

Achieving Reliability in Master-Worker Computing via Evolutionary Dynamics

Evgenia Christoforou, Antonio Fernández Anta, Chryssis Georgiou, Miguel A. Mosteiro, Angel (Anxo) Sánchez

www.cs.ucy.ac.cy/ric

Motivation Framework Evolutionary Mechanism (No Reputation) Background

Computational Tasks

- Increasing demand for processing complex computational tasks
- One-processor machines have **limited** computational resources
- Powerful parallel machines (supercomputers) are **expensive** and are **not globally available**

Internet emerges as a **viable** platform for supercomputing

- P2P, Grid and Cloud computing
- e.g., EGEE Grid, TERA Grid, Amazon's EC2
- Volunteer Master-Worker computing: @home projects
- e.g., SETI@home, AIDS@home, Folding@home, PrimeNet
- Amazon's Mechanical Turk (Contractor-based approach)
- Internet-based Computing**

SETI-like Internet-based Computing

Comparable processing power with top Supercomputers @ a fraction of the cost!

Task 1 Task 2 Task 3 Task 4 Task 5 Task 6 Task 7 Task 8 Task 9 Task 10 Task 11 Task 12 Task 13 Task 14 Task 15 Task 16 Task 17 Task 18 Task 19 Task 20 Task 21 Task 22 Task 23 Task 24 Task 25 Task 26 Task 27 Task 28 Task 29 Task 30 Task 31 Task 32 Task 33 Task 34 Task 35 Task 36 Task 37 Task 38 Task 39 Task 40 Task 41 Task 42 Task 43 Task 44 Task 45 Task 46 Task 47 Task 48 Task 49 Task 50 Task 51 Task 52 Task 53 Task 54 Task 55 Task 56 Task 57 Task 58 Task 59 Task 60 Task 61 Task 62 Task 63 Task 64 Task 65 Task 66 Task 67 Task 68 Task 69 Task 70 Task 71 Task 72 Task 73 Task 74 Task 75 Task 76 Task 77 Task 78 Task 79 Task 80 Task 81 Task 82 Task 83 Task 84 Task 85 Task 86 Task 87 Task 88 Task 89 Task 90 Task 91 Task 92 Task 93 Task 94 Task 95 Task 96 Task 97 Task 98 Task 99 Task 100

REDUNDANCY

Great potential limited by untrustworthy entities

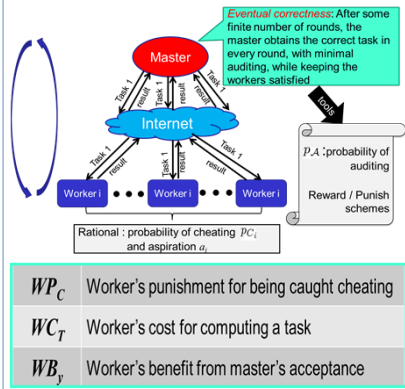
Prior Work

- Rational workers: act upon their best interest, i.e., choose the strategy that maximizes their own benefit! (Jonathan Parkes 03)
- Workers with predefined behavior:
 - Malicious: fabricate and report a bogus result
 - Altruistic: compute and truthfully report correct result
- In Internet-based master-worker task computation
 - Honest: compute and report correct result
 - Cheat: fabricate and return a bogus result
- Mechanisms with reward/punish schemes that provide incentives to workers to be honest
- One shot: in each round a task is performed and no knowledge is forwarded to the next round (Jurek et al. 2005, Fernandez et al. 2006)

Can the repeated interaction between the master and the workers be exploited effectively?

Our Approach

- We introduce the concept of **evolutionary dynamics** under the biological and social perspective and relate them to Internet-based master-worker task computing
- Employ **reinforcement learning** both on Master and Workers (Camerer 03, Szepesvari 10)
- Objective:** Develop a **reliable computation platform** where the master obtains the correct task results



Conditions for Eventual Correctness (No Reputation)

For the system to achieve **eventual correctness**, it is **necessary and sufficient** to set:

$$W_B \geq a_i + W_C, \forall i \in Z, |Z| > n/2$$

Given that $P_A > 0$

Lemma 1. Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z: W_B \geq a_i$. If $|Z| > n/2$, then the set of states $S = \{(P_A, PC_1, \dots, PC_n) | (P_A = 0) \wedge (\forall w \in Z: PC_w = 1)\}$, is a closed cheat set.

→ Lemma 1: Motivates the necessity of $P_A > 0$

Lemma 2. If there exists a set of workers $Z \subseteq W$ such that $|Z| > n/2$ and $\forall i \in Z: W_B < a_i + W_C$ and $\forall i \in W: a_i + W_C < 1$, then no honest set is closed.

→ Lemma 2: Motivates the necessity of a covered majority

Lemma 3. Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z: W_B \geq a_i + W_C$ and $\forall j \notin Z: W_B < a_j + W_C$. If $|Z| > n/2$, then the set of states $S = \{(P_A, PC_1, \dots, PC_n) | \forall w \in Z: PC_w = 0\}$, is a closed set.

→ Lemma 3: Proves that there exists at least one honest closed set

Lemma 4. Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z: W_B \geq a_i + W_C$ and $\forall j \notin Z: W_B < a_j + W_C$. Then, for any set of states $S = \{(P_A, PC_1, \dots, PC_n) | \exists Y \subseteq W: (|Y| > n/2) \wedge (\forall w \in Y: PC_w = 0) \wedge (Z \not\subseteq Y)\}$, S is not a closed set.

Lemma 5. Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z: W_B \geq a_i + W_C$ and $\forall j \notin Z: W_B < a_j + W_C$. If $|Z| > n/2$ and $P_A > 0$, then for any set of states $S = \{(P_A, PC_1, \dots, PC_n) | \exists Y \subseteq W: (|Y| > n/2) \wedge (\forall w \in Y: PC_w > 0)\}$, S is not a closed set.

→ Lemma 4-5: Proves that all closed sets are honest and that one honest closed set is reachable from any initial state

Master's Protocol

Set initial P_A (e.g., 0.5)

Repeat

Send a task to all n workers

Upon receiving all answers do

Audit the answers with probability P_A

If the answers were not audited then

Accept the value returned by the majority

Else

$P_A \leftarrow P_A + \alpha_m \cdot \left(\frac{\text{cheaters}}{n} - \tau \right)$

Give appropriate payoff I_i to each worker i

α_m : learning rate (tunes the extent of change)

τ : tolerance (desirable fraction of cheaters, e.g., 0.5)

Protocol for Worker i

Set initial PC_i (e.g., 0.5)

Repeat

Receive a task from the master

Set $S_i = -1$ with probability PC_i , $S_i = 1$ otherwise

If $S_i = 1$ then compute the task and send the result

Else send an arbitrary result

Get payoff I_i

$PC_i \leftarrow PC_i - \alpha_w \cdot (|I_i - a_i| \cdot S_i)$

α_w : learning rate (tunes the extent of change)

We analyze the evolution of the master-worker system as a Markov chain:

ROUND r

$(P_A^r, PC_1^r, PC_2^r, \dots, PC_n^r)$

Master's audit prob. update: $P_A^r = P_A^{r-1} + \alpha_m (|I|/n - \tau)$

worker's cheating prob. update: cheating worker: $PC_i^r = PC_i^{r-1} - \alpha_w (a_i + W_C)$ honest worker: $PC_i^r = PC_i^{r-1} + \alpha_w (a_i - (W_B - W_C))$

restricted to $[0, 1]$

Set of cheaters r

ROUND $r+1$

$(P_A^{r+1}, PC_1^{r+1}, PC_2^{r+1}, \dots, PC_n^{r+1})$

cheating prob. update: $PC_i^r \leftarrow PC_i^r - \alpha_w \cdot a_i$

$|I| > n/2$

$|I| < n/2$

Audit

No Audit

Cheating worker

Honest worker

$PC_i^r \leftarrow PC_i^r + \alpha_w (W_B - a_i)$

$PC_i^r \leftarrow PC_i^r + \alpha_w (a_i + W_C)$

$PC_i^r \leftarrow PC_i^r - \alpha_w \cdot a_i$

$PC_i^r \leftarrow PC_i^r + \alpha_w (a_i - (W_B - W_C))$

Evolutionary Dynamics

- Evolutionary dynamics applied first in biology
 - Tool to study the mathematical principles according to which life is evolving
 - Inspiration for many fields: sociology, economics, artificial intelligence (multi-agent systems) etc.
- Inspired by dynamics of evolution as a mean to model workers adaptation to a truthful behavior

Evolutionary Stable Strategy

Evolutionary Game Theory

In biological terms: the application of game theory to evolving populations of life forms

Our aim: **Evolutionary Stable Strategy**

A strategy is called evolutionary stable if, when the whole population is using this strategy, any group of mutants (mutants) using a different strategy will eventually die over multiple generations (evolutionary rounds)

Reinforcement Learning

PLAYER

Environment

Percepts

Actions

Decision

Action A

Action B

Notion of Aspiration

- Bush and Mosteller's model, **aspiration based**
 - player's adapt by comparing their experience with an aspiration level
 - an aspiration a_i for player i
 - the minimum benefit it expects to obtain in an interaction

Reputation

- Accumulated information about an entity
- Induce **learning** by signaling the true abilities of involved entities
- Computer science:
 - On-line community exchange, eBay
 - Buyers, sellers: positive, negative and neutral rating
 - P2P systems
 - >adTorrent: increase reputation by uploading => increase download speed (if for lat)
 - >Gnutella (who to download from?)
- A reputation system can be:
 - Centralized
 - Decentralized

Evolutionary Mechanism with Reputation

Reputation

- Master maintains a **reputation** for each worker
- Workers are **ignorant** towards the reputation scheme
- Update** when master audits
- Calculated based on:
 - number of audits up to round r : $audit(r)$
 - number of times worker i was **honest** when master audited up to round r : $correct_audit(r)$
- Reputation types:
 - Type 1: $p_i = \frac{correct_audit(r) + 1}{audit(r) + 2}$ (based on Sonnet et al. 2007)
 - Type 2: $p_i = e^{(audit(r) - correct_audit(r))}$ where $\epsilon \in (0, 1)$
- Need to cope with **malicious** workers

Master's Protocol

Set initial P_A (e.g., 0.5)

Repeat

Send a task to all n workers

Upon receiving all answers do

Audit the answers with probability P_A

If the answers were not audited then

Accept the value returned by the majority

Else

$\sum_{i \in W} p_i > \sum_{i \in W} W_C$

Update worker's i reputation, and

$P_A = P_A + \alpha_m \cdot \left(\frac{\sum_{i \in W} p_i}{n} - \tau \right)$

Give appropriate payoff I_i to each worker i

If master audited then update the n most reputable workers

α_m : learning rate

τ : tolerance

W_C : set of workers

W_P : set of cheaters

W_P : workers in W with the same answer

Simulations

- Choose parameters likely to be encountered:
 - 9 workers (e.g., SETI@home3 workers)
 - initial $PC_i = 0.5$
 - initial $P_A = 0.5$
 - $\tau = 0.5$ (master does not tolerate a majority of cheaters)
 - aspiration $a_i = 0.1$ for each worker
 - $\alpha = \alpha_m = \alpha_w, \alpha \in \{0.1, 0.01\}$
 - $W_B \in \{1, 2\}$ set as our normalizing
 - $W_C = 0.1$
 - $W_P = 0$
- Auditing probability of the master as a function of time
- Workers cheating probability as a function of time
- Workers cheating probability as a function of time
- Auditing probability of the master as a function of time
- Extreme case where initial $PC_i = 1$
- Master's mechanism with no reputation cannot cope efficiently with malicious workers

Evolutionary Mechanism with Reputation

Choosing from a Pool

- There is a **pool** of N workers available to the master
- In each round the master selects n out of these workers, the **most reputable** ones
- In a round
 - If the master **audits**
 - updates the **reputation** of the workers
 - updates the **set** of the n most reputable workers to be used in the next round
 - Otherwise, it uses the same set of workers for the next round
- In our simulations: Take 5 most reputable workers out of 9

Master's Protocol

Set initial P_A (e.g., 0.5)

Choose n most reputable workers

Repeat

Send a task to all n workers

Upon receiving all answers do

Audit the answers with probability P_A

If the answers were not audited then

Accept the value returned by the majority

Else

$\sum_{i \in W} p_i > \sum_{i \in W} W_C$

Update worker's i reputation, and

$P_A = P_A + \alpha_m \cdot \left(\frac{\sum_{i \in W} p_i}{n} - \tau \right)$

Give appropriate payoff I_i to each worker i

If master audited then update the n most reputable workers

α_m : learning rate

τ : tolerance

W_C : set of workers

W_P : set of cheaters

W_P : workers in W with the same answer

Malicious and Rational Workers

4 malicious, 5 rational workers, reputation type 2

Workers cheating probability as a function of time

Workers cheating probability as a function of time

Dealing with communication uncertainty

- Reputation is based both on **auditing** and **responsiveness**
- Master's protocol outline:
 - Send task to the n most reputable workers
 - Wait for time T :
 - For workers that do not reply reduce their reputation (using reputation type)
 - Then proceed as usual on the workers that have replied



ΚΥΠΡΙΑΚΗ ΔΗΜΟΚΡΑΤΙΑ

This work is supported by the Cyprus Research Promotion Foundation grant ΤΠΕ/ΠΑΗΡΟ/0609(ΒΕ)/05