

Designing Mechanisms for Reliable Internet-based Computing*

Antonio Fernández Anta
LADyR, GSyC,
Univ. Rey Juan Carlos
anto@gsync.es

Chryssis Georgiou
Computer Science Dept.,
Univ. of Cyprus
chryssis@cs.ucy.ac.cy

Miguel A. Mosteiro
LADyR, GSyC, Univ. Rey Juan Carlos,
Computer Science Dept., Rutgers Univ.
mosteiro@cs.rutgers.edu

Abstract

In this work, using a game-theoretic approach, cost-sensitive mechanisms that lead to reliable Internet-based computing are designed. In particular, we consider Internet-based master-worker computations, where a master processor assigns, across the Internet, a computational task to a set of potentially untrusted worker processors and collects their responses. Several game-theoretic models that capture the nature of the problem are analyzed and mechanisms that, for each given set of cost and system parameters, achieve high reliability are designed. Additionally, two specific realistic system scenarios are studied. These scenarios are a system of volunteering computing like SETI, and a company that buys computing cycles from Internet computers and sells them to its customers in the form of a task-computation service. Notably, under certain conditions, non redundant allocation yields the best trade-off between cost and reliability.

1 Introduction

Motivation. As traditional one-processor machines have limited computational resources, and powerful parallel machines are very expensive to obtain and maintain, the Internet is emerging as the computational platform of choice 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 [5, 27] and the “@home” projects [2], such as SETI [17] (a classical example of *volunteering computing*). Although the potential is great, the use of Internet-based computing (also referred as P2P computing–P2PC [9, 28]) is limited by the untrustworthi-

ness nature of the platform’s components [2, 11]. Let us take SETI as an example. In SETI, data are distributed for processing to millions of voluntary machines around the world. At a conceptual level, in SETI there is a machine (which we can call the *master*) that sends jobs, across the Internet, to these computers (which we can call 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. Usually, 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 [2]).

In this paper, we consider Internet-based master-worker computations from a game-theoretic point of view. Specifically, we model these computations as games where each worker chooses whether to be *honest* (that is, compute and return the correct task result) or a *cheater* (that is, fabricate a bogus result and return it to the master). We design cost-sensitive mechanisms that provide the necessary incentive for the workers to truthfully compute and report the correct result. The objective is to maximize the probability of the master of obtaining the correct result while minimizing its cost (or alternatively, increasing its benefit).

Additionally, we identify and propose mechanisms for two paradigmatic applications. Namely, a volunteering computing system as the aforementioned SETI where computing processors are altruistic, and a second scenario where a company distributes computing tasks among contractor processors that get an economic reward in exchange.

Background and Prior/Related Work. Prior examples of game theory in distributed computing include work on Internet routing [18, 19, 25], resource/facility location and sharing [10, 13], containment of viruses spreading [20], secret sharing [1, 15] and task computations [28]. For more discussion on the connection between game theory and computing we refer the reader to the survey by Halpern [14] and the book by Nisan et al. [23].

In traditional distributed computing, the behavior of the

*This research was supported in part by Comunidad de Madrid grant S-0505/TIC/0285; Spanish MEC grant TIN2005-09198-C02-01; EU IST #15964 (AEOLUS); EU Marie Curie European Reintegration Grant IRG 210021; and NSF grant CCF 0632838.

The authors appear in alphabetical order.

system components (i.e., processors) is characterized a priori as either good or bad, depending on whether they follow the prescribed protocol or not. In game theory, processors are assumed to act on their own *self-interest* and they do not have an a priori established behavior. Such processors are usually referred as *rational* [1, 11]. In other words, the processors decide on how to act in an attempt to increase their own benefit, or alternatively to lower their own cost.

In *algorithmic mechanisms design* [1, 6, 22, 24], games are designed to provide the necessary incentives so that processors’ interests are best served by acting “correctly.” The usual practice is to provide some reward (resp. penalty) should the processors (resp. do not) behave as desired. The design objective is to force a desired unique *Nash equilibrium* (NE) [21], i.e., a strategy choice by each game participant such that none of them has incentive to change it.

In [7, 16] reliable master-worker computations have been considered by redundant task-allocation. In these works probabilistic guarantees of obtaining the correct result while minimizing the cost (number of workers chosen to perform the task) are also shown. However, a traditional distributed computing approach is used, in which the behavior of each worker is pre-defined. In this paper, much richer payoff parameters are studied and the behavior of each worker is not pre-defined, introducing new challenges that naturally drive to a game-theoretic approach.

Master-worker computations in a game-theoretic model have been studied before [28]. In that paper, the master can audit the results returned by rational workers with a tunable probability. Bounds for that audit probability are computed to guarantee that workers have incentives to be honest in three scenarios: redundant allocation with and without collusion¹, and single-worker allocation. They conclude that, in their model, single-worker allocation is a cost-effective mechanism specially in presence of collusion. Although our model comprises a weaker type of collusion, similar conclusions are reached here under certain system parameters. Additionally, our paper complements their work in various ways, such as studying more algorithms and games, including a richer payoff model, or considering probabilistic cheating. Also, ours are one-round protocols and we show useful trade-offs between the benefit of the master and the probability of accepting a wrong result.

A somewhat related work is [3] in which they face the problem of bootstrapping a P2P computing system, in the presence of rational peers. The goal is to incentive peers to join the system, for which they propose a scheme that mixes lottery psychology and multilevel marketing. In our setting, the master could use their scheme to recruit workers. We assume in this paper that enough workers are willing to participate in the computation.

Framework. We consider a distributed system consisting of a master processor that assigns a task to a set of workers to compute and return the task result. We assume that the master has the possibility of verifying whether the value returned by a worker is the correct result of the task. We also assume that verifying an answer is more efficient than computing the task [12] (e.g., *NP*-complete problems if $P \neq NP$), but the correct result of the computation might not be obtained. Therefore, the master by verifying does not necessarily obtain the correct answer (e.g., when all workers cheat). As in [3, 28], workers are assumed to be rational and seek to maximize their benefit, i.e., they are not destructively malicious. We note that this assumption can conceptually be justified by the work of Shneidman and Parks [26] where they reason on the connection of rational players—of algorithmic mechanism designs—and workers in realistic P2P systems. Furthermore, we do not consider non-intentional errors produced by hardware or software problems.

The general protocol used by master and workers is the following. The master process assigns the task to n workers. Each worker processor i cheats with probability $p_c^{(i)}$ and the master processor verifies the answers with some probability p_v . If the master processor verifies, it rewards the honest workers and penalizes the cheaters. If the master does not verify, it accepts the answer returned by the majority of workers. However, it does not penalize any worker given that the majority can be actually cheating. Instead, the master rewards workers according to one of the three following models. Either the master rewards the majority only (*Reward Model* \mathcal{R}_m), or the master rewards all workers independently of the returned value (*Reward Model* \mathcal{R}_a), or the master does not reward at all (*Reward Model* \mathcal{R}_\emptyset).

Our model comprises the following form of collusion. Workers decide whether to cheat independently, but all cheaters collude in returning the same incorrect answer. Since the master accepts the majority, this behavior maximizes the chances of cheating the master. Being this the worst case, it subsumes models where cheaters do not necessarily return the same answer. We also assume that if a worker does not perform the task, then it is (almost) impossible to guess the correct answer (i.e., the probability is negligible).

Given the protocol above, the game is defined by a set of parameters that include rewards to the workers that return the correct value, punishments to the workers that cheated (that is, returned the incorrect result and “got caught”). Hence, the game is played between the master and the workers, where the first wants to obtain the correct result with a desired probability, while obtaining a desired utility value (in expectation), and the workers decide whether to be honest or cheaters, depending on their expected utility gain or loss. In this paper, we design several games and

¹Cooperation among various workers concealed from the master.

study the conditions under which unique Nash equilibria are achieved. Each NE results in a different benefit for the master and a different probability of accepting an incorrect result. Thus, the master can choose some game conditions so that a unique NE that best fits its goals is achieved.

Contribution. The main contributions of this paper are:

1. The identification of a collection of realistic payoff parameters that allow to model Internet-based master-worker computational environments in game theoretic terms.
2. The definition of four different games that the master can force to be played: (a) A game between the master and a single worker, (b) a game between the master and a worker, played n times (with different workers), (c) a game with a master and n workers, and (d) a game of n workers in which the master participates indirectly. Together with the three reward models defined above, we have overall defined twelve games among which the master can choose the most convenient to use in each specific context.
3. The analyses of all the games under general payoff models, and the characterization of conditions under which a unique Nash Equilibrium point is reached for each game and each payoff-model. These analyses lead to mechanisms that the master can run to trade cost and reliability.
4. To demonstrate the utility of the analysis, we design mechanisms for two specific realistic scenarios. These scenarios reflect, in their fundamental elements, (a) a system of volunteering computing like SETI, and (b) a company that buys computing cycles from Internet computers and sells them to its customers in the form of a task-computation service. Our results show that for (a) the best choice is non-redundant allocation, even for our weak model of collusion. Furthermore, in this case we show that to obtain always the correct answer it is enough to verify with arbitrarily small probability. As an example of the results obtained in (b), if the master only chooses the number of workers n , we show that, again, the best choice is non-redundant allocation. However, in order to achieve correctness always, the required probability of verifying can now be large.

Paper Structure. In Section 2 we provide basic definitions to be used throughout the paper. In Section 3 we present and analyze the games proposed. In Section 4 the mechanisms for the two realistic scenarios are designed.

2 Definitions

Game Definition. Game participants are referred as workers and master, or simply as *players*, indistinctively. In order to define the game played in each case, we follow the customary notation used in game theory. Given that this notation is repeatedly used throughout the paper,

WP_C	worker's punishment for being caught cheating
WC_T	worker's cost for computing the task
WB_A	worker's benefit from master's acceptance
MP_W	master's punishment for accepting a wrong answer
MC_A	master's cost for accepting the worker's answer
MC_V	master's cost for verifying worker's answers
MB_R	master's benefit from accepting the right answer

Table 1. Payoffs

we summarize it in Table 2 for clarity. We assume that the master always chooses an odd number of worker players n . Whenever needed, we will express a strategy profile as a group of individual strategy choices and two sets of workers, grouped by their strategy choice. For instance, $(s_i = C, s_M = V, F, T)$ means a strategy profile s where worker i chooses strategy C , the master chooses strategy V , a set F of $|F|$ workers choose strategy C , and a set T of $|T|$ workers choose strategy \bar{C} . For games with one worker and the master, the strategy profile is composed only by their choices. For example, m_{CV} stands for the master's payoff in the case that the worker cheated and the master verified. We require that, for each worker i , $p_C^{(i)} = 1 - p_{\bar{C}}^{(i)}$ and, for the master, $p_V = 1 - p_{\bar{V}}$. For games where we only have one worker or all workers use the same probability, we will express $p_C^{(i)}$ (resp. $p_{\bar{C}}^{(i)}$) simply by p_C (resp. $p_{\bar{C}}$). Whenever the strategy is clear from the context, we will refer to the expected utility of a worker as U_i , and for the master as U_M . Unless otherwise stated, the games studied are games of complete information, i.e., that the worker processors and the master know the algorithm and the parameters involved.

Equilibrium Definition. We define now precisely the conditions for the equilibrium among players. For any finite game, a mixed strategy profile σ^* is a *mixed-strategy Nash equilibrium* (MSNE) if, and only if, for each player i

$$\begin{aligned}
 U_i(s_i, \sigma_{-i}^*) &= U_i(s_j, \sigma_{-i}^*), \forall s_i, s_j \in \text{supp}(\sigma_i^*), & (1) \\
 U_i(s_i, \sigma_{-i}^*) &\geq U_i(s_k, \sigma_{-i}^*), \\
 \forall s_i, s_k &: s_i \in \text{supp}(\sigma_i^*), s_k \notin \text{supp}(\sigma_i^*). & (2)
 \end{aligned}$$

In words, given a MSNE with mixed-strategy profile σ^* , for each player i , the expected utility, assuming that all other players do not change their choice, is the same for each pure strategy that the player can choose with positive probability in σ^* , and it is not less than the expected utility of any pure strategy with probability zero of being chosen in σ^* . A *fully* MSNE is an equilibrium with mixed strategy profile σ where, for each player i , $\text{supp}(\sigma_i) = \mathcal{S}_i$.

Payoffs Definition. We detail in Table 1 the payoff definitions that will be used throughout the paper. All the parameters in this table are non-negative. Notice that we split the

$W = \{1, 2, \dots, n\}$	set of assigned workers
M	master processor
\mathcal{S}_i	set of pure strategies available to player i
$\{\mathcal{C}, \bar{\mathcal{C}}\}$	set of pure strategies of a worker
$\{\mathcal{V}, \bar{\mathcal{V}}\}$	set of pure strategies of the master
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
$w_s^{(i)}$	payoff of worker i for the strategy profile s
m_s	payoff of the master for the strategy profile s
$p_{s_i}^{(i)}$	probability that worker i uses strategy s_i
p_{s_M}	probability that the master uses strategy s_M
σ	mixed strategy profile (a mapping from players to prob. distrib. over pure strategies)
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 worker i with mixed strategy profile σ
$U_M(s_M, \sigma_{-M})$	expected utility of master with mixed strategy profile σ
$supp(\sigma_i)$	set of strategies of player i with probability > 0 in σ

Table 2. Game notation

reward to a worker in WB_A and MC_A to model the fact that the cost of the master might be different than the benefit of a worker. In fact, in some models they may be completely unrelated. Among the parameters involved, we assume that the master has the freedom of choosing the cheater penalty WP_C and the worker reward for computing MC_A . By tuning these parameters and choosing n , the master achieves the desired trade-off between correctness and cost.

3 Equilibria Analysis

In the following sections, different games are studied according with the participants involved. In order to identify the parameter conditions for which there is a NE, Equations 1 and 2 of the MSNE definition are instantiated in each particular game, without making any assumptions on the payoffs. We call this the general payoffs model. From these instantiations, we obtain conditions on the parameters (payoffs and probabilities) that would make such equilibrium unique. Finally, we introduce the reward models described before on those conditions, so that we can compare among all games and models in Section 4.

3.1 Game 1:1: One Master - One Worker

We start the analysis by considering the game between the master and one worker.

General Payoffs Model. In order to evaluate all possible equilibria, all the different mixes have to be considered. In other words, according with the range of values

that p_C and p_V can take, we can have fully MSNE, partially MSNE, or pure-strategies NE. More specifically, both p_C and p_V can take values either 0, 1 or in the open interval $(0, 1)$. Depending on these values the different conditions in Equations 1 and 2 have to be achieved in order to have an equilibrium. Hence, conditions on p_C and p_V for each equilibrium can be obtained from these equations (for details, see [8]). On the other hand, the expected utility of the master and the worker in any equilibrium are $U_M = p_C p_V m_{C_V} + (1 - p_C) p_V m_{\bar{C}_V} + p_C (1 - p_V) m_{C_{\bar{V}}} + (1 - p_C) (1 - p_V) m_{\bar{C}_{\bar{V}}}$ and $U_W = p_C p_V w_{C_V} + p_C (1 - p_V) w_{C_{\bar{V}}} + (1 - p_C) p_V w_{\bar{C}_V} + (1 - p_C) (1 - p_V) w_{\bar{C}_{\bar{V}}}$ respectively, and the probability of accepting the wrong answer can be obtained as $P_{wrong} = (1 - p_V) p_C$.

Reward Model \mathcal{R}_m . Recall that in this model we assume that when the master does not verify, rewards only the majority. Given that there is only one worker, in this case the master rewards always. Under the payoff model detailed in Table 1, the payoffs are

$$\begin{aligned}
m_{C_V} &= -MC_V & w_{C_V} &= -WP_C \\
m_{\bar{C}_V} &= MB_R - MC_V - MC_A & w_{\bar{C}_V} &= WB_A - WC_T \\
m_{C_{\bar{V}}} &= -MP_W - MC_A & w_{C_{\bar{V}}} &= WB_A \\
m_{\bar{C}_{\bar{V}}} &= MB_R - MC_A & w_{\bar{C}_{\bar{V}}} &= WB_A - WC_T
\end{aligned}$$

Replacing appropriately, we obtain the conditions for equilibrium, probability of accepting the wrong answer and utilities for each case.

Reward Model \mathcal{R}_a . In this model we assume that if the master does not verify, it rewards all workers independently of the answer. Hence, the analysis is identical to the previous section.

Reward Model \mathcal{R}_\emptyset . Recall that in this model we assume that if the master does not verify, it does not reward the worker. Hence, under the payoff model detailed in Table 1, the payoffs are

$$\begin{aligned} m_{C\mathcal{V}} &= -MC_{\mathcal{V}} & w_{C\mathcal{V}} &= -WP_C \\ m_{\bar{C}\mathcal{V}} &= MB_{\mathcal{R}} - MC_{\mathcal{V}} - MC_{\mathcal{A}} & w_{\bar{C}\mathcal{V}} &= WB_{\mathcal{A}} - WC_T \\ m_{C\bar{\mathcal{V}}} &= -MP_{\mathcal{W}} & w_{C\bar{\mathcal{V}}} &= 0 \\ m_{\bar{C}\bar{\mathcal{V}}} &= MB_{\mathcal{R}} & w_{\bar{C}\bar{\mathcal{V}}} &= -WC_T \end{aligned}$$

Replacing appropriately, we obtain the conditions for equilibrium, probability of accepting the wrong answer and utilities for each case, as we will see in the next Section.

3.2 Game 1:1ⁿ: n Games One to One

Given the equilibria computed in Section 3.1, the master runs n instances of that game, one with each of the n workers, choosing to verify or not with probability $p_{\mathcal{V}}$ only once. Additionally, when paying while not verifying, the master rewards all or none according with the one-to-one game.

General Payoffs Model. Since in this game the master runs n instances of the same one-to-one game, under the payoff model detailed in Table 1, the conditions for equilibria and the utility of a worker are the same as in Section 3.1. Although the expected utility of the master and the probability of accepting the wrong result change. In order to give those expressions, we define the following notation. Let \mathcal{W} be the set of partitions in two subsets (F, T) of W , i.e., $\mathcal{W} = \{(F, T) | F \cap T = \emptyset, F \cup T = W\}$. F is the set of workers that cheat and T the set of honest workers. We also define master payoff functions $m_s : \{0, 1, \dots, n\} \rightarrow \mathbb{R}$, that reflect the fact that the cost may include some fixed amount for unique verification or unique cost of being wrong but they are still function of the number of workers that cheat or not. For the sake of clarity, we will denote the probability that the majority cheats as \mathbf{P}_C . Then, the probability that the majority cheats, the probability of being wrong, and the master's utility are $\mathbf{P}_C = \sum_{(F, T) \in \mathcal{W}} \prod_{j \in F} p_C^{(j)} \prod_{k \in T} (1 - p_C^{(k)})$, $\mathbf{P}_{wrong} = (1 - \frac{|F| > |T|}{|F| + |T|})$, $p_{\mathcal{V}} \mathbf{P}_C$ and $U_M = p_{\mathcal{V}} \sum_{(F, T) \in \mathcal{W}} \prod_{j \in F} p_C^{(j)} \prod_{k \in T} (1 - p_C^{(k)}) m_{\mathcal{V}} + (1 - p_{\mathcal{V}}) \sum_{(F, T) \in \mathcal{W}} \prod_{j \in F} p_C^{(j)} \prod_{k \in T} (1 - p_C^{(k)}) m_{\bar{\mathcal{V}}}$ respectively, where $m_{\mathcal{V}} = m_{C\mathcal{V}}(|F|) + m_{\bar{C}\mathcal{V}}(|T|)$ and $m_{\bar{\mathcal{V}}} = m_{C\bar{\mathcal{V}}}(|F|) + m_{\bar{C}\bar{\mathcal{V}}}(|T|)$.

Reward Models. In this game, we assume that the cost of verification $MC_{\mathcal{V}}$ is independent of the number of workers and that, as long as some worker is honest, upon verification the master obtains the correct result. It is important to note that, under this assumption, the probability of obtaining the correct result is not $1 - \mathbf{P}_{wrong}$, given that if the master

verifies but all workers cheat, the master does not obtain the correct result. Recall that the master plays n instances of a one-to-one game, thus, depending on the model, it must reward every worker if not verifying independently of majorities. We summarize the probability of accepting the wrong result, the master utility for each case, the conditions for equilibrium, and the workers utility for the reward models \mathcal{R}_m and \mathcal{R}_\emptyset in Tables 3 and 4 respectively (Tables 3 and 4 give also these values for Game 1:1 replacing appropriately $n = 1$).

3.3 Game 0:n: No Master in the Game

Another natural generalization of the game of Section 3.1 is to consider a game in which the master assigns the task to n workers that play the game among them. Intuitively, it can be seen that workers will compete to be in the majority to persuade the master in case of not verifying. The question of how the master participating also in the game would affect the results obtained in this section is addressed in Section 3.4.

General Payoffs Model. In order to analyze this game, it is convenient to partition the set of workers W in three disjunct sets F, T, R , such that $F \cup T \cup R = W$ as follows. F is the set of workers that choose to cheat as a pure strategy, i.e., $F = \{i | i \in W \wedge p_C^{(i)} = 1\}$. T is the set of workers that choose not to cheat as a pure strategy, i.e., $T = \{i | i \in W \wedge p_C^{(i)} = 0\}$. R is the set of workers that randomize their choice, i.e., $R = \{i | i \in W \wedge p_C^{(i)} \in (0, 1)\}$. Let \mathcal{R}_i be the set of partitions in two subsets (R_F, R_T) of $R \setminus \{i\}$, i.e., $\mathcal{R}_i = \{(R_F, R_T) | R_F \cap R_T = \emptyset \wedge R_F \cup R_T = R \setminus \{i\}\}$. Let $\mathbf{E}[w_s^{(i)}]$ be the expected payoff of worker i for the strategy profile s , taking the expectation over the pure strategies of the master. More precisely, $\mathbf{E}[w_s^{(i)}] = p_{\mathcal{V}} w_{s_{M=\mathcal{V}}}^{(i)} + (1 - p_{\mathcal{V}}) w_{s_{M=\bar{\mathcal{V}}}}^{(i)}$. Then, for each worker $i \in W$, we have

$$U_i(R, F, T, s_i = C) = \sum_{(R_F, R_T) \in \mathcal{R}_i} \prod_{j \in R_F} p_C^{(j)} \prod_{k \in R_T} (1 - p_C^{(k)}) \mathbf{E}[w_{\substack{F \cup R_F, \\ T \cup R_T, \\ s_i = C}}^{(i)}], \quad (3)$$

$$U_i(R, F, T, s_i = \bar{C}) = \sum_{(R_F, R_T) \in \mathcal{R}_i} \prod_{j \in C} p_C^{(j)} \prod_{k \in R_T} (1 - p_C^{(k)}) \mathbf{E}[w_{\substack{F \cup R_F, \\ T \cup R_T, \\ s_i = \bar{C}}}^{(i)}]. \quad (4)$$

Given that, the normal strategic form of each worker is the same, facing the same scenario, every worker obtains the same payoff. In other words, if a given worker choose to cheat (resp. be honest), receives some payoff, and it

turns out that the majority of nodes cheated (resp. was honest), another worker that also cheats (resp. is honest) receives the same payoff. Additionally, all even splits of the other workers in cheaters/not-cheaters are the same from the point of view of i because what matters is the majority. If this is true for every worker, we fold these cases as follows. For each worker i and each partition $(R_F, R_T) \in \mathcal{R}_i$, $\Delta w_C = \mathbf{E}[w_{s_i=C}^{(i)}] - \mathbf{E}[w_{s_i=\bar{C}}^{(i)}]$, when $|F \cup R_F| > |T \cup R_T|$; $\Delta w_{\bar{C}} = \mathbf{E}[w_{s_i=C}^{(i)}] - \mathbf{E}[w_{s_i=\bar{C}}^{(i)}]$, when $|F \cup R_F| < |T \cup R_T|$; and $\Delta w_X = \mathbf{E}[w_{s_i=C}^{(i)}] - \mathbf{E}[w_{s_i=\bar{C}}^{(i)}]$, when $|F \cup R_F| = |T \cup R_T|$. Also, for convenience, we define $\Delta U_i(R, F, T) = U_i(R, F, T, s_i = C) - U_i(R, F, T, s_i = \bar{C})$. Replacing Equations 3 and 4 in this equation we obtain

$$\begin{aligned} \Delta U_i(R, F, T) = & \\ \Delta w_C \sum_{\substack{(R_F, R_T) \in \mathcal{R}_i \\ |F \cup R_F| > |T \cup R_T|}} & \prod_{j \in R_F} p_C^{(j)} \prod_{k \in R_T} (1 - p_C^{(k)}) + \\ \Delta w_X \sum_{\substack{(R_F, R_T) \in \mathcal{R}_i \\ |F \cup R_F| = |T \cup R_T|}} & \prod_{j \in R_F} p_C^{(j)} \prod_{k \in R_T} (1 - p_C^{(k)}) + \\ \Delta w_{\bar{C}} \sum_{\substack{(R_F, R_T) \in \mathcal{R}_i \\ |F \cup R_F| < |T \cup R_T|}} & \prod_{j \in R_F} p_C^{(j)} \prod_{k \in T} (1 - p_C^{(k)}). \quad (5) \end{aligned}$$

Restating the equilibrium conditions of Equations 1 or 2 in terms of Equation 5, for each worker $i \in R$ that does not choose a pure strategy, the equilibrium condition is $\Delta U_i(R, F, T) = 0$; for each worker $i \in F$, i.e., that chooses to cheat as a pure strategy, the condition is $\Delta U_i(R, F, T) \geq 0$; and for each $i \in T$, it must hold that $\Delta U_i(R, F, T) \leq 0$. In the following sections we show conditions to obtain an equilibrium for each reward model. The following lemma will be useful to show uniqueness.

Lemma 1. *Given a game as defined in this section, if $\Delta w_C \geq \Delta w_X \geq \Delta w_{\bar{C}}$, there is no unique equilibrium where $R \neq \emptyset$.*

Proof. For the sake of contradiction, assume there is a unique equilibrium σ for which $R \neq \emptyset$ and $\Delta w_C \geq \Delta w_X \geq \Delta w_{\bar{C}}$. Then, $\Delta U_i(R, F, T) = 0$ must be solvable for every worker $i \in R$. Given that $R \neq \emptyset$ the probabilities given by the summations in $\Delta U_i(R, F, T) = 0$ are all strictly bigger than zero. Therefore, given that the equation is solvable, either $\Delta w_C = \Delta w_X = \Delta w_{\bar{C}} = 0$, or some of these values have different signs. If $\Delta w_C \geq 0$, there would be also an equilibrium when all workers choose to cheat and σ would not be unique. So, it must hold that $\Delta w_C < 0$. Then, given that the equation is solvable, either $\Delta w_X \geq 0$ or $\Delta w_{\bar{C}} \geq 0$ or both, which is a contradiction. \square

In the following sections, we study conditions to obtain unique equilibria under different models. In all these models it holds that $\Delta w_C \geq \Delta w_X \geq \Delta w_{\bar{C}}$. Then, by Lemma 1, there is no unique equilibrium where $R \neq \emptyset$. Regarding

equilibria where $R = \emptyset$, unless the task assigned has a binary output, a unique equilibrium where all workers choose to cheat is not useful. Then, we make $\Delta w_C < 0$, $\Delta w_X < 0$ and $\Delta w_{\bar{C}} < 0$ so that $\Delta U_i(R, F, T) \geq 0$ has no solution and no worker can choose to cheat as a pure strategy. Thus, the only equilibrium is for all the workers to choose to be honest, which solves $\Delta U_i(R, F, T) \leq 0$. Therefore, $\forall i, p_C^{(i)} = 0$ and hence $\mathbf{P}_{wrong} = 0$.

Reward Model \mathcal{R}_m . Replacing appropriately the payoffs detailed in Table 1, we obtain $\Delta w_C = p_V(-WP_C - 2WB_A) + WB_A + WC_T$, $\Delta w_X = p_V(-WP_C - WB_A) + WC_T$, and $\Delta w_{\bar{C}} = p_V(-WP_C) - WB_A + WC_T$. The condition on p_V to obtain the aforementioned unique equilibrium is then $\Delta w_C < 0$, yielding $p_V > (WB_A + WC_T)/(WP_C + 2WB_A)$. The expected utilities are then $U_M = MB_R - p_V MC_V - nMC_A$ and $U_W = WB_A - WC_T$.

Reward Model \mathcal{R}_a . Again, replacing appropriately the payoffs detailed in Table 1, we have $\Delta w_C = p_V(-WP_C - WB_A) + WC_T$, $\Delta w_X = p_V(-WP_C - WB_A) + WC_T$ and $\Delta w_{\bar{C}} = p_V(-WP_C - WB_A) + WC_T$. Then, the condition to obtain the unique equilibrium and the expected utilities are $p_V > (WC_T)/(WP_C + WB_A)$, $U_M = MB_R - p_V MC_V - nMC_A$ and $U_W = WB_A - WC_T$.

Reward Model \mathcal{R}_\emptyset . Replacing appropriately the payoffs detailed in Table 1, $\Delta w_C = p_V(-WP_C - WB_A) + WC_T$, $\Delta w_X = p_V(-WP_C - WB_A) + WC_T$ and $\Delta w_{\bar{C}} = p_V(-WP_C - WB_A) + WC_T$. The condition to obtain the unique equilibrium and the expected utilities are $p_V > (WC_T)/(WP_C + WB_A)$, $U_M = MB_R - p_V(MC_V + nMC_A)$ and $U_W = p_V WB_A - WC_T$.

3.4 Game 1:n: One Master - n Workers

We now observe how the conditions obtained in the previous game are modified if the master also participates as a player. The equilibria analysis regarding the workers follows the same lines as in Section 3.3. However, now Equations 1 and 2 have to be applied to the master also, as follows.

General Payoffs Model. Recall that R is the set of workers that randomize their choice. Let \mathcal{R} be the set of partitions in two subsets (R_F, R_T) of R , i.e., $\mathcal{R} = \{(R_F, R_T) | R_F \cap R_T = \emptyset \wedge R_F \cup R_T = R\}$. Then, for the master,

$$U_M(R, F, T, s_M = \mathcal{V}) = \sum_{(R_F, R_T) \in \mathcal{R}} \prod_{j \in R_F} p_C^{(j)} \prod_{k \in R_T} (1 - p_C^{(k)}) m_{\substack{F \cup R_F, \\ T \cup R_T, \\ s_M = \mathcal{V}}}$$

$$U_M(R, F, T, s_M = \bar{\mathcal{V}}) = \sum_{(R_F, R_T) \in \mathcal{R}} \prod_{j \in R_F} p_C^{(j)} \prod_{k \in R_T} (1 - p_C^{(k)}) m_{F \cup R_F, T \cup R_T, s_M = \bar{\mathcal{V}}}$$

From Equation 1, if $p_V \in (0, 1)$, the MSNE condition is $U_M(R, F, T, s_M = \mathcal{V}) = U_M(R, F, T, s_M = \bar{\mathcal{V}})$. From Equation 2, if $p_V = 0$ the condition is $U_M(R, F, T, s_M = \mathcal{V}) \leq U_M(R, F, T, s_M = \bar{\mathcal{V}})$, and if $p_V = 1$ the condition is $U_M(R, F, T, s_M = \mathcal{V}) \geq U_M(R, F, T, s_M = \bar{\mathcal{V}})$.

The MSNE conditions for the workers are the same as in Section 3.3. Hence, the conditions obtained for each of the reward models are the same. However, additional conditions are obtained from the master-utility conditions as follows. As in Section 3.3, the desired unique MSNE occurs when $p_C = 0$. Using that, in the master-utility conditions we get for the reward model \mathcal{R}_m that if $p_V < 1$, $MB_{\mathcal{R}} - MC_V - nMC_A = MB_{\mathcal{R}} - nMC_A$, and if $p_V = 1$, $MB_{\mathcal{R}} - MC_V - nMC_A \geq MB_{\mathcal{R}} - nMC_A$. Therefore, in any case it must hold $MC_V = 0$. For the reward model \mathcal{R}_a , the master-utility conditions give, if $p_V < 1$, $MB_{\mathcal{R}} - MC_V - nMC_A = MB_{\mathcal{R}} - nMC_A$ and if $p_V = 1$, $MB_{\mathcal{R}} - MC_V - nMC_A \geq MB_{\mathcal{R}} - nMC_A$. Therefore, again, $MC_V = 0$. Finally, for the reward model 3, the master-utility conditions give if $p_V < 1$, $MB_{\mathcal{R}} - MC_V - nMC_A = MB_{\mathcal{R}}$ and if $p_V = 1$, $MB_{\mathcal{R}} - MC_V - nMC_A \geq MB_{\mathcal{R}}$. Therefore, $MC_V = MC_A = 0$. Hence, to achieve the goal of forcing the workers to be honest, in this game, verifying must be free for the master.

4 Mechanism Design

In this section two realistic scenarios in which the master-worker model considered could be naturally applicable are proposed. For these scenarios, we determine appropriate games and parameters to be used by the master to maximize its benefit.

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). Hence, we assume that $WB_A > WC_T = 0$. The master incurs in a (possibly small) cost MC_A 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 verify the values returned by the workers, at a cost $MC_V > 0$. We also assume that the master obtains a benefit $MB_{\mathcal{R}} > MC_A$ if

it accepts the correct result of the task, and suffers a cost $MP_W > MC_V$ if it accepts an incorrect value.

Under these constraints, the equilibria for games 1:1 and 1:1ⁿ collapse to one single equilibrium point. Also, since game 1:n requires $MC_V = 0$ for the equilibrium to be unique, it cannot be used in this scenario. The different applicable cases are summarized in Table 5. In this table it can be observed that in games 1:1 and 1:1ⁿ the equilibrium is achieved with any value of p_C in an interval. The master has no way to force the specific value of p_C that a worker uses within the interval. And, in particular, it cannot force $p_C = 0$ (i.e., $\mathbf{P}_{wrong} = 0$). Additionally, looking at the master utility, all games have $U_M < MB_{\mathcal{R}}$. However, in game (0:n, \mathcal{R}_\emptyset) the master can make U_M arbitrarily close to $MB_{\mathcal{R}}$ by setting p_V arbitrarily small. (Notice that the utility of a worker will be arbitrarily small likewise, but given that workers are volunteering this is not a problem.) In conclusion, the game (0:n, \mathcal{R}_\emptyset) with $n = 1$ and very small p_V is the best choice in this scenario, since it satisfies $\mathbf{P}_{wrong} = 0$ and $U_M \approx MB_{\mathcal{R}}$.

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. In this case the company pays the users an amount $S = WB_A = MC_A$ for using their computing capabilities, and charges the consumers another amount $MB_{\mathcal{R}} > MC_A$ for the provided service. Since the users are not altruistic in this scenario, we assume that computing a task is not free for them (i.e., $WC_T > 0$), and they must have incentives to participate (i.e., $U_W > 0$). As in the previous case, we assume that the master verifies and has a cost for accepting a wrong value, such that $MP_W > MC_V > 0$. Again, under these assumptions, the equilibria for games 1:1 and 1:1ⁿ collapse to unique equilibria and game 1:n can not be used. The different cases are summarized in Table 6. Observe that there are cases in this table in which the worker has negative expected utility U_W . Given that in this case workers are not altruistic, they will not accept to participate in such a game. This fact immediately rules out games (1:1, \mathcal{R}_\emptyset) and (1:1ⁿ, \mathcal{R}_\emptyset). Similarly, this restriction forces the master to use a value of $p_V > WC_T / WB_A$ in game (0:n, \mathcal{R}_\emptyset). Finally, comparing games (0:n, \mathcal{R}_m) and (0:n, \mathcal{R}_a), it can be seen that the master would never choose the former, because the lower bound of p_V is smaller in the latter while the rest of expressions are the same, which leads to a larger master utility.

In this scenario, beyond choosing the game and number of workers as in the previous one, we assume that the master can also choose the reward WB_A to the workers for correctly computing the task, and the punishment WP_C if they are caught returning an incorrect value. All possible

Equilibrium p_C, p_V	Conditions	\mathbf{P}_{wrong}	U_M	U_W
$\frac{MC_V}{MC_A+MP_W}, \frac{WC_T}{WB_A+WP_C}$		$(1-p_V)\mathbf{P}_C$	$p_V((1-p_C^n)MB_{\mathcal{R}} - MC_V - (1-p_C)nMC_A) + (1-p_V)(MB_{\mathcal{R}}(1-\mathbf{P}_C) - MP_W\mathbf{P}_C - nMC_A)$	$WB_A - WC_T$
$0, \frac{WC_T}{WB_A+WP_C} \leq p_V < 1$ $0 < p_V$	$MC_V = 0$	0	$MB_{\mathcal{R}} - nMC_A$	$WB_A - WC_T$
$1, 0 < p_V \leq \frac{WC_T}{WB_A+WP_C}$ $p_V < 1$	$MC_V = MP_W + MC_A$	$1-p_V$	$-p_V MC_V - (1-p_V)(MP_W + nMC_A)$	$(1-p_V)WB_A - p_V WP_C$
$0 \leq p_C \leq \frac{MC_V}{MC_A+MP_W}, 0$ $p_C < 1$	$WC_T = 0$	\mathbf{P}_C	$MB_{\mathcal{R}}(1-\mathbf{P}_C) - MP_W\mathbf{P}_C - nMC_A$	WB_A
$\frac{MC_V}{MC_A+MP_W} \leq p_C < 1, 1$ $0 < p_C$	$WC_T = WB_A + WP_C$	0	$(1 - \prod_{j \in W} p_C^{(j)})MB_{\mathcal{R}} - MC_V - \sum_{(W_F, W_T) \in W} \prod_{j \in W_F} p_C^{(j)} \cdot \prod_{k \in W_T} (1 - p_C^{(k)}) W_T MC_A$	$-WP_C$
1, 1	$MC_V \leq MP_W + MC_A$ $WC_T \geq WB_A + WP_C$	0	$-MC_V$	$-WP_C$
0, 1	$MC_V = 0$ $WC_T \leq WB_A + WP_C$	0	$MB_{\mathcal{R}} - nMC_A$	$WB_A - WC_T$
1, 0	$MC_V \geq MP_W + MC_A$	1	$-MP_W - nMC_A$	WB_A

Table 3. Game 1:1ⁿ, Models \mathcal{R}_m and \mathcal{R}_a (and Game 1:1 for $n = 1$)

Equilibrium p_C, p_V	Conditions	\mathbf{P}_{wrong}	U_M	U_W
$\frac{MC_V+MC_A}{MC_A+MP_W}, \frac{WC_T}{WB_A+WP_C}$		$(1-p_V)\mathbf{P}_C$	$p_V((1-p_C^n)MB_{\mathcal{R}} - MC_V - (1-p_C)nMC_A) + (1-p_V)(MB_{\mathcal{R}}(1-\mathbf{P}_C) - MP_W\mathbf{P}_C)$	$-p_V WP_C$
$0, \frac{WC_T}{WB_A+WP_C} \leq p_V < 1$ $0 < p_V$	$MC_A = MC_V = 0$	0	$MB_{\mathcal{R}}$	$p_V WB_A - WC_T$
$1, 0 < p_V \leq \frac{WC_T}{WB_A+WP_C}$ $p_V < 1$	$MC_V = MP_W$	$1-p_V$	$-MC_V$	$-p_V WP_C$
$0 \leq p_C \leq \frac{MC_V+MC_A}{MC_A+MP_W}, 0$ $p_C < 1$	$WC_T = 0$	\mathbf{P}_C	$MB_{\mathcal{R}}(1-\mathbf{P}_C) - MP_W\mathbf{P}_C$	0
$\frac{MC_V+MC_A}{MC_A+MP_W} \leq p_C < 1, 1$ $0 < p_C$	$WC_T = WB_A + WP_C$	0	$(1 - \prod_{j \in W} p_C^{(j)})MB_{\mathcal{R}} - MC_V - \sum_{(W_F, W_T) \in W} \prod_{j \in W_F} p_C^{(j)} \cdot \prod_{k \in W_T} (1 - p_C^{(k)}) W_T MC_A$	$-WP_C$
1, 1	$MC_V \leq MP_W$ $WC_T \geq WB_A + WP_C$	0	$-MC_V$	$-WP_C$
0, 1	$MC_V = MC_A = 0$ $WC_T \leq WB_A + WP_C$	0	$MB_{\mathcal{R}}$	$WB_A - WC_T$
1, 0	$MC_V \geq MP_W$	1	$-MP_W$	0

Table 4. Game 1:1ⁿ, Model \mathcal{R}_θ (and Game 1:1 for $n = 1$)

combined variations of these parameters yield a huge number of cases to be considered. In this work, we assume that the master only can choose one of these parameters, while the rest are predefined. A study of richer combinations is left for future work.

The following notation is used for clarity. Whenever a parameter may be different among different games being compared, a super-index indicates the game to which the parameter belongs. For instance, $U_M^{(i,j)}$ is the utility of the master for game (i, j) . MC_A and WB_A are referred to as simply S .

A simple observation of games $(0:n, \mathcal{R}_a)$ and $(0:n, \mathcal{R}_\theta)$ leads to find that in both cases it is convenient for the master to choose the smallest possible value of p_V . For this reason, in the following we assume in these games values $p_V^{(0:n, \mathcal{R}_a)} = \frac{WC_T}{C_P+S} + \gamma^{(0:n, \mathcal{R}_a)}$ and $p_V^{(0:n, \mathcal{R}_\theta)} =$

$\frac{WC_T}{S} + \gamma^{(0:n, \mathcal{R}_\theta)}$, for arbitrarily small $\gamma^{(0:n, \mathcal{R}_a)} > 0$ and $\gamma^{(0:n, \mathcal{R}_\theta)} > 0$.

Tunable n : Regarding games $(1:1, \mathcal{R}_m)$ and $(1:1^n, \mathcal{R}_m)$, in this case the master has no control over p_C or p_V , since they are completely defined by the application parameters. Hence, the probability of accepting a wrong answer might be arbitrarily close to 1, even for game $(1:1^n, \mathcal{R}_m)$, because \mathbf{P}_C grows with n if $p_C > 1/2$ as shown in Claim 2. Given that we want to design a mechanism that can be applied to any setting, we rule out these games for this case. In the case that n is tunable, the benefit of the master in games $(0:n, \mathcal{R}_a)$ and $(0:n, \mathcal{R}_\theta)$ decreases as n increases. Hence for these games the master chooses $n = 1$. Additionally, these games provide $\mathbf{P}_{wrong} = 0$. Out of these games, $(0:n, \mathcal{R}_a)$ is better iff $WC_T + WC_T MC_V/S > S + WC_T MC_V/(WP_C + S)$.

(Game,Model)	Equilibrium p_C, p_V	\mathbf{P}_{wrong}	U_M	U_W
(1:1, \mathcal{R}_m), (1:1, \mathcal{R}_a)	$0 \leq p_C \leq \frac{MC_V}{MC_A+MP_W}, p_C < 1, p_V = 0$	p_C	$MB_{\mathcal{R}} - p_C(MB_{\mathcal{R}} + MP_W) - MC_A$	WB_A
(1:1, \mathcal{R}_\emptyset)	$0 \leq p_C \leq \frac{MC_V+MC_A}{MC_A+MP_W}, p_C < 1, p_V = 0$	p_C	$MB_{\mathcal{R}} - p_C(MB_{\mathcal{R}} + MP_W)$	0
(1:1 ⁿ , \mathcal{R}_m), (1:1 ⁿ , \mathcal{R}_a)	$0 \leq p_C \leq \frac{MC_V}{MC_A+MP_W}, p_C < 1, p_V = 0$	\mathbf{P}_C	$MB_{\mathcal{R}} - \mathbf{P}_C(MB_{\mathcal{R}} + MP_W) - nMC_A$	WB_A
(1:1 ⁿ , \mathcal{R}_\emptyset)	$0 \leq p_C \leq \frac{MC_V+MC_A}{MC_A+MP_W}, p_C < 1, p_V = 0$	\mathbf{P}_C	$MB_{\mathcal{R}} - \mathbf{P}_C(MB_{\mathcal{R}} + MP_W)$	0
(0:n, \mathcal{R}_m)	$p_C = 0, \frac{WB_A}{WP_C+2WB_A} < p_V \leq 1$	0	$MB_{\mathcal{R}} - p_V MC_V - nMC_A$	WB_A
(0:n, \mathcal{R}_a)	$p_C = 0, 0 < p_V \leq 1$	0	$MB_{\mathcal{R}} - p_V MC_V - nMC_A$	WB_A
(0:n, \mathcal{R}_\emptyset)	$p_C = 0, 0 < p_V \leq 1$	0	$MB_{\mathcal{R}} - p_V(MC_V + nMC_A)$	$p_V WB_A$

Table 5. SETI-like Scenario

(Game,Model)	Equilibrium p_C, p_V	\mathbf{P}_{wrong}	U_M	U_W
(1:1, \mathcal{R}_m), (1:1, \mathcal{R}_a)	$\frac{MC_V}{MC_A+MP_W}, \frac{WC_T}{WB_A+WP_C}$	$(1-p_V)p_C$	$MB_{\mathcal{R}} - p_C(MB_{\mathcal{R}} + MP_W) - MC_A$	$WB_A - WC_T$
(1:1, \mathcal{R}_\emptyset)	$\frac{MC_V+MC_A}{MC_A+MP_W}, \frac{WC_T}{WB_A+WP_C}$	$(1-p_V)p_C$	$MB_{\mathcal{R}} - p_C(MB_{\mathcal{R}} + MP_W)$	$-p_V WP_C$
(1:1 ⁿ , \mathcal{R}_m), (1:1 ⁿ , \mathcal{R}_a)	$\frac{MC_V}{MC_A+MP_W}, \frac{WC_T}{WB_A+WP_C}$	$(1-p_V)\mathbf{P}_C$	$(p_V(1-p_C^n) + (1-p_V)(1-\mathbf{P}_C))MB_{\mathcal{R}} - p_V MC_V - (1-p_V)\mathbf{P}_C MP_W - (1-p_V p_C)nMC_A$	$WB_A - WC_T$
(1:1 ⁿ , \mathcal{R}_\emptyset)	$\frac{MC_V+MC_A}{MC_A+MP_W}, \frac{WC_T}{WB_A+WP_C}$	$(1-p_V)\mathbf{P}_C$	$(p_V(1-p_C^n) + (1-p_V)(1-\mathbf{P}_C))MB_{\mathcal{R}} - p_V MC_V - (1-p_V)\mathbf{P}_C MP_W - p_V(1-p_C)nMC_A$	$-p_V WP_C$
(0:n, \mathcal{R}_m)	$0, \frac{WB_A+WC_T}{WP_C+2WB_A} < p_V \leq 1$	0	$MB_{\mathcal{R}} - p_V MC_V - nMC_A$	$WB_A - WC_T$
(0:n, \mathcal{R}_a)	$0, \frac{WC_T}{WP_C+WB_A} < p_V \leq 1$	0	$MB_{\mathcal{R}} - p_V MC_V - nMC_A$	$WB_A - WC_T$
(0:n, \mathcal{R}_\emptyset)	$0, \frac{WC_T}{WP_C+WB_A} < p_V \leq 1$	0	$MB_{\mathcal{R}} - p_V(MC_V + nMC_A)$	$p_V WB_A - WC_T$

Table 6. Contractor Scenario

Tunable WP_C : Comparing games (0:n, \mathcal{R}_a) and (0:n, \mathcal{R}_\emptyset), $U_M^{(0:n,\mathcal{R}_a)} = MB_{\mathcal{R}} - p_V^{(0:n,\mathcal{R}_a)} MC_V - nS = MB_{\mathcal{R}} - WC_T MC_V / (S + WP_C^{(0:n,\mathcal{R}_a)}) - nS - \gamma^{(0:n,\mathcal{R}_a)} MC_V$ and $U_M^{(0:n,\mathcal{R}_\emptyset)} = MB_{\mathcal{R}} - p_V^{(0:n,\mathcal{R}_\emptyset)} MC_V - p_V^{(0:n,\mathcal{R}_\emptyset)} nS = MB_{\mathcal{R}} - WC_T MC_V / S - nWC_T - \gamma^{(0:n,\mathcal{R}_\emptyset)} MC_V - \gamma^{(0:n,\mathcal{R}_\emptyset)} nS$. Thus, game (0:n, \mathcal{R}_\emptyset) is better iff $n > WC_T MC_V / (S(S - WC_T))$ for small enough $\gamma^{(0:n,\mathcal{R}_\emptyset)}$. Otherwise, (0:n, \mathcal{R}_a) is better for small enough $\gamma^{(0:n,\mathcal{R}_a)}$ and large enough $WP_C^{(0:n,\mathcal{R}_a)}$. As argued in the previous case, in this case the master has no control over p_C . Although the master can reduce WP_C to increase p_V , it can not make p_V arbitrarily close to 1 to reduce \mathbf{P}_{wrong} in case p_C is big (and consequently \mathbf{P}_C). Then, some cases might lead to a big probability of accepting the wrong answer. Thus, games (1:1, \mathcal{R}_m) and (1:1ⁿ, \mathcal{R}_m) are ruled out from consideration.

Tunable S in $(WC_T, MB_{\mathcal{R}})$: In this case n is fixed, and given that we do not make any assumptions about its magnitude, we evaluate game 1 : 1 while evaluating game 1 : 1ⁿ for an arbitrary n . Using calculus, the utility of the master for game (0 : n, \mathcal{R}_a) is maximum when $S_{\max}^{(0:n,\mathcal{R}_a)} = \pm \sqrt{MC_V WC_T / n} - WP_C$. Due to the aforementioned constraints, only values in the interval $(WC_T, MB_{\mathcal{R}})$ are valid for S . Assuming then that $WC_T < S_{\max}^{(0:n,\mathcal{R}_a)} < MB_{\mathcal{R}}$, the utilities are $U_M^{(0:n,\mathcal{R}_a)}(S = S_{\max}^{(0:n,\mathcal{R}_a)}) = MB_{\mathcal{R}} - 2\sqrt{nMC_V WC_T} + nWP_C$ and $U_M^{(0:n,\mathcal{R}_\emptyset)} = MB_{\mathcal{R}} - WC_T MC_V / S^{(0:n,\mathcal{R}_\emptyset)} - nWC_T - \gamma^{(0:n,\mathcal{R}_\emptyset)}(MC_V + nS^{(0:n,\mathcal{R}_\emptyset)})$. Since $U_M^{(1:1^n,\mathcal{R}_m)} \leq MB_{\mathcal{R}}$,

game (0 : n, \mathcal{R}_a) is better than game (1 : 1ⁿ, \mathcal{R}_m) whenever $n > 4MC_V WC_T / WP_C^2$. On the other hand, game (0 : n, \mathcal{R}_\emptyset) is better than game (0 : n, \mathcal{R}_a) if $MB_{\mathcal{R}} > WC_T MC_V / (2\sqrt{nMC_V WC_T} - n(WP_C + WC_T))$, for small enough $\gamma^{(0:n,\mathcal{R}_\emptyset)}$ and $S^{(0:n,\mathcal{R}_\emptyset)}$ arbitrarily close to $MB_{\mathcal{R}}$. In order to show a scenario where game (1:1ⁿ, \mathcal{R}_m) is better, we assume now that $MP_W \geq 2MC_V$. Then, under this assumption, $p_C \leq 1/2$. The following claim that makes use of this fact will be useful.

Claim 2. For game 1:1ⁿ, if $\mathbf{P}_C(n)$ is the probability that the majority out of n workers cheat, then, if the probability that a worker cheats $p_C \leq \frac{1}{2}$, $\mathbf{P}_C(n+2) \leq \mathbf{P}_C(n)$.

Proof. Let $\mathbf{P}_C(n, > 1)$ be the probability that, out of n workers, the number of cheaters exceed the number of honest workers by more than one (i.e., at least 3 given that we consider only odd number of workers), $\mathbf{P}_C(n, = 1)$ by exactly one, and $\mathbf{P}_{\bar{C}}(n, = 1)$ be the probability that the number of honest workers exceed the number of cheaters by exactly one. Then, $\mathbf{P}_C(n+2) = \mathbf{P}_C(n, > 1)(p_C^2 + (1-p_C)^2) + \mathbf{P}_C(n, = 1)(p_C^2 + 2p_C(1-p_C)) + \mathbf{P}_{\bar{C}}(n, = 1)p_C^2$. Bounding p_C the claim follows. \square

From the previous claim, given that $\mathbf{P}_C = 1/2$ for $p_C = 1/2$, we conclude that $\mathbf{P}_C \leq 1/2$. Using that $p_C \leq 1/2$, $\mathbf{P}_C \leq 1/2$, and $MP_W > 2MC_V$, the utility of the master

for game $(1:1^n, \mathcal{R}_m)$ is

$$\begin{aligned}
U_M^{(1:1^n, \mathcal{R}_m)} &\geq \frac{1}{2} MB_{\mathcal{R}} - p_V^{(1:1^n, \mathcal{R}_m)} MC_V - \\
&\quad \frac{1}{2} (1 - p_V^{(1:1^n, \mathcal{R}_m)}) MP_W - nS^{(1:1^n, \mathcal{R}_m)} \\
&= \frac{1}{2} MB_{\mathcal{R}} - p_V^{(1:1^n, \mathcal{R}_m)} MC_V - \\
&\quad \frac{1}{2} MP_W + \frac{1}{2} p_V^{(1:1^n, \mathcal{R}_m)} MP_W - nS^{(1:1^n, \mathcal{R}_m)} \\
&\geq \frac{1}{2} (MB_{\mathcal{R}} - MP_W) - nS^{(1:1^n, \mathcal{R}_m)}.
\end{aligned}$$

As shown before, game $(0:n, \mathcal{R}_a)$ is better than game $(0:n, \mathcal{R}_\emptyset)$ when $MB_{\mathcal{R}} < WC_T MC_V / (2\sqrt{nMC_V WC_T} - n(WP_C + WC_T))$. Comparing games $(1:1^n, \mathcal{R}_m)$ and $(0:n, \mathcal{R}_a)$ when $WC_T < \sqrt{MC_V WC_T / n} - WP_C < MB_{\mathcal{R}}$, we have $(MB_{\mathcal{R}} - MP_W) / 2 - nS^{(1:1^n, \mathcal{R}_m)} \geq MB_{\mathcal{R}} - 2\sqrt{nMC_V WC_T} + nWP_C$. Therefore, game $(1:1^n, \mathcal{R}_m)$ is better whenever

$$\begin{aligned}
WC_T \leq S^{(1:1^n, \mathcal{R}_m)} &\leq 2\sqrt{\frac{MC_V WC_T}{n}} - \\
&\quad \frac{1}{2n} (MB_{\mathcal{R}} + MP_W) - WP_C \quad (6)
\end{aligned}$$

All three conditions are feasible simultaneously for big enough MC_V , therefore there exists a scenario for which game $(1:1^n, \mathcal{R}_m)$ is better. Notice that under the aforementioned condition, for game $(0:n, \mathcal{R}_a)$ to be better, i.e., $n > 4MC_V WC_T / WP_C^2$, it must be true that $WP_C > 2\sqrt{MC_V WC_T / n}$ and the inequality 6 does not hold.

References

- [1] 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, 2006.
- [2] D. Anderson. BOINC: A system for public-resource computing and storage. In *proc. of GRID 2004*, pp. 4–10, 2004.
- [3] J. R. Douceur and T. Moscibroda. Lottery trees: motivational deployment of networked systems. In *proc. of SIGCOMM 2007*, pp. 121–132, 2007.
- [4] J. R. Douceur. The Sybil attack. In *proc. of IPTPS 2002*, pp. 251–260, 2002.
- [5] “Enabling Grids for E-science”, <http://www.eu-egee.org>.
- [6] J. Feigenbaum, Ch. Papadimitriou, R. Sami, and S. Shenker. A BGP-based mechanism for lowest-cost routing. *Distributed Computing*, 18(1):61–72, 2005.
- [7] A. Fernández, C. Georgiou, L. Lopez, and A. Santos. Reliably executing tasks in the presence of untrusted processors. In *proc. of SRDS 2006*, pp. 39–50, 2006.
- [8] A. Fernández Anta, C. Georgiou, and M. A. Mosteiro. Cost-sensitive Mechanisms for Internet-based Computing. Technical Report UCY-TR-08-01, Dept. of Computer Science, University of Cyprus, March 2008.
- [9] 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, 2003.
- [10] D. Fotakis. Memoryless facility location in one pass. In *proc. of STACS 2006*, pp. 608–620, 2006.
- [11] P. Golle and I. Mironov. Uncheatable distributed computations. In *proc. of CT-RSA 2001*, pp. 425–440, 2001.
- [12] P. Golle and S. Stubblebine. Secure distributed computing in a commercial environment. In *proc. of Financial Cryptography*, 2001.
- [13] M. Halldorsson, J.Y. Halpern, L. Li, and V. Mirrokni. On spectrum sharing games. In *proc. of PODC 2004*, pp. 107–114, 2004.
- [14] J.Y. Halpern. Computer science and game theory: A brief survey. *Palgrave Dictionary of Economics*, to appear.
- [15] J.Y. Halpern and V. Teague. Rational secret sharing and multiparty computation. In *proc. of STOC 2004*, pp. 623–632, 2004.
- [16] K.M. Konwar, S. Rajasekaran, and A.A. Shvartsman. Robust network supercomputing with malicious processes. In *proc. of DISC 2006*, pp. 474–488, 2006.
- [17] 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.
- [18] E. Koutsoupias and Ch. Papadimitriou. Worst-case equilibria. In *proc. of STACS 1999*, pp. 404–413, 1999.
- [19] M. Mavronicolas and P. Spirakis. The price of selfish routing. *Algorithmica*, 48(1):91–126, 2007.
- [20] 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, 2006.
- [21] J.F. Nash. Equilibrium points in n -person games. *National Academy of Sciences*, 36(1):48–49, 1950.
- [22] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.
- [23] N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [24] A. Lopez-Ortiz. Algorithmic foundations of the Internet. ACM SIGACT News Distributed Computing Column 18, 2005.
- [25] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of ACM*, 49(2):236–259, 2002.
- [26] J. Shneidman and D.C. Parkes. Rationality and self-interest in P2P networks. In *proc. of IPTPS 2003*, pp. 139–148, 2003.
- [27] “TeraGrid”, <http://www.teragrid.org>.
- [28] 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, 2005.