 model and for a fix value of $S = 0.8$ (taken in analogy with MB_R); we vary $p_\mu \in [0, 0.5]$ and $WC_T \in [0, S]$. Notice that the workers’ cost for computing the task can not exceed their payment. In Figure 7(a) we fix $n=7$, in Figure 7(b) we fix $n=15$ and in Figure 7(c) we fix $n=75$. For each of these plots we have two planes, one for each value of $MB_R = \{1, 4\}$ and a reference surface plane $U_M = 0$ (similarly to the plots for the reliable communication case in the SETI-like setting).

Observe that a threshold point exists where the master changes its strategy from auditing with some probability (that guaranties the utility of the rational workers is positive) to auditing. We generally observe that (not surprisingly) for values of p_μ and WC_T close to zero we get the highest utility.

In all plots in Figure 7 when the master audits with some probability (before the threshold point) observe that as WC_T increases, the utility of the master decreases for every p_μ . This is a classical example of the trade-off between reliability and cost. The larger WC_T is, the higher the probability of p_A should be to guarantee correctness, thus the utility of the master decreases.

Another observation (especially in Figure 7(c)), is that before the threshold value, as p_μ increases, the utility of the master increases, and then decreases for every value of WC_T (except when close to $WC_T = 0$ and $WC_T = S$). When p_μ is increasing, the number of truthful workers decreases thus the master has to reward less honest workers and so its utility

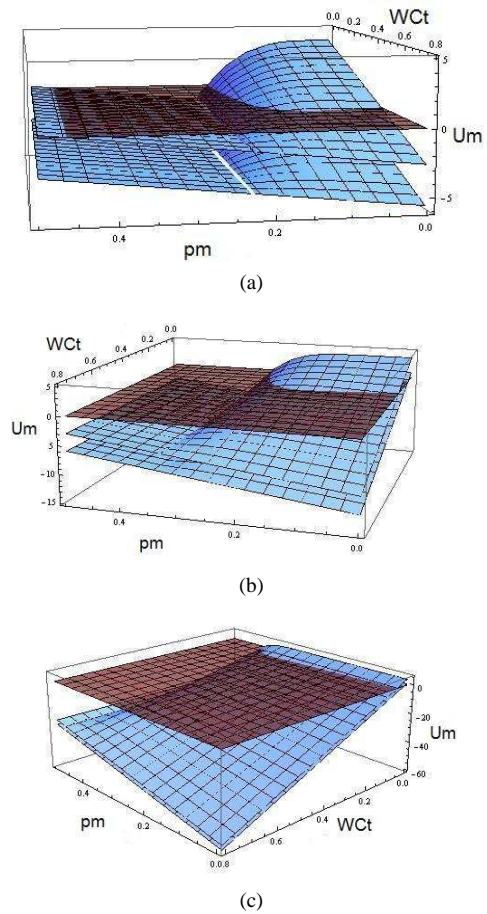


Fig. 7: Contractor Scenario plots for fixed S and $d = 1$. The upper plane corresponds to $MB_R = 4$ the lower plane to $MB_R = 1$ and the red flat plane to $U_M = 0$. (a) $n = 7$. (b) $n = 15$. (c) $n = 75$.

increases; recall that the master audits the answers with some probability. On the other hand, when the value of p_μ increases even more, the probability of having a majority of incorrect answers is very large. So it is quite probable since the master audits with some probability to get an incorrect result; thus its utility decreases.

Naturally when the master audits, for every value of WC_T , as p_μ increases so does its utility. Notice again that having larger MB_R does not affect the shape of the plots; the utility of the master increases uniformly. For similar reasons as in the SETI-like setting, the threshold value (p_μ value) increases for larger number of workers. Finally, observe the big decrease in the master’s utility as the number of workers grows. This is due to the large payments that the master has to give to large groups of workers to guarantee reliability.

5 DISCUSSION

In this paper we have combined a classical distributed computing approach (voting) with a game-theoretic one (cost-based incentives and payoffs). This has led to designing and analyzing two mechanisms that enable a master process to reliably obtain a result despite the co-existence of malicious,

altruistic and rational workers, and the underlying communication unreliability.

Several future directions emanate from this work. For example, in this work we have considered a cost-free, weak version of worker collusion (all rational cheaters and malicious workers return the same incorrect result). It would be interesting to study more involved collusions, as the ones studied in [2] or [13]. A possible approach would be to analyze a game among groups of colluders as in [22]. In this work, we have considered a single-task one-shot protocol, in which the master decides which result to accept in one round of message exchange with the workers. It would be interesting to consider several task waves over multiple rounds assuming that the workers' behavior changes over time, that is, view the computation as an *Evolutionary Game* [33], [58]. The master could use the knowledge gained in the previous rounds to increase its utility and its probability of success in future rounds. Issues such as worker *aspiration level* [10] could be taken into account.

REFERENCES

- [1] I. Abraham, L. Alvisi, and J.Y. Halpern. Distributed computing meets game theory: Combining insights from two fields. *ACM SIGACT News: Distributed Computing Column*, 42(2):69–76, 2011.
- [2] I. Abraham, D. Dolev, R. Goden, and J.Y. Halpern. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *proc. of PODC 2006*, pp. 53–62.
- [3] A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J. Martin, and C. Porth. BAR fault tolerance for cooperative services. In *proc. of SOSIP 2005*, pp. 45–58.
- [4] N. Alon, Y. Emek, M. Feldman, and M. Tennenholtz. Bayesian ignorance. In *proc. of PODC 2010*, pp. 384–391.
- [5] Amazon's Mechanical Turk, <https://www.mturk.com>.
- [6] D. Anderson. BOINC: A system for public-resource computing and storage. In *proc. of GRID 2004*, pp. 4–10.
- [7] M. Babaioff, M. Feldman, and N. Nisan. Combinatorial agency. In *proc. of ACM EC 2006*, pp. 18–28.
- [8] M. Babaioff, M. Feldman, and N. Nisan. Free riding and free labor in combinatorial agency. In *proc. of SAGT 2009*.
- [9] M. Babaioff, M. Feldman, and N. Nisan. Mixed Strategies in Combinatorial Agency. In *proc. of WINE 2006*, pp. 353–364.
- [10] J. Bendor, D. Mookherjee and D. Ray. Aspiration-based reinforcement learning in repeated interaction games: An overview. *International Game Theory Review*, 3:159–174, 2001.
- [11] S. Chien and A. Sinclair. Convergence to approximate Nash equilibria in congestion games. In *proc. of SODA 2007*, pp. 169–178.
- [12] G. Christodoulou and E. Koutsoupias. Mechanism design for scheduling. *Bulletin of the EATCS*, 97:39–59, 2009.
- [13] J. R. Douceur. The Sybil attack. In *proc. of IPTPS 2002*, pp. 251–260.
- [14] W. Du, J. Jia, M. Mangal, and M. Murugesan. Uncheatable grid computing. In *proc. of ICDCS 2004*, pp. 4–11.
- [15] R. Eidenbenz and S. Schmid. Combinatorial agency with audits. In *proc. of GameNets 2009*.
- [16] "Einstein@home", <http://einstein.phys.uwm.edu>.
- [17] "Enabling Grids for E-science", <http://www.eu-egee.org>.
- [18] K. Eliaz. Fault tolerant implementation. *Review of Economic Studies*, 69:589–610, 2002.
- [19] T. Estrada, M. Tauber, and D. P. Anderson. Performance prediction and analysis of BOINC projects: An empirical study with EmBOINC. *Journal of Grid Computing*, 7(4):537–554, 2009.
- [20] A. Fernández, Ch. Georgiou, L. Lopez, and A. Santos. Reliably executing tasks in the presence of untrusted processors. In *proc. of SRDS 2006*, pp. 39–50.
- [21] A. Fernández Anta, Ch. Georgiou, and M. A. Mosteiro. Designing mechanisms for reliable Internet-based computing. In *proc. of NCA 2008*, pp. 315–324.
- [22] A. Fernández Anta, C. Georgiou, and M. A. Mosteiro. Algorithmic Mechanisms for Reliable Internet-based Computing under Collusion. Technical Report RoSaC-2009-5, Univ. Rey Juan Carlos, 2009. <http://gsyc.es/tr-docs/RoSAC-2009-5.pdf>
- [23] I.T. Foster and A. Iamnitchi. On death, taxes, and the convergence of P2P and grid computing. In *proc. of IPTPS 2003*, pp. 118–128.
- [24] D. Fotakis. Memoryless facility location in one pass. In *proc. of STACS 2006*, pp. 608–620.
- [25] M. Gairing. Malicious Bayesian congestion games. In *proc. of WAOA 2008*, pp. 119–132.
- [26] P. Golle and I. Mironov. Uncheatable distributed computations. In *proc. of CT-RSA 2001*, pp. 425–440.
- [27] M. Halldorsson, J.Y. Halpern, L. Li, and V. Mirrokni. On spectrum sharing games. In *proc. of PODC 2004*, pp. 107–114.
- [28] J.Y. Halpern. Computer science and game theory: A brief survey. *Palgrave Dictionary of Economics*, 2007.
- [29] J.Y. Halpern and V. Teague. Rational secret sharing and multiparty computation. In *proc. of STOC 2004*, pp. 623–632.
- [30] J. C. Harsanyi. Games with incomplete information played by Bayesian players, I, II, III. *Management Science*, 14:159–182, 320–332, 468–502, 1967.
- [31] E.M. Heien, D.P. Anderson, and K. Hagihara. Computing low latency batches with unreliable workers in volunteer computing environments. *Journal of Grid Computing*, 7:501–518, 2009.
- [32] M. Hoefer and A. Skopalik. Altruism in atomic congestion games. In *proc. of ESA 2009*, pp. 179–189.
- [33] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
- [34] W. G. Horner. A new method of solving numerical equations of all orders by continuous approximation. *Philos. Trans. Roy. Soc. London* 109:308–335, 1819.
- [35] D. Kondo, F. Araujo, P. Malecot, P. Domingues, L. Silva, G. Fedak, and F. Cappello. Characterizing result errors in Internet desktop grids. In *proc. of Euro-Par 2007*, pp. 361–371.
- [36] D. Kondo, H. Casanova, E. Wing, and F. Berman. Models and scheduling mechanisms for global computing applications. In *proc. of IPDPS 2002*, pp. 79–86.
- [37] K.M. Konwar, S. Rajasekaran, and A.A. Shvartsman. Robust network supercomputing with malicious processes. In *proc. of DISC 2006*, pp. 474–488.
- [38] E. Korpela, D. Werthimer, D. Anderson, J. Cobb, and M. Lebofsky. SETI@home: Massively distributed computing for SETI. *Computing in Science and Engineering*, 3(1):78–83, 2001.
- [39] E. Koutsoupias and Ch. Papadimitriou. Worst-case equilibria. In *proc. of STACS 1999*, pp. 404–413.
- [40] M. Kuhn, S. Schmid, and R. Wattenhofer. Distributed asymmetric verification in computational grids. In *proc. of IPDPS 2008*, pp. 1–10.
- [41] P. Kuznetsov and S. Schmid. Towards network games with social preferences. In *proc. of SIROCCO 2010*, pp. 14–28, 2010.
- [42] H. C. Li, A. Clement, M. Marchetti, M. Kapritsos, L. Robison, L. Alvisi, and M. Dahlin. FlightPath: Obedience vs Choice in Cooperative Services. In *proc. of USENIX OSDI 2008*, pp. 355–368.
- [43] H. C. Li, A. Clement, E. L. Wong, J. Napper, I. Roy, L. Alvisi, and M. Dahlin. BAR gossip. In *proc. of OSDI 2006*, pp. 191–204.
- [44] A. Mass-Colell, M. Whinton, and J. Green. *Microeconomic Theory*, Oxford University Press, 1995.
- [45] M. Mavronicolas and P. Spirakis. The price of selfish routing. *Algorithmica*, 48(1):91–126, 2007.
- [46] D. Monderer and M. Tennenholtz. k-Implementation. In *proc. of EC 2003*, pp. 19–28.
- [47] T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: byzantine players in a virus inoculation game. In *proc. of PODC 2006*, pp. 35–44.
- [48] J.F. Nash. Equilibrium points in n -person games. *National Academy of Sciences*, 36(1):48–49, 1950.
- [49] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.
- [50] N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [51] M. J. Osborne. *An Introduction to Game Theory*. Oxford University Press, 2003.
- [52] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of ACM*, 49(2):236–259, 2002.
- [53] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [54] L. Sarmenta. Sabotage-tolerance mechanisms for volunteer computing systems. *Future Generation Computer Systems*, 18(4):561–572, 2002.
- [55] J. Shneidman and D.C. Parkes. Rationality and self-interest in P2P networks. In *proc. of IPTPS 2003*, pp. 139–148.

- [56] M. Tauber, D. Anderson, P. Cicotti, and C. L. Brooks. Homogeneous redundancy: a technique to ensure integrity of molecular simulation results using public computing. In *proc. of IPDPS 2005*.
- [57] "TeraGrid", <http://www.teragrid.org>.
- [58] J.W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, 1995.
- [59] S. Wong. An authentication protocol in web-computing. In *proc. of IPDPS 2006*.
- [60] M. Yurkewych, B.N. Levine, and A.L. Rosenberg. On the cost-ineffectiveness of redundancy in commercial P2P computing. In *proc. of CCS 2005*, pp. 280–288.

AUTHORS BIOGRAPHIES



Miguel A. Mosteiro is an Assistant Professor in the Department of Computer Science at Kean University, USA. He is also affiliated with the Department of Telematic Systems and Computer Science at Universidad Rey Juan Carlos, Spain. Previously, he was a Research Professor in the Department of Computer Science at Rutgers University, USA. He received his MS and PhD degrees in Computer Science from Rutgers University in 2003 and 2006 respectively. His current research interests lie in the analysis of algorithmic problems in limited-resource environments such as Radio Networks, and the application of Algorithmic Game Theory to Distributed Computing.



Evgenia Christoforou is a Research Associate at the Department of Computer Science in the University of Cyprus, Cyprus. She received her BSc and MSc degrees in Computer Science from the University of Cyprus, Cyprus, in 2010 and 2012 respectively. Her current research interests lie in Algorithmic and Evolutionary Game Theory.



Antonio Fernández Anta is Senior Researcher at the Institute IMDEA Networks in Madrid, Spain, on leave from his position of Professor at the Universidad Rey Juan Carlos. Previously, he was a member of the faculty at the Universidad Politécnica de Madrid. He graduated in Computer Science from the Universidad Politécnica de Madrid in 1991. He got his PhD in Computer Science from the University of Louisiana at Lafayette in 1994 and was a postdoc at MIT from 1995 to 1997. He is Senior Member of the

IEEE and the ACM. His research interests include computer networks, distributed processing, algorithms, and discrete and applied mathematics.



Chryssis Georgiou is an Assistant Professor of Computer Science at the University of Cyprus, Cyprus. He received a BSc in Mathematics from the University of Cyprus in 1998, and MSc and PhD degrees in Computer Science and Engineering from the University of Connecticut, USA, in 2002 and 2003 respectively. His research interests span the Theory and Practice of Fault-Tolerant Distributed and Parallel Computing with a focus on Algorithms and Complexity.