Achieving Reliability in Master-Worker Computing via Evolutionary Dynamics

Evgenia Christoforou — University of Cyprus
Antonio Fernandez Anta — Institute IMDEA Networks
Chryssis Georgiou — University of Cyprus
Miguel A. Mosteiro — Kean University and URJC
Angel Sanchez — Univ. Carlos III de Madrid

August 30th, 2012
Rhodes, Greece

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

Internet-Based Task Computing

- 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
  - Grid and Cloud computing
    - e.g., EGEE Grid, TERA Grid, Amazon's EC2
  - Master-Worker volunteer computing: @home projects
    - e.g., SETI@home, AIDS@home, Folding@home, PrimeNet

SETI-Like Internet-Based Computing

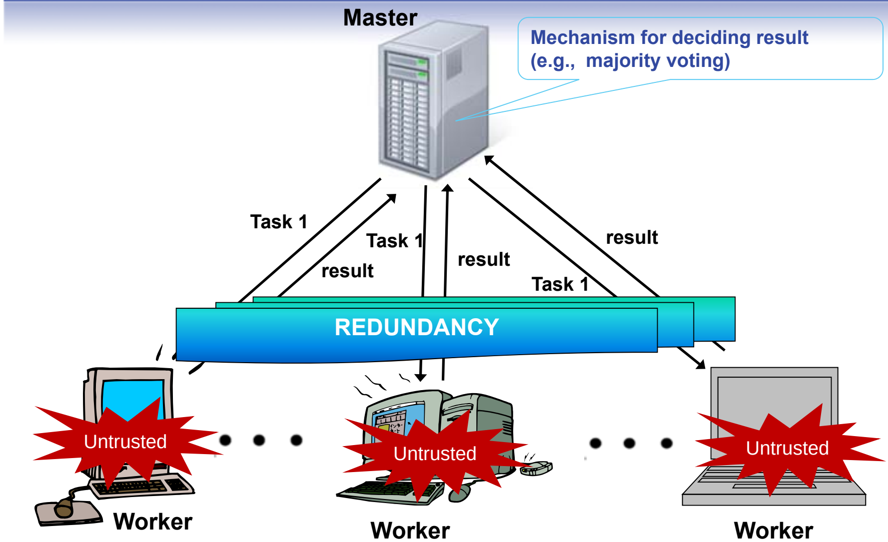

Prior Work

- Rational workers: act upon their best interest, i.e., choose the strategy that maximizes their own benefit
  [Shneidman Parkers 03]
- 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
  [Yurkewych et al 2005, Fernandez et al 2008]

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

Background: Evolutionary Dynamics

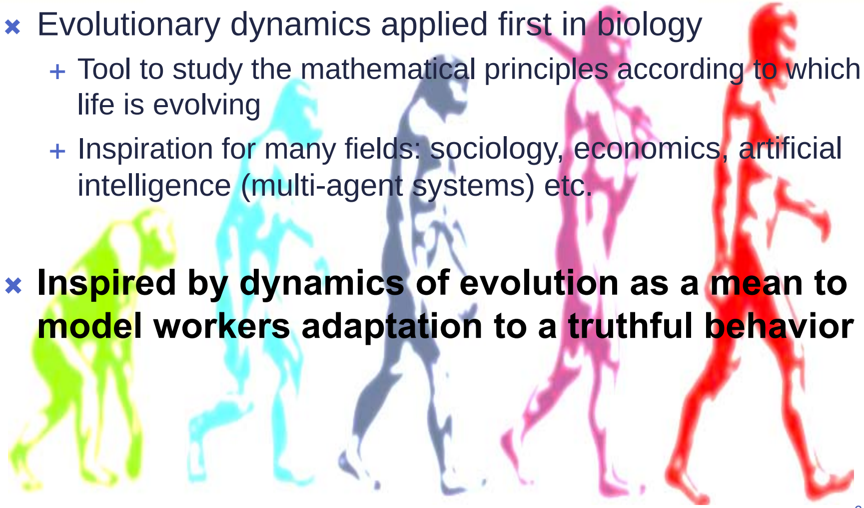

- Evolutionary dynamics applied first in biology
  - Tool to study the mathematical principles according to which life is evolving
  - Inspiration for many fields: sociology, economics, artificial intelligence (multi-agent systems) etc.
- **Inspired by dynamics of evolution as a mean to model workers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

Background: 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 invaders (mutants) using a different strategy will eventually die over multiple generations (evolutionary rounds).

[Gintis 2000]

Background: Reinforcement Learning

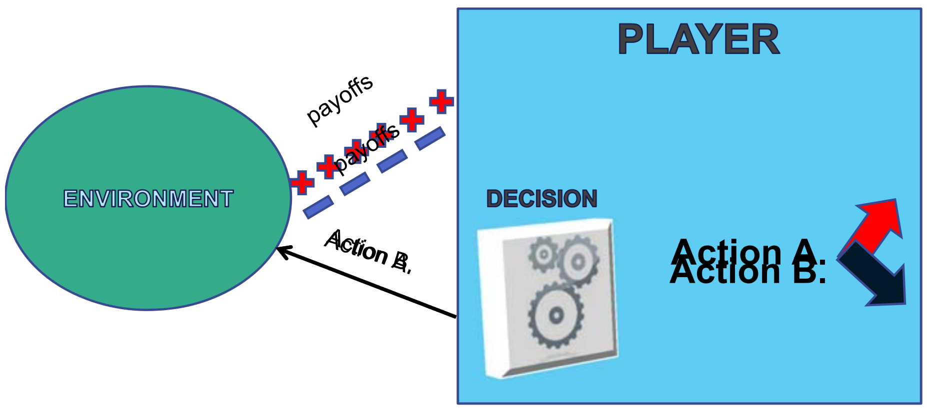

Background: Notion of Aspiration

- Bush and Mosteller's model, aspiration based
  - player's adapt by comparing their experience with an aspiration level
    [Bush Mosteller 55]
  - an aspiration a_i for player *i*
    - the minimum benefit it expects to obtain in an interaction

Contributions (i)

Initiate the study of the evolutionary dynamics of Internet-based master-worker computations through reinforcement learning :

- Develop and analyze a mechanism based on reinforcement learning to be used by the master and the workers

Contributions (ii)

- Show necessary and sufficient conditions under which the mechanism ensures eventual correctness (EC)
- *Convergence time*: The number of rounds to achieve eventual correctness
  - We show, both in expectation and with high probability, that our mechanism reaches convergence time quickly
  - Complement our analysis with simulations

Framework

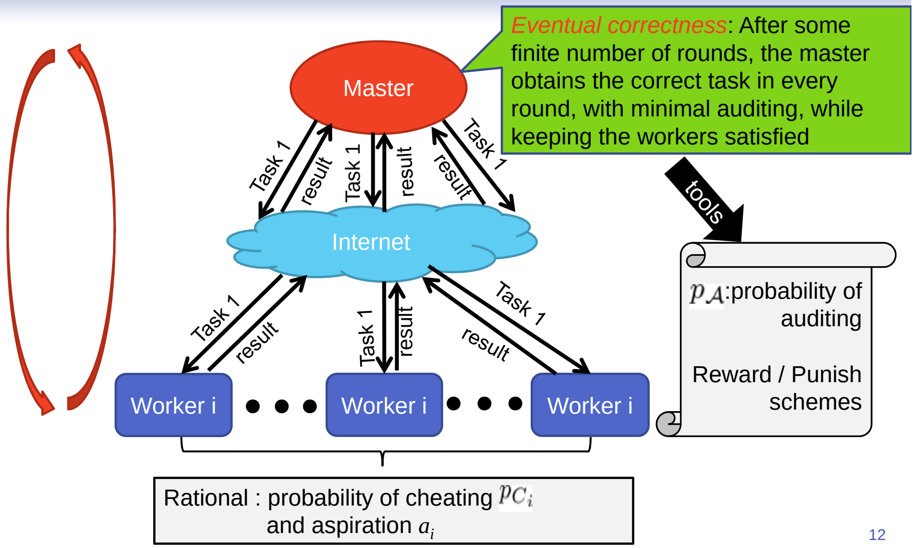

Payoffs

WP_C	Worker's punishment for being caught cheating
WC_T	Worker's cost for computing a task
WB_y	Worker's benefit from master's acceptance

Master's Protocol

Set initial $p_{\mathcal{A}}$ (e.g., 0.5)
Repeat
  Send a task to all *n* workers
  Upon receiving all answers **do**
    Audit the answers with probability $p_{\mathcal{A}}$
    If the answers were *not* audited **then**
      Accept the value returned by the majority
    Else
$$p_{\mathcal{A}} \leftarrow p_{\mathcal{A}} + \alpha_m \cdot \left(\frac{cheaters}{n} - \tau \right)$$
Give appropriate payoff Π_i to each worker i

α_m : learning rate (tunes the extent of change)
τ : tolerance (tolerable fraction of cheaters, e.g., 0.5)

Protocol for Worker i

Set initial p_{C_i} (e.g., 0.5)
Repeat
  Receive a task from the master
  Set S_i = -1 *with* probability p_{C_i} , S_i = 1 otherwise
  If S_i = 1 then **compute** the task and **send** the result
  Else send an arbitrary result
  Get payoff Π_i
$$p_{C_i} \leftarrow p_{C_i} - \alpha_w \cdot (\Pi_i - a_i) \cdot S_i$$

α_w : learning rate (tunes the extent of change)

Conditions for Eventual Correctness

- We analyze the evolution of the master-worker system as a *Markov chain* and we show:

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

$$WB_y \geq a_i + WC_T, \ \ \forall i \in Z, \ |Z| > n/2$$

Given that $p_{\mathcal{A}} > 0$

Master-Worker System as Markov Chain

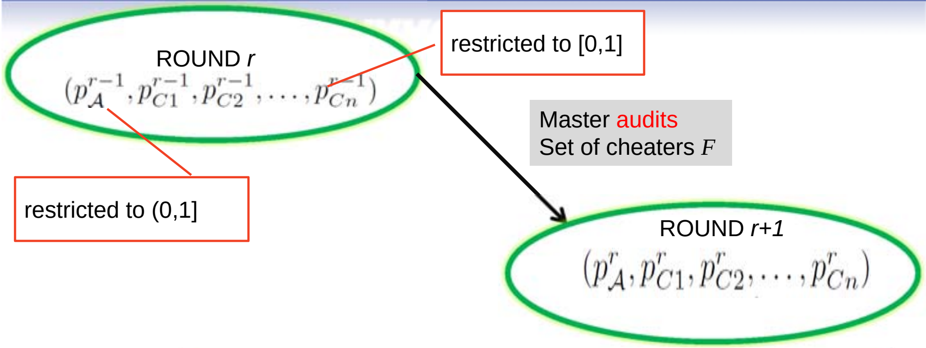

Master's audit prob. update: $p_{\mathcal{A}}^{r} = p_{\mathcal{A}}^{r-1} + \alpha_m(|F|/n - \tau)$

worker's cheating prob. update:

- cheating worker: $p_{Ci}^{r} = p_{Ci}^{r-1} - \alpha_w(a_i + WP_{\mathcal{C}})$
- honest worker: $p_{Ci}^{r} = p_{Ci}^{r-1} + \alpha_w(a_i - (WB_{\mathcal{Y}} - WC_{\mathcal{T}}))$

Master-Worker System as Markov Chain

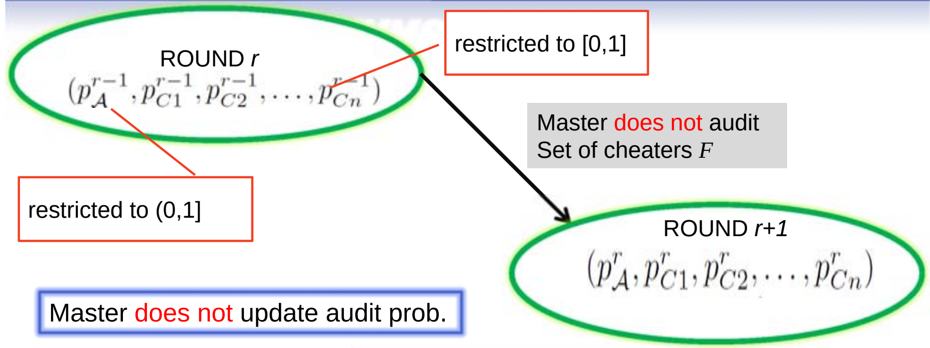

worker's cheating prob. update

	Cheating worker	Honest worker
$\lvert F\rvert > n/2$	$p_{Ci}^{r-1} + \alpha_w(WB_{\mathcal{Y}} - a_i)$	$p_{Ci}^{r-1} + \alpha_w(a_i + WC_{\mathcal{T}})$
$\lvert F\rvert < n/2$	$p_{Ci}^{r-1} - \alpha_w \cdot a_i$	$p_{Ci}^{r-1} + \alpha_w(a_i - (WB_{\mathcal{Y}} - WC_{\mathcal{T}}))$

Terminology

- Covered worker is one that receives at least its aspiration a_i and the computing $WC_{\mathcal{T}}$ cost
- In any given round r, honest worker is one for which $p_C^{r-1} = 0$
- Honest state is one where the majority of workers are honest
- Honest set is any set of honest states
- Opposite cases: uncovered worker, cheater worker, cheat state, and cheat set respectively
- 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 $s' \notin S$

Eventual Correctness Proof Roadmap

- To show eventual correctness, we must show eventual convergence to a closed honest set
- We need to show
  - that there exists at least one such closed honest set
  - that all closed sets are honest
  - that one honest closed set is reachable from any initial state

Eventual Correctness Proof Roadmap

Lemma 1. *Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z : WB_{\mathcal{Y}} \geq a_i$. If $|Z| > n/2$, then the set of states*

$$S = \{(p_{\mathcal{A}}, p_{C1}, \ldots, p_{Cn}) | (p_{\mathcal{A}} = 0) \wedge (\forall w \in Z : p_{Cw} = 1)\},$$

is a closed cheat set.

Lemma 1: Motivates the necessity of $p_{\mathcal{A}} > 0$

Lemma 2. *If there exists a set of workers $Z \subseteq W$ such that $|Z| > n/2$ and $\forall i \in Z : WB_{\mathcal{Y}} < a_i + WC_{\mathcal{T}}$, 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 $\forall i \in Z : WB_{\mathcal{Y}} \geq a_i + WC_{\mathcal{T}}$ and $\forall j \notin Z : WB_{\mathcal{Y}} < a_j + WC_{\mathcal{T}}$. If $|Z| > n/2$, then the set of states*

$$S = \{(p_{\mathcal{A}}, p_{C1}, \ldots, p_{Cn}) | \forall w \in Z : p_{Cw} = 0\},$$

is a closed set.

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

Eventual Correctness Proof Roadmap

Lemma 4. *Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z : WB_{\mathcal{Y}} \geq a_i + WC_{\mathcal{T}}$ and $\forall j \notin Z : WB_{\mathcal{Y}} < a_j + WC_{\mathcal{T}}$. Then, for any set of states*

$$S = \{ (p_{\mathcal{A}}, p_{C1}, \ldots, p_{Cn}) | \exists Y \subseteq W : (|Y| > n/2) \wedge (\forall w \in Y : p_{Cw} = 0) \wedge (Z \nsubseteq Y)\},$$

S is not a closed set.

Lemma 5. *Consider any set of workers $Z \subseteq W$ such that $\forall i \in Z : WB_{\mathcal{Y}} \geq a_i + WC_{\mathcal{T}}$ and $\forall j \notin Z : WB_{\mathcal{Y}} < a_j + WC_{\mathcal{T}}$. If $|Z| > n/2$ and $p_{\mathcal{A}} > 0$, then for any set of states*

$$S = \{(p_{\mathcal{A}}, p_{C1}, \ldots, p_{Cn}) | \exists Y \subseteq W : (|Y| > n/2) \wedge (\forall w \in Y : p_{Cw} > 0)\},$$

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

Eventual Correctness Proof Roadmap

Theorem 1. *If $p_{\mathcal{A}} > 0$ 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* $WB_{\mathcal{Y}} \geq a_i + WC_{\mathcal{T}}$ *for all $i \in Z$ in some set $Z \subseteq W$ such that $|Z| > n/2$.*

Examples of Convergence

- *Under certain conditions, the expected convergence time is*

$$\left(\alpha_w \cdot (WB_{\mathcal{Y}} - WC_{\mathcal{T}} - \max_i\{a_i\}) \cdot \varepsilon\right)^{-1}$$

where

$$\varepsilon \in (0, 1 - (WC_{\mathcal{T}} + \max_i\{a_i\})/WB_{\mathcal{Y}}).$$

Examples of Convergence

- *Under certain conditions, the convergence time is at most*

$$\ln(1/\varepsilon)/p + 1/dec$$

with probability at least

$$(1-\varepsilon)(1-e^{-n/96})(1-e^{-n/36})^{1/dec}$$

where

$$dec = \min_i\{\alpha_w \cdot \min\{a_i,\ WB_{\mathcal{Y}} - WC_{\mathcal{T}} - a_i\}\}, \text{and } \varepsilon \in (0,1)$$

Simulations

- We created our own simulation setup by implementing our mechanism
- Choose parameters likely to be encountered:
  - 9 workers (e.g. SETI@home 3 workers)
  - initial p_{C_i} = 0.5
  - initial $p_{\mathcal{A}}$ = 0.5
  - τ = 0.5 (master does not tolerate a majority of cheaters)
  - aspiration a_i = 0.1 for each worker
  - $\alpha = \alpha_m = \alpha_w \quad \alpha \in \{0.1, 0.01\}$
  - $WB_{\mathcal{Y}} \in \{1, 2\}$ set as our normalizing
  - $WC_{\mathcal{T}} = 0.1$
  - $WP_{\mathcal{C}} \in \{0, 1, 2\}$

Simulations

Cheating probability for the workers as a function of evolutionary rounds

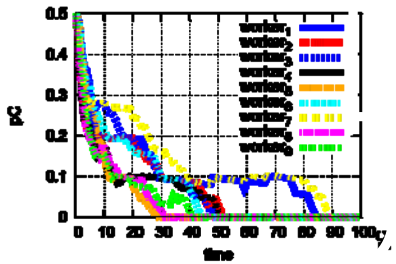

$\alpha = 0.1 \quad WB_{\mathcal{Y}} = 2$
$WP_{\mathcal{C}} = 0$

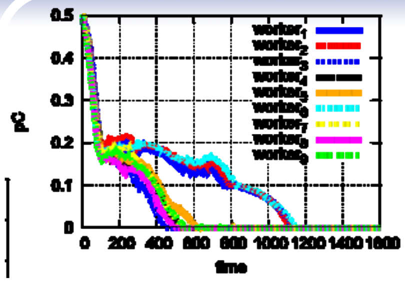

$\alpha = 0.01$
$WB_{\mathcal{Y}} = 1$
$WP_{\mathcal{C}} = 0$

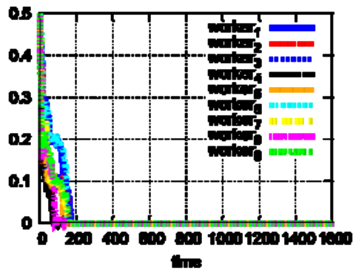

$\alpha = 0.1$
$WB_{\mathcal{Y}} = 1$
$WP_{\mathcal{C}} = 0$

Simulations

Cheating probability for the workers as a function of evolutionary rounds

$\alpha = 0.1 \quad WB_{\mathcal{Y}} = 2$

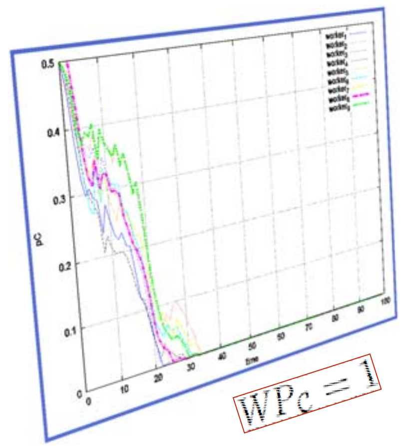

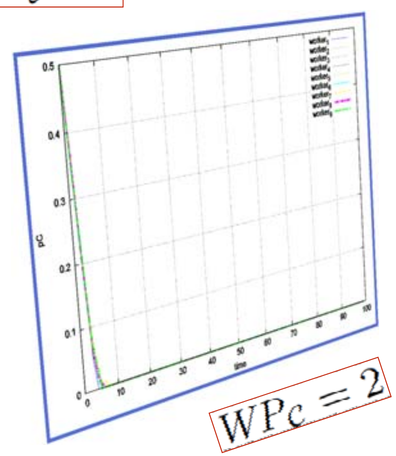

Simulations

Cheating probability for the workers on the 5000th round of evolution as a function of the tolerance

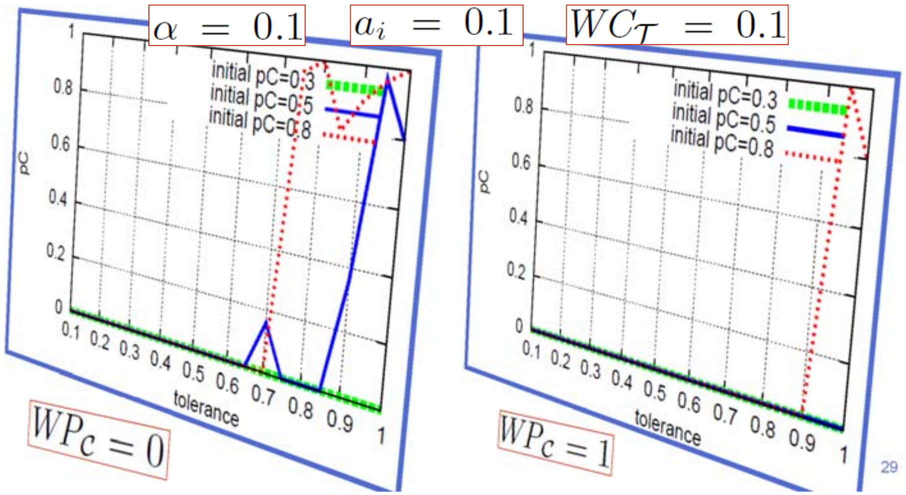

Summary

Initiate the study of the evolutionary dynamics of Internet-based master-worker computations through reinforcement learning:

- Develop and analyze our mechanism
- Under necessary and sufficient conditions the master reaches eventual convergence
- Our analysis shows that eventual convergence can be reached quickly
  - Complement our analysis with simulations

Future Work: Study the implications of adding a reputation system to our mechanism

Thank you!

Evgenia Christoforou

evgenia.christoforou@cs.ucy.ac.cy