



Universidad
Carlos III de Madrid



This document is published in:

Concurrency and Computation: Practice and Experience 25 (2013) 17,
pp.2363-2380
DOI:10.1002/cpe.3104

© 2013 John Wiley & Sons, Ltd.

Applying the dynamics of evolution to achieve reliability in master–worker computing

Evgenia Christoforou¹, Antonio Fernández Anta¹, Chryssis Georgiou^{2,*†},
Miguel A. Mosteiro^{3,4} and Angel Sánchez^{5,6}

¹*Institute IMDEA Networks, Madrid, Spain*

²*Department of Computer Science, University of Cyprus, Nicosia, Cyprus*

³*Department of Computer Science, Kean University, Union, NJ, USA*

⁴*GSyC, Universidad Rey Juan Carlos, Madrid, Spain*

⁵*GISC, Department of Mathematics, Universidad Carlos III de Madrid, Madrid, Spain*

⁶*BIFI Institute, Zaragoza, Spain*

SUMMARY

We consider Internet-based master–worker task computations, such as SETI@home, where a master process sends tasks, across the Internet, to worker processes; workers execute and report back some result. However, these workers are not trustworthy, and it might be at their best interest to report incorrect results. In such master–worker computations, the behavior and the best interest of the workers might change over time. We model such computations using evolutionary dynamics, and we study the conditions under which the master can reliably obtain task results. In particular, we develop and analyze an algorithmic mechanism based on reinforcement learning to provide workers with the necessary incentives to eventually become truthful. Our analysis identifies the conditions under which truthful behavior can be ensured and bounds the expected convergence time to that behavior. The analysis is complemented with illustrative simulations. Copyright © 2013 John Wiley & Sons, Ltd.

KEY WORDS: performing tasks; Internet-based computing; evolutionary dynamics; reinforcement learning; algorithmic mechanism design

1. INTRODUCTION

1.1. Motivation

As an alternative to expensive supercomputing parallel machines, Internet provides a feasible computational platform for processing complex computational jobs. Several Internet-based applications operate on top of this global computation infrastructure. Examples are volunteer-based ‘@home’ projects [1] such as SETI [2] and profit-seeking computation platforms such as Amazon’s Mechanical Turk [3].

Although the potential is great, the use of Internet-based computing is limited by the untrustworthy nature of the platform’s components [1, 4, 5]. In SETI for example, there is a machine, call it the *master*, that sends tasks, across the Internet, to volunteers’ computers, call them *workers*, that execute and report back some result. However, these workers may not be are not trustworthy and it might be at their best interest to report incorrect results; that is, workers (or their owners) can be viewed as *rational* [4, 6, 7]. In SETI, the master attempts to attenuate the impact of these bogus

*Correspondence to: Chryssis Georgiou, Department of Computer Science, University of Cyprus, Nicosia, Cyprus.

†E-mail: chryssis@cs.ucy.ac.cy

results by assigning the same task to several workers and comparing their outcomes (i.e., redundant task allocation is employed [1, 8]).

Prior work [9–11] has shown that it is possible to design algorithmic mechanisms with reward/punishment schemes so that the master can reliably obtain correct task results. We view these mechanisms as one-shot in the following sense: In a round, the master sends a task to be computed to a collection of workers, and the mechanism, using auditing and reward/punishment schemes, guarantees (with high probability) that the master obtains the correct task result. For another task to be computed, the process is repeated (with the same or different collection of workers) but without taking advantage of the knowledge gained.

Given a long running computation (such as SETI-like master–worker computations), it can be the case that the best interest, and hence the behavior of the workers, might change over time. The question then arises: Would it be possible to design a mechanism for performing many tasks, over the course of a possibly infinite computation, that could benefit from the repeated interaction between a master and the same collection of workers?

1.2. Our approach

In this work, we provide a positive answer to the earlier question. To do so, we introduce the concept of *evolutionary dynamics* (widely used under the biological and social perspective) and apply it to Internet-based master–worker task computing. More specifically, we employ *reinforcement learning* [12, 13] to model how system entities, or learners, interact with the environment to decide upon a strategy, and use their experience to select or avoid actions according to the consequences observed. Positive payoffs increase the likelihood of reusing the strategy just chosen, and negative payoffs reduce it. Payoffs are seen as parameterizations of players’ responses to their experiences. Empirical evidence [14, 15] suggests that reinforcement learning is more plausible with players that have information only on the payoffs they receive; that is, they do not have knowledge of the strategies involved. This model of learning fits nicely in our master–worker computation problem: each worker has no information about the master and the other workers’ strategies, and it does not know the set of strategies that led to the payoff it receives. The workers have only information about the strategies they choose at each round and their own received payoffs. The master also has minimal information about the workers and their intentions (to be truthful or not). Thus, we employ reinforcement learning for both the master and the workers in an attempt to build a reliable computational platform.

1.3. Our contributions

- To the best of our knowledge, this is the first work that studies the evolutionary dynamics of Internet-based master–worker task computing through reinforcement learning. We develop and analyze a mechanism based on reinforcement learning to be used by the master and the workers. In particular, in each round, the master allocates a task to the workers and decides whether to audit their responses with a certain probability p_A . Depending on whether it audits or not, it applies a different reward/punishment scheme, and adjusts the probability p_A for the next round (also known as the next task execution). Similarly, in a round, each worker i decides, with a certain probability p_{Ci} , whether it will report an incorrect result or it will truthfully compute and report the correct task result. Depending on the outcome of its decision, measured by the increase or the decrease of the worker’s *utility*, the worker adjusts its probability p_{Ci} for the next round.
- We show necessary and sufficient conditions under which the mechanism ensures *eventual correctness*. That is, we establish the conditions under which, after some finite number of rounds, the master obtains the correct task result in every round, with minimal auditing, while keeping the workers satisfied (with respect to their utility). Eventual correctness can be viewed as a form of *evolutionary stable strategy* [16, 17] as studied in Evolutionary Game Theory (EGT) [18]: even if a ‘mutant’ worker decides to change its strategy to cheating, it will soon be brought back to an honest strategy.
- Finally, we show that our mechanism, when adhering to the aforementioned conditions, reaches eventual correctness quickly. In particular, we show analytically probabilistic bounds on the

convergence time, as well as bounds on the expected convergence time. Our analysis is complemented with simulations, for a variety of parameter combinations likely to occur in practice.

1.4. Background and related work

An increasing number of works that apply game theory to distributed computing exists, including research on Internet routing, resource/facility location and sharing, containment of virus spreading, secret sharing, P2P services, and task computations. For more discussion on the connection between game theory and distributed computing, we refer the reader to the surveys by Halpern [19] and by Abraham *et al.* [20], and to the book by Nisan *et al.* [21].

The problem of achieving reliability in master-worker computations has recently been studied under two different views: from a ‘classical’ distributed computing view [8, 22, 23] and from a game-theoretic view [9, 11]. 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 always compute and return the correct task result. 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 are designed.

Under the classical game-theoretic view, workers act on their own *self-interest*, and they do not have an a priori established behavior (malicious or altruistic). They are assumed to be *rational* [4, 6]. In other words, the workers decide on whether they will be *honest* and report the correct task result, or *cheat* and report a bogus result, depending on which strategy increases their benefit or *utility*. Under this view, algorithmic mechanisms [6, 24, 25] are employed, where games are designed to provide 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) [26], that is, a strategy choice by each worker such that none of them has an incentive to change it. That NE is the one in which the master achieves a desired probability of obtaining the correct task result.

In a previous work by Fernández Anta *et al.* [9], a game-theoretic approach was used to achieve reliability in a similar Internet-based master-worker computation. A master processor assigns, across the Internet, a single computational task to a set of potentially untrusted worker processors and collects their responses. A weak form of collusion, as in this work, is assumed where all workers that cheat return the same incorrect result. Game-theoretic models for the problem were designed and analyzed, for example, one master and one worker, and several workers. For each of those models, a mechanism that achieves high reliability under the same payoff system that we use was designed. In [9], besides rewarding the majority of workers, as we do here, other reward models were considered, for example, rewarding all workers and rewarding none of the workers. The work in [9] was later extended [10] to also consider the possibility of workers being malicious.

In a work by Yurkewych *et al.* [11], a similar game-theoretic approach is studied. 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 collusion, and single-worker allocation. They conclude that, in their model, single-worker allocation is a cost-effective mechanism, specially in presence of collusion. As opposed to [9], it is possible that a decision is not reached in a round, in which case a different set of workers is selected, and the process is repeated (but without using any knowledge of the previous unsuccessful round).

The works in [9] and [11] construct a mechanism assuming a number of parameters are known to the master and the workers. In addition, the computation between master and workers is viewed as one-shot: each round of task executions is independent of the previous round. That is, their work does not consider the advantage of repeated communication with the same set of workers. Moreover it is expected that in a master-worker computation the behavior and the best interest of the workers

change over time. A dynamic environment like that could be analyzed using evolutionary dynamics, originally introduced in biology as a tool to study the mathematical principles according to which life is evolving [27]. Since then, a number of fields were inspired by the principles of evolutionary dynamics (e.g., sociology, and economics, artificial intelligence), and a variety of mechanisms was developed, aiming to accurately model the process of evolution. Our work is inspired by dynamics of evolution as a mean to model workers adaptation to a truthful behavior.

The dynamics of evolution have been studied under the principles of EGT. Maynard-Smith and Price [28, 29] introduced the concept of EGT in an effort to apply game theoretical ideas to understand evolving populations of lifeforms. This made Maynard-Smith [28] adjust the traditional concept of strategy, equilibrium, and the nature of a player’s interaction, so that a player would learn how to optimize its behavior and maximize its return. However, while in traditional game theory players choose a strategy from their strategy sets, EGT in biology is dealing with species inheriting possibly mutated strategies. When referring to social entities such as humans, evolution is understood as a learning process akin to ‘cultural evolution’ [30]. Cultural evolution implies an analogy between learning and biological evolution. For the existing analogies between learning, at the individual level, and biological evolution, we refer the reader to a paper by Borgers and Sarin [31].

In EGT, instead of the NE, Maynard-Smith, and Price used the evolutionarily stable strategy (ESS) concept [16–18, 32]. A strategy is called evolutionarily stable if, when the whole population is using this strategy, any group of invaders (mutants) using a different strategy will eventually die away over multiple generations (evolutionary rounds). All ESS are Nash equilibria, but the reverse is not true. Our work is driven by the concept of ESS, and we wish to have a similar stable strategy among workers that would guarantee reliability. Even if a mutant worker decides to change its strategy to cheating, it will be soon brought back to an honest strategy. Instead of the one-shot and repeated games of classical game theory, EGT assumes that the game is played repeatedly by players randomly drawn from large populations, uninformed of the preferences of opponents. In our work, we do not wish to change the set of players as evolutionary rounds progress, but we rather talk about ‘cultural evolution’, where workers change their strategies as a process of learning, rather than being replaced themselves.

While evolution operates on the global distribution of strategies within a given population, reinforcement learning [13] operates on the individual level of each member of the population. A well-known model of reinforcement learning is Bush and Mosteller’s model [12]. In this model, the players have limited information, and they play in discrete time repeatedly the same normal-form game. At each point in time, the players are characterized by a probability distribution over their strategy sets. Players’ choices are random, because they are affected by some unpredictable ‘psychological’ factor. This probability distribution is adjusted over time in response to experience. This experience is gained through repeated interactions of the players with the system, based on their strategies and the received payoffs. Positive payoffs reinforce the strategy just chosen, and negative payoffs discourage the use of that strategy.

Specifically, Bush and Mosteller’s model is an aspiration-based reinforcement learning model: Players adapt by comparing their experience with an *aspiration* level. There are several models of how aspirations are formed and adjusted over time, formally described in a study by Bendor *et al.* [14]. In the present paper, we use a simple model where aspiration is fixed by the workers and does not change during the evolutionary process (as in [33]). For more information on the different reinforcement learning models and comparisons between them, we refer the reader to a paper by Laslier *et al.* [34] and a study by Izquierdo and Izquierdo [35].

A survey by Phelps *et al.* [36] and an article by Conitzer and Sandholm [37] take a new approach on Mechanism Design by introducing the concept of evolutionary mechanism design. Evolutionary mechanism design assumes an engineering approach, based on an incremental process that creates a partly automated mechanism design. The evolutionary mechanism has a continuous interaction and feedback from the current mechanism, as opposed to classical mechanism design, which after the mechanism is introduced in the system, remains in the same NE forever. Looking at it from a different perspective, evolutionary mechanism design is analogous to EGT. Just as players may be forced to gradually adjust their strategies, in an analogous manner mechanisms are gradually

making adjustments in their rules with respect to what strategies are currently in play. In some way, our mechanism can be seen as an evolutionary mechanism, because the probability of auditing of the master and the probability of cheating of the workers change, which is similar to changing the mechanism.

Distributed computation in the presence of selfishness was studied within the scope of combinatorial agencies in Economics [38–41]. The basic model considered is a combinatorial variant of the classical principal-agent problem [42]: 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. The principal-agent approach deals in general with designing contracts between the principal and the workers that allow the principal to get the most out of the workers without knowing a priori what their actual capabilities are. One difference with respect to our master–worker framework is that, the worker’s actions cannot really be viewed as *hidden* in our setting. Another important difference is that 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 (either a worker truthfully performs a task or not).

In the work of Rose and Willemain [43], the principal-agent problem is extended to evolutionary learning, and bounded rationality of the agents is assumed. Players’ learning is simulated with a genetic algorithm that roughly mimics selection and mutations in biological evolution. Changes in the system are externally induced through the use of incentives. The agents’ learning is aided by the principal’s incentives that are used to adjust the learning, according to the output the principal desires. The principal is also able to use an artificial selection procedure to identify high performing agents for its own benefit.

Compared with the work of Rose and Willemain [43], in our line of work the learning model is different (in addition to the differences our work has with the principal-agent model). We assume that the learning procedure of the players remains the same through the evolutionary process. In contrast, in the more general model of Rose and Willemain, the learning procedure of the players may change over time and players can experience mutations. In both works, incentives are used to impose a desired behavior over time. But in [43], bounded rationality has been used, while in our work no cognitive limitation of the workers is assumed. As a future direction, our model could be enriched by including ‘biological learning’ through replicator dynamics. While in [43], the principal artificially selects the agents that increase performance, in our context the master could, by using a reputation technique, exclude from the computation low performing workers.

2. MODEL AND DEFINITIONS

2.1. Master–worker framework

We consider a distributed system consisting of a master processor that assigns, over the Internet, computational tasks to a set W of n workers (with out loss of generality, we assume that n is odd). In particular, the computation is broken into rounds. In each round, the master sends a task to be computed to the workers, and the workers return the task result. The master, based on the workers’ replies, must decide on the value it believes is the correct outcome of the task in the same round. The tasks considered in this work are assumed to have a unique solution; although such limitation reduces the scope of application of the presented mechanism [44], there are plenty of computations where the correct solution is unique: for example, any mathematical function. Note that in this work, we do not focus on any specific application. Therefore, applying our mechanism in practice may require to pay attention to details depending on the specific application considered. In this work, security issues are not considered. Security can be achieved by cryptographic means, as performed in BOINC [45], which allows for encrypting communication, authenticating master and workers, signing the code of tasks and executing tasks in sandboxes.

Following Abraham *et al.* [6] and Shneidman and Parkes [7], we assume that workers are *rational*, that is, they are selfish in a game-theoretic sense and their aim is to maximize their benefit (utility) under the assumption that other workers do the same. In the context of this paper, a worker is *honest* in a round when it truthfully computes and returns the task result, and it *cheats* when it returns some incorrect value. So, a worker decides to be honest or to cheat depending on which strategy maximizes its utility. We denote by p_{Ci}^r the probability of a worker i cheating in round r . This probability is not fixed, as the worker adjusts it over the course of the computation.

While it is assumed that workers make their decision individually and with no coordination, it is assumed that all the workers that cheat in a round return the same incorrect value (as performed, e.g., in [8, 9, 22]). This assumption yields a worst case scenario (and hence analysis) for the master with respect to 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).

2.2. Auditing, payoffs, rewards, and aspiration

To ‘persuade’ workers to be honest, the master employs, when necessary, *auditing* and *reward/punishment* schemes. The master, in a round, 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 honest. We denote by p_A the probability of the master auditing the responses of the workers. The master can change this auditing probability over the course of the computation. Unless otherwise stated, we assume that there is a value $p_A^{min} > 0$ so that at all times $p_A \geq p_A^{min}$. Furthermore, the master can reward and punish workers, which can be used (possibly combined with auditing) to encourage workers to be honest. 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 it rewards only the majority. We refer to this as the \mathcal{R}_m reward scheme (as presented in previous works).

The payoff parameters considered in this work are detailed in Table I. 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. 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.

We assume that each worker i has an aspiration a_i (the same in all rounds), which is the minimum benefit it expects to obtain in a round. In order to motivate the worker to participate in the computation, the master must ensure that $WB_Y \geq a_i$; in other words, the worker has the potential of its aspiration to be covered. We assume that the master knows the aspirations. This information can be included, for example, in a contract the master and the worker agree upon, prior to the start of the computation.

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 tries to achieve the goal of eventual correctness (see later). All other parameters can either be fixed because they are system parameters or may also be chosen by the master (except the aspiration, which is a parameter set by each worker).

Table I. Payoffs. The parameters are nonnegative.

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
MB_R	Master’s benefit from accepting the right answer

2.3. Eventual correctness

The goal of the master is to eventually obtain a reliable computational platform. In other words, after some finite number of rounds, the system must guarantee that the master obtains the correct task results in every round with probability 1. We call such property *eventual correctness*.

Algorithm 1 Master's Algorithm

```

 $p_A \leftarrow x$ , where  $x \in [p_A^{min}, 1]$ 
for  $r \leftarrow 1$  to  $\infty$  do
  send a task  $T$  to all workers in  $W$ 
  upon receiving all answers do
    audit the answers with probability  $p_A$ 
    if the answers were not audited then
      accept the majority
    else
       $p'_A \leftarrow p_A + \alpha_m(\text{cheaters}(r)/n - \tau)$ 
       $p_A \leftarrow \min\{1, \max\{p_A^{min}, p'_A\}\}$ 
       $\forall i \in W$  : pay/charge  $\Pi_i$  to worker  $i$ 

```

Algorithm 2 Algorithm for Worker i

```

 $p_{Ci} \leftarrow y$ , where  $y \in [0, 1]$ 
for  $r \leftarrow 1$  to  $\infty$  do
  receive a task  $T$  from the master
  set  $S_i \leftarrow -1$  with probability  $p_{Ci}$ , and
     $S_i \leftarrow 1$  otherwise
  if  $S_i = 1$  then  $\sigma \leftarrow \text{compute}(T)$ 
  else  $\sigma \leftarrow \text{arbitrary solution}$ 
  send response  $\sigma$  to the master
  get payoff  $\Pi_i$ 
   $p'_{Ci} \leftarrow p_{Ci} - \alpha_w(\Pi_i - a_i)S_i$ 
   $p_{Ci} \leftarrow \max\{0, \min\{1, p'_{Ci}\}\}$ 

```

3. ALGORITHMIC MECHANISM

We now detail the algorithms run by the master and the workers.

3.1. Master's algorithm (Algorithm 1)

The master's algorithm begins by choosing the initial probability of auditing. After that, at each round, the master sends a task to all workers and, after all answers are received (a reliable network is assumed), the master audits the answers with probability p_A . In the case the answers are not audited, the master accepts the value contained in the majority of answers and continues to the next round with the same probability of auditing. In the case the answers are audited, the value p_A of the next round is reinforced (i.e., modified according to the outcome of the round). Then, the master rewards/penalizes the workers accordingly.

The master initially has scarce or no information about the environment (e.g., workers initial p_C). The initial probability of auditing will be set according to the information the master possesses. For example, if it has no information about the environment, a safe approach may be to initially set $p_A = 0.5$.

Observe that, when the answers are not audited, the master has no information about the number of cheaters in the round. Thus, the probability p_A remains the same as in the previous round. When the answers are audited, the master can determine the number of cheaters; we denote by $\text{cheaters}(r)$ the number of cheaters in round r . Then, the master adapts the auditing probability p_A according to this number. Observe that the algorithm guarantees $p_A \geq p_A^{min}$. This, combined with the property $p_A^{min} > 0$, will prevent the system to fall in a permanent set of 'bad' states where $p_A = 0$ and $p_C > 0$. A discount factor, which we call *tolerance* and denote by τ , expresses the master's tolerable ratio of cheaters (typically, we will assume $\tau = 1/2$). Hence, if the proportion of cheaters is larger than τ , p_A will be increased, and otherwise, p_A will be decreased. The amount by which p_A changes depends on the change in the number of cheaters, modulated by a *learning rate* α_m . This latter value determines to what extent the newly acquired information will override the old information. (For example, if $\alpha_m = 0$ the master will never adjust p_A).

3.2. Workers' algorithm (Algorithm 2)

The workers' algorithm begins with each worker i deciding an initial probability of cheating p_{Ci} . At each round, each worker receives a task from the master and, with probability $1 - p_{Ci}$ calculates

the task and replies to the master with the correct answer. If the worker decides to cheat, it fabricates an answer and sends the incorrect response to the master. (We use a flag S_i to model the decision of a worker i to cheat or not.) After receiving its payoff (detailed in the analysis section), each worker i changes its p_{Ci} according to the payoff Π_i received, the chosen strategy S_i , and its aspiration a_i . Observe that the workers' algorithm guarantees $0 \leq p_{Ci} \leq 1$.

Workers have a learning rate α_w . We assume that all workers have the same learning rate, that is, they learn in the same manner (see the discussion in [13]; the learning rate is called step-size there); note that our analysis can be adjusted to accommodate also workers with different learning rates. We choose the value of α_w so that $\alpha_w(a_i + WP_C) < 1$, $\forall i \in W$. Otherwise, the system could enter in an oscillating condition where some nodes alternate p_C between 0 and 1 never converging to a stable state, which is necessary to guarantee reliability.

4. ANALYSIS

In this section, we analyze the mechanism presented in Section 3. We model the evolution of the mechanism as a Markov chain, and we prove necessary and sufficient conditions for achieving eventual correctness. We provide analytical evidence that convergence to eventual correctness can be reached rather quickly. Observe in Algorithms 1 and 2 that there are a number of variables that may change in each round. We will denote the value of a variable X after a round r with a superindex r , as X^r .

4.1. The mechanism as a Markov chain

We analyze the evolution of the master-workers system as a Markov chain. To do so, we first define the set of states and the transition function as follows.

Let the state of the Markov chain be given by the vector of probabilities $(p_A, p_{C1}, p_{C2}, \dots, p_{Cn})$. Then, we denote the state after round r by $(p_A^r, p_{C1}^r, p_{C2}^r, \dots, p_{Cn}^r)$. Observe from Algorithms 1 and 2 that any state $(p_A, p_{C1}, p_{C2}, \dots, p_{Cn})$ in which $p_A \in [p_A^{min}, 1]$ and $p_{Ci} \in [0, 1]$ for each worker i , is a possible initial state of the Markov chain. The workers' decisions, the number of cheaters, and the payoffs in round r are the stochastic outcome of the probabilities used in round r . Then, restricted to $p_A^r \in [p_A^{min}, 1]$ and $p_{Ci}^r \in [0, 1]$, we can describe the transition function of the Markov chain in detail. For each subset of workers $F \subseteq W$, $P(F) = \prod_{j \in F} p_{Cj}^{r-1} \prod_{k \notin F} (1 - p_{Ck}^{r-1})$ is the probability that the set of cheaters is exactly F in round r . Then, we have the following.

- With probability $p_A^{r-1} \cdot P(F)$, the master audits when the set of cheaters is F , and then,
 - (0) the master updates p_A as $p_A^r = p_A^{r-1} + \alpha_m(|F|/n - \tau)$, and
 - (1) each worker $i \in F$ updates p_{Ci} as $p_{Ci}^r = p_{Ci}^{r-1} - \alpha_w(a_i WP_C)$,
 - (2) each worker $i \notin F$ updates p_{Ci} as $p_{Ci}^r = p_{Ci}^{r-1} + \alpha_w(a_i - (WB_Y - WC_T))$.
- With probability $(1 - p_A^{r-1})P(F)$, the master does not audit when F is the set of cheaters. Then, the master does not change p_A and the workers update p_{Ci} as follows. For each $i \in F$,
 - (3) if $|F| > n/2$ then $p_{Ci}^r = p_{Ci}^{r-1} + \alpha_w(WB_Y - a_i)$,
 - (4) if $|F| < n/2$ then $p_{Ci}^r = p_{Ci}^{r-1} - \alpha_w \cdot a_i$,
 and for each $i \notin F$,
 - (5) if $|F| > n/2$ then $p_{Ci}^r = p_{Ci}^{r-1} + \alpha_w(a_i + WC_T)$,
 - (6) if $|F| < n/2$ then $p_{Ci}^r = p_{Ci}^{r-1} + \alpha_w(a_i - (WB_Y - WC_T))$.

The following terminology will be used throughout. Let a *covered worker* be one that is paid at least its aspiration a_i and the computing cost WC_T . In any given round r , let an *honest worker* be one for which $p_C^{r-1} = 0$. Let an *honest state* be one where the majority of workers are honest. Let an *honest set* be any set of honest states. We refer to the opposite cases as *uncovered worker*, *cheater worker* ($p_C^{r-1} = 1$), *cheat state*, and *cheat set*, respectively.

4.2. Conditions for eventual correctness

We show the conditions under which the system can guarantee eventual correctness. We begin with some terminology. Let a set of states S be called *closed* if, once the chain is in any state $s \in S$, it

will not move to any state $s' \notin S$. (A singleton closed set is called an *absorbing* state.) For any given set of states S , we say that the chain *reaches* (resp. *leaves*) the set S if the chain reaches some state $s \in S$ (resp. reaches some state $s \notin S$).

In order to show eventual correctness, we must show eventual convergence to a closed honest set. Thus, we need to show (i) that there exists at least one such closed honest set; (ii) that all closed sets are honest; and (iii) that one honest closed set is reachable from any initial state. Lemma 1 shows that, if $p_A = 0$ then some cheat set is closed. Given (ii), the necessity of $p_A^{min} > 0$ is motivated by this claim. Hence, $p_A > 0$ is assumed for the rest of the analysis. Lemma 2 shows that, if the majority of workers is uncovered, no honest set is closed. Given (i), the necessity of a covered majority is motivated. Hence, it is assumed that the majority of workers are covered for the rest of the analysis. Lemma 3 shows that the honest set including all the states in which all covered workers are honest is closed, which proves (i). Lemma 4 shows that any honest set where some covered worker is not honest is not closed, and Lemma 5 shows that any set that is not honest is not closed. Together, they prove (ii), and also (iii) because, if only honest sets are closed, there is a way of going from non-honest sets to one of them. The overall result is established in Theorem 6.

Lemma 1

Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z : WB_Y \geq a_i$. If $|Z| > n/2$, then the set of states $S = \{(p_A, p_{C1}, \dots, p_{Cn}) | (p_A = 0) \wedge (\forall w \in Z : p_{Cw} = 1)\}$, is a closed cheat set.

Proof

Observe first that each state in S is a cheat state, because the master does not audit and a majority of workers cheat. From transition (3) it can be seen that, if the chain is in a state of the set S before round r , for each worker $i \in Z$, $p_{Ci}^r \geq p_{Ci}^{r-1} = 1$ holds. In addition, p_A does not change. Hence, once the chain has reached a state in the set S , it will move only to states in the set S . \square

As already mentioned, from this lemma, it is concluded that $p_A > 0$ is required for eventual correctness. From now on, it is assumed that in all rounds $p_A \geq p_A^{min} > 0$.

Lemma 2

If there exists a set of workers $Z \subseteq W$ such that $|Z| > n/2$ and $\forall i \in Z : WB_Y < a_i + WC_T$, then no honest set is closed.

Proof

Recall that we choose the value of α_w so that $\forall i \in W : n\alpha_w(a_i + WP_C) < 1$. Consider any starting state, which by assumption is an honest state, $S = \{(p_A, p_{C1}, \dots, p_{Cn}) | \exists Y \subseteq W : (|Y| > n/2) \wedge (\forall w \in Y : p_{Cw} = 0)\}$.

Let the set Z be divided in three sets depending on whether the workers are honest ($p_C = 0$), cheaters ($p_C = 1$), or cheat with a probability between zero and one ($0 < p_C < 1$). We denote these sets by Z_0 , Z_1 , and Z_b , respectively. In the next round, the master audits (possible because $p_A > 0$), workers in Z_0 and Z_b do not cheat and workers in Z_1 cheat. Then from transition (2), all workers in Z_0 and Z_b increase their probability of cheating. From transition (1), all workers in Z_1 decrease their cheating probability by $\alpha_w(a_i + WP_C)$. Because all workers in Z_1 are cheater workers ($p_C = 1$) and $\alpha_w(a_i + WP_C) < 1$, after this round their cheating probability is larger than 0. Hence, for all workers in Z their cheating probability is larger than 0 and the new state is not honest. \square

Lemma 3

Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z : WB_Y \geq a_i + WC_T$ and $\forall j \notin Z : WB_Y < a_j + WC_T$. If $|Z| > n/2$, then the set of states $S = \{(p_A, p_{C1}, \dots, p_{Cn}) | \forall w \in Z : p_{Cw} = 0\}$, is an honest closed set.

Proof

Consider any round r before which the state of the chain is $s \in S$. Given that $|Z| > n/2$, at round r we have $cheaters(r) < n/2$. Then, for all workers in Z , the transition function is either (2) or (6),

depending on whether the master audits or not. Then, given that $WB_Y \geq a_i + WC_T$ for all workers in Z , their probability of cheating after round r is still 0. Hence, the claim follows. \square

Lemma 4

Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z : WB_Y \geq a_i + WC_T$ and $\forall j \notin Z : WB_Y < a_j + WC_T$. Then, for any set of states $S = \{(p_A, p_{C1}, \dots, p_{Cn}) | \exists Y \subseteq W : (|Y| > n/2) \wedge (\forall w \in Y : p_{Cw} = 0) \wedge (Z \not\subseteq Y)\}$, S is not a closed set.

Proof

For the sake of contradiction, assume that S is a closed set. Then, after a round r when the state is $s \in S$, the chain remains in S forever. Given that $|Y| > n/2$ and the assumption that S is a closed set, at all rounds $r' > r$, we must have $cheaters(r') < n/2$. Then, given that $\forall i \in Z : WB_Y \geq a_i + WC_T$, from the transition function, it can be seen that the probability of cheating of all workers in Z decreases in every round, independently of whether the master audits or not. But then, at some round $r' > r$, for all $i \in Z$, $p_{Ci}^{r'} = 0$ must hold. Then, $Z \subseteq Y$ at round r' , showing that S is not closed. \square

Lemma 5

Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z : WB_Y \geq a_i + WC_T$ and $\forall j \notin Z : WB_Y < a_j + WC_T$. If $|Z| > n/2$ and $p_A > 0$, then for any set of states $S = \{(p_A, p_{C1}, \dots, p_{Cn}) | \exists Y \subseteq W : (|Y| > n/2) \wedge (\forall w \in Y : p_{Cw} > 0)\}$, S is not a closed set.

Proof

To prove this claim, it is enough to show that, after a round r when the state is $s \in S$, with some positive probability the chain moves out of S . Assume that, starting at some round $r' \geq r$, the master audits in all rounds in $[r', r'']$, for a suitable r'' . Such assumption is valid because $p_A > 0$. Then, given that $\forall i \in Z : WB_Y \geq a_i + WC_T$, from the transition function it can be seen that the probability of cheating of all workers in Z decreases in every round by an amount not smaller than $\alpha_w \min\{WB_Y - a_i - WC_T, WP_C + a_i\}$. But then, at some round $r'' > r$, for all $i \in Z$, we have $p_{Ci}^{r''} = 0$. Therefore, $|Y| < n/2$ at round r'' showing that S is not closed. \square

The following theorem shows that there is a positive probability of reaching some state after which correctness can be guaranteed, as long as for a chosen majority of workers, the payment is enough to cover their aspiration and cost of performing the task. Its proof follows directly from Lemmas 3–5.

Theorem 6

If $p_A > 0$ and for all $i \in W : \alpha_w(a_i + WP_C) < 1$ then, in order to guarantee with positive probability that, after some finite number of rounds, the system achieves eventual correctness, it is *necessary* and *sufficient* to set $\boxed{WB_Y \geq a_i + WC_T}$ for all $i \in Z$ in some set $Z \subseteq W$ such that $|Z| > n/2$.

Remark 1

From Algorithm 1, it is easy to see that once the closed set $S = \{(p_A, p_{C1}, \dots, p_{Cn}) | \forall w \in Z : p_{Cw} = 0\}$ is reached, eventually $p_A = p_A^{min}$ and stays such forever.

4.3. Convergence time

Theorem 6 gives necessary and sufficient conditions to achieve eventual correctness. However, in order to have a practical system, it is necessary to bound the time taken to achieve it, which we call the *convergence time*. In other words, starting from any initial state, we want to compute the number of rounds that the Markov chain takes to reach an honest closed set. In this section, we show bounds on the convergence time.

4.3.1. Expected convergence time. Let C be the set of all covered workers. We assume, as required by Theorem 6, that $|C| > n/2$. From transitions (1) and (2) in the Markov chain definition, it can be seen that it is enough to have a consecutive sequence of $1/(\alpha_w \min\{WB_Y - a_i - WC_T, WP_C + a_i\})$ audits to enforce $p_C = 0$ for all covered workers $i \in C$. This gives the following upper bound on the convergence time.

Theorem 7

The expected convergence time is at most $\rho/(p_A^{min})^\rho$, where $\rho = 1/(\alpha_w \min_{i \in C}\{WB_Y - a_i - WC_T, WP_C + a_i\})$, and C is the set of covered workers.

Proof

The expected convergence time is upper bounded by the expected time for ρ consecutive audits. Consider the time divided in phases of ρ rounds. Let a phase where the master audits in all rounds be called *successful*. The expected time for ρ consecutive audits is at most the expected time for a successful phase. The probability of success in any given phase is at least $(p_A^{min})^\rho$. Consider the probability distribution of the number X of phases needed to have success, each with probability $(p_A^{min})^\rho$. This distribution is geometric, and the expectation of X is $1/(p_A^{min})^\rho$. Given that each phase has ρ rounds, the claim follows. \square

The upper bound shown in Theorem 7 may be too pessimistic for certain values of the parameters. The following theorem provides a tighter bound under certain conditions.

Theorem 8

Let us define, for each worker i , $dec_i \triangleq \alpha_w \min\{WP_C + a_i, WB_Y - WC_T - a_i\}$, $inc_i \triangleq \alpha_w \max\{WB_Y - a_i, WC_T + a_i\}$. Let C be the set of covered workers. If $p_A^{min} = \max_{i \in C}\{inc_i/(inc_i + dec_i)\} + \varepsilon$, for some $0 < \varepsilon < 1 - \max_{i \in C}\{inc_i/(inc_i + dec_i)\}$, then the expected convergence time is $1/(\varepsilon(\min_{i \in C}\{dec_i\} + \max_{i \in C}\{inc_i\}))$.

Proof

Let us define a potential function ϕ over the rounds as follows. Initially $\phi(0) = \max_{i \in C}\{p_{Ci}^0\}$. Then, for each round $r > 0$, $\phi(r) = \phi(r-1)$ if $\phi(r-1) = 0$. If $\phi(r-1) > 0$, then $\phi(r) = \max(0, \phi(r-1) - \min_{j \in C}\{dec_j\})$ if the master audits in round $r > 0$, and $\phi(r) = \phi(r-1) + \max_{j \in C}\{inc_j\}$ otherwise.

Consider any worker $i \in C$, and observe that, in the extreme cases, p_{Ci} decreases by $\min_{j \in C}\{dec_j\}$ when the master audits and increases by $\max_{j \in C}\{inc_j\}$ when the master does not audit. Also, once all workers in C have $p_{Ci} = 0$, this value does not change, because there is a majority of honest workers. Hence, it is clear that for all $r \geq 0$, $\phi(r) \leq \max_{i \in C}\{p_{Ci}^r\}$.

As a worst case, assume that $\phi(0) = 1$. We compute the expected number of rounds needed to obtain $\phi = 0$ as follows. We need $1/dec_i$ audits for each $1/inc_i$ non-audit rounds to compensate for the increase in potential. Setting $p_A^{min} = inc_i/(inc_i + dec_i)$ the master achieves at least that ratio in expectation for any period. (Omitting that time is discrete for clarity.) Additionally, in order to compensate for the initial $\phi(0) = 1$, $1/dec_i$ additional audits are needed. Making $p_A^{min} = inc_i/(inc_i + dec_i) + \varepsilon$, for some $0 < \varepsilon < 1 - inc_i/(inc_i + dec_i)$, the expected convergence time is $1/(\varepsilon(\min_{i \in C}\{dec_i\} + \max_{i \in C}\{inc_i\}))$. \square

The following corollary is derived from the previous theorem for a suitable scenario.

Corollary 9

If $WP_C + a_i \geq WB_Y - WC_T - a_i$ and $WB_Y - a_i \leq WC_T + a_i$, $\forall i \in C$, and if $p_A^{min} = \frac{WC_T + \max_{i \in C} a_i}{WB_Y} + \varepsilon$, where C is the set of covered workers and $0 < \varepsilon < 1 - (WC_T + \max_{i \in C} a_i)/WB_Y$, then the expected convergence time is ρ/ε , where $\rho = 1/(\alpha_w WB_Y)$.

4.3.2. *Probabilistic bound on the number of rounds for convergence.* We show now that, under certain conditions on the parameters of the system, it is possible to bound the probability to achieve convergence and the number of rounds to do so. Assume that $p_{\mathcal{A}}^0 > 0$. Because $p_{\mathcal{A}}$ is not changed unless the master audits, we have the following.

Lemma 10

Let $p_{\mathcal{A}}^0 = p > 0$. Then, the master audits in the first $\rho = \ln(1/\varepsilon_1)/p$ rounds with probability at least $1 - \varepsilon_1$, for any $\varepsilon_1 \in (0, 1)$.

Proof

The master audits in the first ρ rounds with probability $1 - (1 - p)^\rho \geq 1 - \exp(-\rho \cdot p) = 1 - \varepsilon_1$. \square

Let us assume that the system parameters are such that, for all workers i , $\alpha_w(WP_C + a_i) \in [0, 1]$ and $\alpha_w(WB_Y - WC_T - a_i) \in (0, 1]$ (all workers are covered). Let us define $dec_cheater \triangleq \alpha_w \min_i \{WP_C + a_i\}$ and $dec_honest \triangleq \alpha_w \min_i \{WB_Y - WC_T - a_i\}$. From transitions (1) and (2), we derive the following lemma.

Lemma 11

Let r be a round in which the master audits, and F be the set of cheaters in round r . Then,

$$p_{Ci}^r \leq 1 - \alpha_w(WP_C + a_i) \leq 1 - dec_cheater, \forall i \in F$$

$$p_{Cj}^r \leq 1 - \alpha_w(WB_Y - WC_T - a_j) \leq 1 - dec_honest, \forall j \notin F.$$

Let us denote the sum of all cheating probabilities before a round r as $P^{r-1} \triangleq \sum_i p_{Ci}^{r-1}$.

Lemma 12

Let r be a round in which the master audits such that $P^{r-1} > n/3$. If $dec_cheater \geq dec_honest$ and $dec_cheater + 3 \cdot dec_honest \geq 8/3$, then $P^r \leq n/3$ with probability at least $1 - \exp(-n/96)$.

Proof

Let F be the set of cheaters in round r . Then, using a Chernoff bound $Pr[|F| < (1 - \delta)P^{r-1}] \leq \exp(-\delta^2 P^{r-1}/2)$, for any $\delta \in (0, 1)$. Then, because $P^{r-1} > n/3$, using $\delta = 1/4$, there are at least $(1 - \delta)P^{r-1} > n/4$ cheaters with probability at least $1 - \exp(-\delta^2 P^{r-1}/2) > 1 - \exp(-n/96)$. If that is the case, from Lemma 11 and $dec_cheater \geq dec_honest$, we have that

$$\begin{aligned} P^r &\leq n - |F|dec_cheater - (n - |F|)dec_honest \leq n - (n/4)dec_cheater \\ &\quad - (3n/4)dec_honest = n(1 - (dec_cheater + 3 \cdot dec_honest)/4) \leq n/3, \end{aligned}$$

as desired. \square

Let us now define $dec_i \triangleq \alpha_w \min\{a_i, WB_Y - WC_T - a_i\}$. Let, $dec \triangleq \min_i dec_i$. Assume $WP_C \geq 0$ and $a_i \geq 0$, for all workers.

Lemma 13

Consider a round r such that $P^{r-1} \leq n/3$. Then, with probability at least $1 - \exp(-n/36)$ each worker i has $p_{Ci}^r \leq \max\{0, p_{Ci}^{r-1} - dec\}$, and hence $P^r \leq n/3$.

Proof

Using Chernoff, there is a majority of honest workers with probability at least

$$Pr[\text{majority honest} | P^{r-1} \leq n/3] \geq 1 - \exp(-(1/2)^2(n/3)/3) = 1 - \exp(-n/36).$$

It can be observed in Algorithm 2 that, if there is a majority of honest workers in a round r , then any worker i has $p_{Ci}^r \leq \max\{0, p_{Ci}^{r-1} - dec\}$, independently of whether the master audits. Hence the proof. \square

Theorem 14

Assume $\alpha_w(WP_C + a_i) \in [0, 1]$ and $\alpha_w(WB_Y - WC_T - a_i) \in (0, 1]$ for all workers i . (Observe that all workers are covered.) Let $dec_cheater \triangleq \alpha_w \min_i \{WP_C + a_i\}$, $dec_honest \triangleq \alpha_w \min_i \{WB_Y - WC_T - a_i\}$, and $dec \triangleq \alpha_w \min_i \{a_i, WB_Y - WC_T - a_i\}$. If $p_A^0 = p > 0$, $dec_cheater \geq dec_honest$ and $dec_cheater + 3 \cdot dec_honest \geq 8/3$, then eventual convergence is reached in at most $\ln(1/\varepsilon_1)/p + 1/dec$ rounds, with probability at least $(1 - \varepsilon_1)(1 - \exp(-n/96))(1 - \exp(-n/36))^{1/dec}$, for any $\varepsilon_1 \in (0, 1)$.

Proof

Consider the first round r in which the masters audits. From Lemma 10, r is in the first $\ln(1/\varepsilon_1)/p$ rounds with probability at least $1 - \varepsilon_1$. If so, either $P^{r-1} \leq n/3$ and also $P^r \leq n/3$ (from Algorithm 2 and the fact that the master audits in round r , P cannot increase in round r), or $P^{r-1} > n/3$. In this latter case, from Lemma 12, $P^r \leq n/3$ with probability at least $1 - \exp(-n/96)$. Then starting at round $r + 1$, from Lemma 13, with probability at least $(1 - \exp(-n/36))^{1/dec}$, there are $1/dec$ consecutive rounds with majorities of honest workers. Because in each of these rounds the cheating probability of any worker decreases at least by dec (unless it is already zero), at the end of these rounds all workers have zero cheating probability. \square

5. SIMULATIONS

This section complements our analytical results with illustrative simulations. The graphical representation of the data obtained captures the tradeoffs between reliability and cost, a concept hard to view through the analysis. This is important as our analytical upper bounds on convergence time correspond to worst case scenarios. Here, we present simulations for a variety of parameter combinations likely to occur in practice. We have created our own simulation setup by implementing our mechanism (the master's and the workers' algorithms) using the C++ programming language. We have run our simulations on a PC with an Intel Core 2 Duo, 2.80 GHz CPU, 4 GB of RAM and Ubuntu 11.04 OS. Each depicted plot value represents the average over 10 executions of the implementation. The plots of Figure 4 are an exception, and the plotted values represent only 1 execution of the implementation; the purpose is to illustrate the per round cost of the master. We have simulated several scenarios for different parameter values (here we present selected results).

5.1. Simulation parameters

We choose sensible parameter values likely to be encountered in real applications; the choice of the parameters was influenced by statistics obtained from experiments contacted in SETI-like projects [46–48]. In particular, the number of workers has been set to nine (an odd number to accommodate majority voting when the master does not audit). In systems such as Seti@home, typically, each task is assigned to three workers [2]. So, in that context, nine workers seems an appropriate workforce. The initial cheating probability of each worker i is not known; therefore, we have experimented with $p_{Ci} = 0.5$, as a reasonable assumption, and with $p_{Ci} = 1$ as an extreme case. Similarly, we have set $p_A \in \{0.5, 1\}$ as the master's initial probability of auditing. The minimum probability of cheating is set to be $p_A^{min} = 0.01$ and tolerance $\tau = 0.5$, which means that the master will not tolerate a majority of cheaters. Besides this intuition on the value of tolerance, we also carried out a set of experiments to understand the effect of this parameter, on which we will comment later.

The payoffs for the workers are set using $WB_Y \in \{1, 2\}$ as our normalizing parameter, and we take $WP_C \in \{0, 1, 2\}$ and $WC_T = 0.1$ as realistic values (within the same order of magnitude as WB_Y) to explore the effects of these choices. The payoffs for the master are set to $MC_A = 20$, a large value allowing us to notice the impact of auditing, and we take $MC_Y = WB_Y = 1$. We have

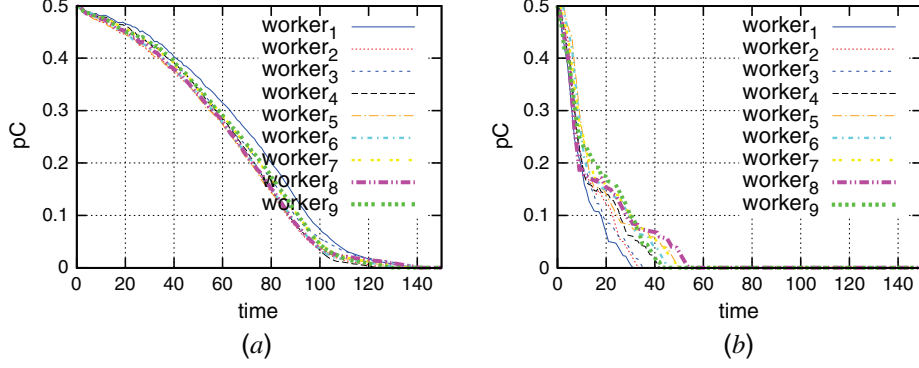


Figure 1. Cheating probability for each worker as a function of time (number of rounds) for parameters $p_C = p_A = 0.5$, $WB_Y = 1$, $WP_C = 0$, $WC_T = 0.1$ and $a_i = 0.1$. (a) $\alpha = 0.01$; (b) $\alpha = 0.1$.




Figure 2. Auditing probability for the master as a function of time (number of rounds) for parameters $p_C = p_A = 0.5$, $WB_Y = 1$, $WP_C = 0$, $WC_T = 0.1$ and $a_i = 0.1$. Left: $\alpha = 0.01$; Right: $\alpha = 0.1$.

set $MB_R = 0$ and $MP_W = 0$ to have a clear view of the master's cost. In the simulations, unless otherwise stated, the master covers all workers and not some majority as assumed in the analysis (as a worst case scenario with respect to the master's cost).

The aspiration is a parameter defined by the workers in an idiosyncratic manner; for simplicity, in these simulations, we consider all workers having the same aspiration level $a_i \in \{0.01, 0.1\}$. We have checked that, when values are assigned randomly around some mean, the results are similar to those presented here, provided the variance is not very large. As for the values for the aspiration and of the workers' cost for computing the task WC_T , they are such that the necessary conditions of Theorem 6 are satisfied, and hence, eventual convergence is reached. Finally, we consider the same learning rate for the master and the workers, that is, $\alpha = \alpha_m = \alpha_w$. For practical reasons [13], it must be set to a small constant value, so we consider $\alpha \in \{0.1, 0.01\}$.

5.2. Convergence time

Figure 1 shows that convergence can be reached very quickly, even without punishing cheaters and with small WB_Y . Note that even if all workers have the same aspiration level and begin with the same initial cheating probability, their evolution in time may be different from each other as it depends on the individual realizations of cheating. In Figure 1, we also notice that a slightly higher value of α can make the convergence time shorter. (As we argued before, the value of α can not be very high because the learning procedure will become unstable and p_C will bounce up and down without reaching convergence.) Similar conclusions can be drawn from Figure 2, where we can notice how quickly p_A drops to $p_A = 0.01$, and also that p_A decreases in the same manner as p_C . Notice however, that p_A decreases at a slower rate; intuitively, this is to ensure that workers will not try to deviate from the desirable behavior.

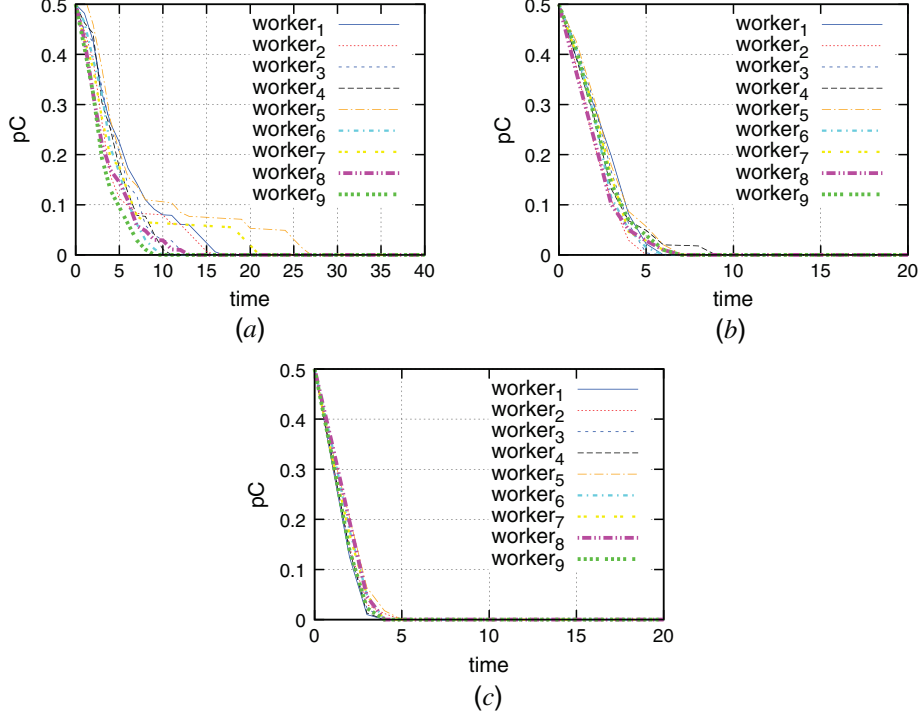


Figure 3. Cheating probability for each worker as a function of time (number of rounds) for parameters $p_C = p_A = 0.5$, $\alpha = 0.1$, $a_i = 0.1$, $WB_Y = 2$ and $WC_T = 0.1$. (a) $WP_C = 0$; (b) $WP_C = 1$; (c) $WP_C = 2$.

5.3. Effects of punishment

Notice that in previous simulations only a positive reinforcement is applied to the workers (i.e., $WP_C = 0$). Now, from Figure 3, we can notice that the larger the punishment we apply (i.e., $WP_C \in \{1, 2\}$), the faster the convergence time is. In fact, we may conclude that applying only punishment is enough to have fast convergence. Comparing Figure 1(b) with Figure 3(a), we observe that, for a specific set of parameter values, a larger WB_Y leads to a shorter convergence time. Interestingly, this observation reveals a trade-off between convergence time and the cost the master has for reaching faster convergence and maintaining it. Thus, the master could choose between different protocols estimating the cost of the auditing until it reaches convergence. But less auditing leads to larger convergence times. So it is not clear initially what is going to be optimal.

5.4. Master's cost

Recall from the simulation parameters that only the cost of auditing and the workers' payment are nonzero, and that the master covers all workers (instead of some majority as in the analysis). We contrasted first a worst case scenario where workers initially cheat with probability $p_C = 1$, against the case where, given the lack of knowledge about worker's behavior, we assume that initially $p_C = p_A = 0.5$ (refer to Figure 2).

The first conclusion we can draw from our simulations is that, even in this unfavorable situation, eventual convergence is still achieved. However, during the process, the master's auditing probability reaches 1 for the system to converge. Of course, this has a direct impact on the cost of convergence, but also on the convergence time. Denote by $p_A(0)$ the master's initial auditing probability. Interestingly, from Figure 4(a2), (a3), (b2), and (b3), we observe that, when $p_A(0) = 1$, the convergence time decreases by 9%, and that even p_A converges to its minimum in 25% less time yielding also a cost reduction. In fact, during the first 10 rounds of evolution, for $p_A(0) = 0.5$, the aggregate cost for the master is 56% smaller (Figure 4(a1)) than when $p_A(0) = 1$ (Figure 4(b1)). However, in the subsequent interval between $p_A^{min} < p_A < 1$, the situation is reversed, and the master's cost for $p_A(0) = 1$ is 19% smaller than the case $p_A(0) = 0.5$. These results show that using

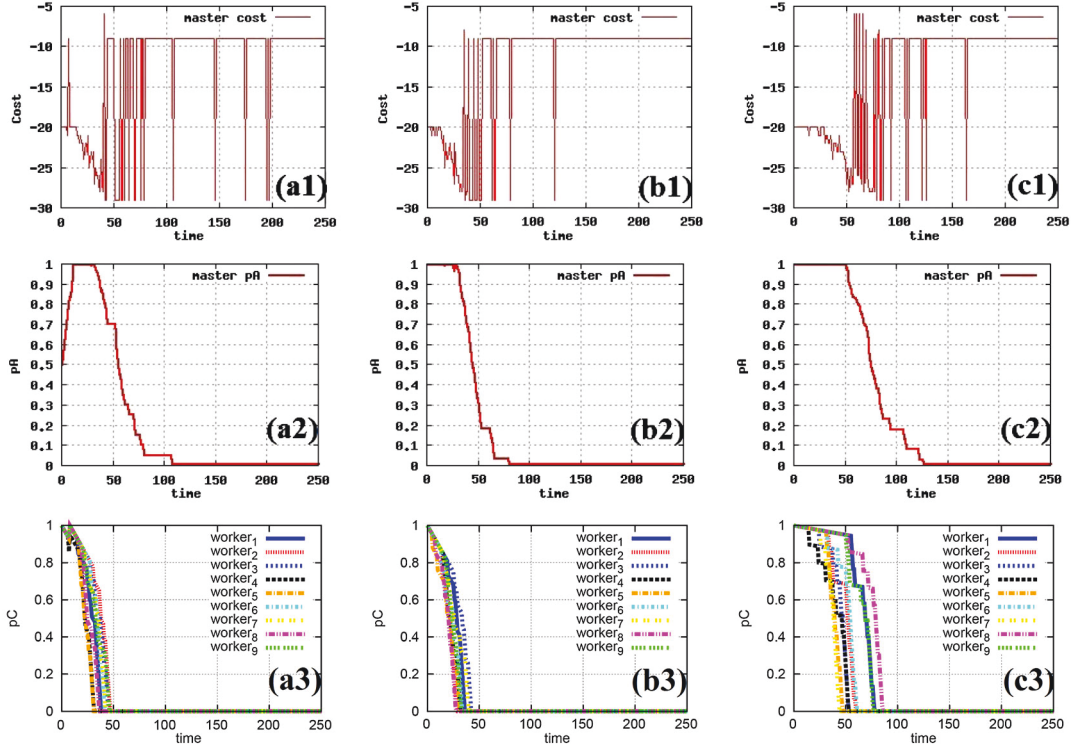


Figure 4. Top panel, master’s cost as a function of time. Middle panel, master’s auditing probability as a function of time. Bottom panel, worker’s cheating probability as a function of time. Parameters in all panels, $p_C = 1$, $WC_T = 0.1$, $WP_C = 0$ and $\alpha = 0.1$. Left panel, $p_A = 0.5$, $a_i = 0.1$. Middle panel $p_A = 1$, $a_i = 0.1$. Right panel $p_A = 1$, $a_i = 0.01$.

$p_A(0) = 1$ does not necessarily increase cost as the intuition might suggest. After convergence is achieved, Figure 4(a1), (b1), and (c1) show that, once the master’s auditing probability has reached its minimum value, the master audits roughly once every hundred rounds, which is expected given that $p_A^{min} = 0.01$. This behavior is observed independently of the initial auditing probability, which is also expected.

Another surprising result, arising from Figure 4(b3) and (c3), is that having workers with larger aspiration values makes the convergence time decrease by more than a half. The reason being that the master initially audits with probability one and all workers cheat, so a larger aspiration causes the workers’ cheating probability to drop at a higher rate. This, in turn, feeds back to the master’s auditing probability and cost, making them decrease faster than the case where workers have a smaller aspiration.

We have also examined the case where only a majority of workers is covered. Specifically, we have run analogous simulations to the ones depicted in Figure 4, but now covering only five out of the nine workers. A first, interesting observation, is that the convergence time for the covered workers is not affected. An even more interesting observation is that the master’s cost, until p_A^{min} was reached, is greater than the case of all covered workers. This is due to the slower rate at which p_A reaches its minimum value. Of course, after this point, the master’s cost is smaller because it rewards fewer workers.

5.5. Tolerance value

Finally, we have also considered the effect of tolerance for achieving eventual convergence using three different initial $p_C \in \{0.3, 0.5, 1\}$. Choosing 5000 iterations as a value large enough to illustrate the problems of convergence for high tolerance (values of the same order of magnitude or larger give basically the same results), we have found that convergence will always be reached for the lower values of p_C . This is due to the fact that with ‘almost’ honest workers, a majority of them

will always compute the answer and will force punishment to the minority of cheaters. However, when the initial p_C equals one and there is no punishment ($WP_C = 0$), the master must respond to a percentage of cheating workers (due to tolerance) to obtain eventual convergence, and as a consequence convergence is not achieved for large tolerance values (for this specific set of parameters, when $\tau > 0.9$). Interestingly, with nonzero punishment ($WP_C = 1$), the master can tolerate all workers cheating and still achieve convergence. Therefore, our simulated examples establish that, based on the rest of the parameter values, the master can appropriately set values for tolerance τ , (its own) learning rate α , punishment WP_C and perhaps WB_Y in such a way that convergence can be swiftly obtained.

6. CONCLUSIONS AND FUTURE WORK

This work applies reinforcement learning techniques to capture the evolution of Internet-based master-worker computations. We show that under necessary and sufficient conditions, the master reaches a state after which the correct task result is obtained at each round, with minimal cost. In addition, we show that such state can be reached ‘quickly’. The convergence analysis is complemented with simulations. Our simulation results suggest that when having a positive reinforcement learning (i.e., $WP_C = 0$) the master can reach fast convergence, while applying punishments (i.e., $WP_C \in \{1, 2\}$) provides even faster convergence. In fact, we may conclude that applying only punishment is enough to have fast convergence. Also, our simulations demonstrated that it is plausible, even if the master begins with an aggressive auditing strategy, to have a smaller cost and convergence time compared with the case of using a less aggressive auditing policy.

The analysis has provided us with provable guarantees of our approach, and the simulations give an insightful view that is hard to detect from the analysis. In view of the potential of our approach, the next logical step is to do real-world experiments, which is the context of our future work. However, such kind of experiment is nontrivial. We are currently in the stage of designing such real-world experiments. We expect to obtain interesting results when comparing our current approach with the observed behavior of real workers. In a follow up work, we additionally consider the presence of malicious workers that always provide the master with an incorrect answer. This will provide additional fault tolerance and security to the system.

ACKNOWLEDGEMENTS

We thank Carlos Diuk for useful discussions on reinforcement learning. This work is supported by the Cyprus Research Promotion Foundation grant TIE/ΠΛΗΠΟ/0609(BE)/05, the National Science Foundation (CCF-0937829, CCF-1114930), Comunidad de Madrid grants S2009TIC-1692 and MODELICO-CM, Spanish PRODIEVO and RESINEE grants and MICINN grant TEC2011-29688-C02-01, and National Natural Science Foundation of China grant 61020106002. Conference version: A preliminary version of this work, entitled ‘Achieving Reliability in Master–Worker Computing via Evolutionary Dynamics’ appears in the proceedings of Euro-Par 2012, pages 451–463.

REFERENCES

1. Anderson D. BOINC: a system for public-resource computing and storage. In *proc. of GRID 2004*; 4–10.
2. Korpela E, Werthimer D, Anderson D, Cobb J, Lebofsky M. SETI@home: massively distributed computing for SETI. *Computing in Science and Engineering* 2001; **3**(1):78–83.
3. Amazon’s Mechanical Turk. <https://www.mturk.com>.
4. Golle P, Mironov I. Uncheatable distributed computations. In *proc. of CT-RSA 2001*; 425–440.
5. Heien EM, Anderson DP, Hagihara K. Computing low latency batches with unreliable workers in volunteer computing environments. *Journal of Grid Computing* 2009; **7**:501–518.
6. Abraham I, Dolev D, Goden R, Halpern JY. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In *proc. of PODC 2006*; 53–62.
7. Shneidman J, Parkes DC. Rationality and self-interest in P2P networks. In *Proc. of IPTPS 2003*; 139–148.
8. Sarmenta L. Sabotage-tolerance mechanisms for volunteer computing systems. *Future Generation Computer Systems* 2002; **18**(4):561–572.
9. Fernández Anta A, Georgiou C, Mosteiro MA. Designing mechanisms for reliable Internet-based computing. In *proc. of NCA 2008*; 315–324.

10. Fernández Anta A, Georgiou C, Mosteiro MA. Algorithmic mechanisms for Internet-based master-worker computing with untrusted and selfish workers. In *proc. of IPDPS 2010*; 1–11.
11. Yurkewych M, Levine BN, Rosenberg AL. On the cost-ineffectiveness of redundancy in commercial P2P computing. In *proc. of CCS 2005*; 280–288.
12. Bush RR, Mosteller F. *Stochastic Models for Learning*. Wiley, 1955.
13. Szepesvári C. *Algorithms for Reinforcement Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool publishers, 2010.
14. Bendor J, Mookherjee D, Ray D. Aspiration-based reinforcement learning in repeated interaction games: an overview. *International Game Theory Review* 2001; **3**(2-3):159–174.
15. Camerer CF. Behavioral game theory: experiments in strategic interaction. *Roundtable Series in Behavioral Economics* 2003.
16. Easley D, Kleinberg J. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, 2010.
17. Gintis MC. *Game Theory Evolving*. Princeton University Press, 2000.
18. Weibull JW. *Evolutionary Game Theory*. MIT Press, 1995.
19. Halpern JY. Computer science and game theory: a brief survey. *Palgrave Dictionary of Economics* 2007.
20. Abraham I, Alvisi L, Halpern J, Y. Distributed computing meets game theory: combining insights from two fields. *ACM SIGACT News: Distributed Computing Column* 2011; **42**(2):69–76.
21. Nisan N, Roughgarden T, Tardos E, Vazirani VV. *Algorithmic Game Theory*. Cambridge University Press, 2007.
22. Fernández A, Georgiou C, Lopez L, Santos A. Reliably executing tasks in the presence of untrusted processors. In *proc. of SRDS 2006*; 39–50.
23. Konwar KM, Rajasekaran S, Shvartsman AA. Robust network supercomputing with malicious processes. In *proc. of DISC 2006*; 474–488.
24. Christodoulou G, Koutsoupias E. Mechanism design for scheduling. *Bulletin of the EATCS* 2009; **97**:39–59.
25. Nisan N, Ronen A. Algorithmic mechanism design. *Games and Economic Behavior* 2001; **35**:166–196.
26. Nash JF. Equilibrium points in n -person games. *National Academy of Sciences* 1950; **36**(1):48–49.
27. Nowak MA. *Evolutionary Dynamics*. Harvard University Press, 2006.
28. Maynard-Smith J. *Evolution and the Theory of Games*. Cambridge University Press, 1982.
29. Maynard-Smith J, Price GR. The logic of animal conflict. *Nature* 1973; **246**:15–18.
30. Henrich J, McElreath R. The evolution of cultural evolution. *Evolutionary Anthropology: Issues, News, and Reviews* 12.3, 2003:123–135.
31. Borgers T, Sarin R. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory* 1997; **77**(1):1–14.
32. Samuelson L. *Evolutionary Games and Equilibrium Selection*. MIT Press, 1998.
33. Bendor J, Mookherjee D, Ray D. Reinforcement learning in repeated interaction games. *Advances in Theoretical Economics* 2001; **1**(1).
34. Laslier J, Topol R, Walliser B. A behavioral learning process in games. *Games and Economic Behavior* 2001; **37**:340–366.
35. Izquierdo LR, Izquierdo SS. Dynamics of the Bush-Mosteller learning algorithm in 2x2 games. *Reinforcement Learning: Theory and Applications* 2008.
36. Phelps S, McBurney P, Parsons S. Evolutionary mechanism design: a review. *Journal of Autonomous Agents and Multi-Agent Systems* 2010.
37. Conitzer V, Sandholm T. Incremental mechanism design. In *Proc. IJCAI 2007*.
38. Babaioff M, Feldman M, Nisan N. Combinatorial agency. In *proc. of ACM EC 2006*; 18–28.
39. Babaioff M, Feldman M, Nisan N. Mixed strategies in combinatorial agency. In *proc. of WINE 2006*; 353–364.
40. Eidenbenz R, Schmid S. Combinatorial agency with audits. In *proc. of GameNets 2009*.
41. Babaioff M, Feldman M, Nisan N. Free riding and free labor in combinatorial agency. In *proc. of SAGT 2009*.
42. Mass-Colell A, Whinston M, Green J. *Microeconomic Theory*. Oxford University Press, 1995.
43. Rose D, Willemain TR. The principal-agent problem with evolutionary learning. *Computational and Mathematical Organization Theory* 1996; **2**:139–162.
44. Taufer M, Anderson D, Cicotti P, Brooks CL. Homogeneous redundancy: a technique to ensure integrity of molecular simulation results using public computing. In *proc. of IPDPS 2005*.
45. BOINC Security. http://boinc.berkeley.edu/wiki/BOINC_Security.
46. Einstein@home. <http://einstein.phys.uwm.edu>.
47. Estrada T, Taufer M, Anderson DP. Performance prediction and analysis of BOINC projects: an empirical study with EmBOINC. *Journal of Grid Computing* 2009; **7**(4):537–554.
48. WUProp@home. <http://wuprop.boinc-af.org/>.