Algorithmic Mechanisms for Reliable Master-Worker Internet-based Task Computing

Part II

Evgenia Christoforou
Dept. of Computer Science
University of Cyprus

October 30th, 2012
Nicosia, Cyprus

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

Outline

- Motivation and preliminaries
- Framework and contributions
- Evolutionary mechanism
  - Rational workers
  - Malicious workers
  - Communication failures
- Conclusions

MOTIVATION AND PRELIMINARIES

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

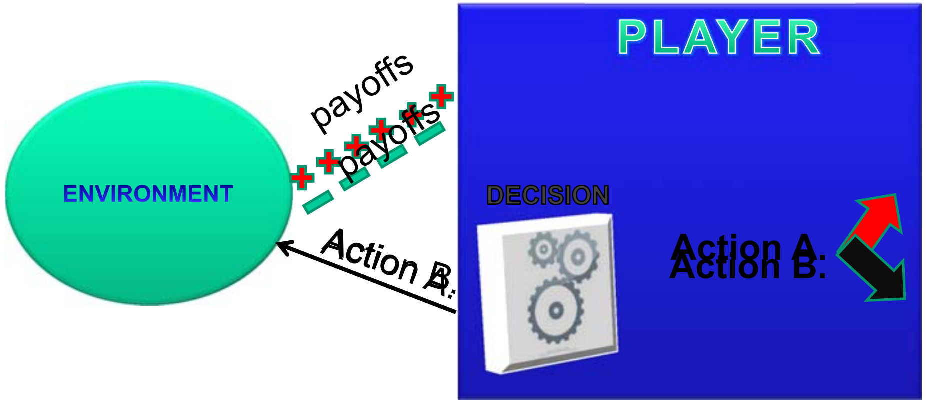

Background: Notion of Aspiration

- 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

FRAMEWORK AND CONTRIBUTIONS

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
- Add reputation to deal with malicious workers
- Choose the n out of N most reputable workers

Framework

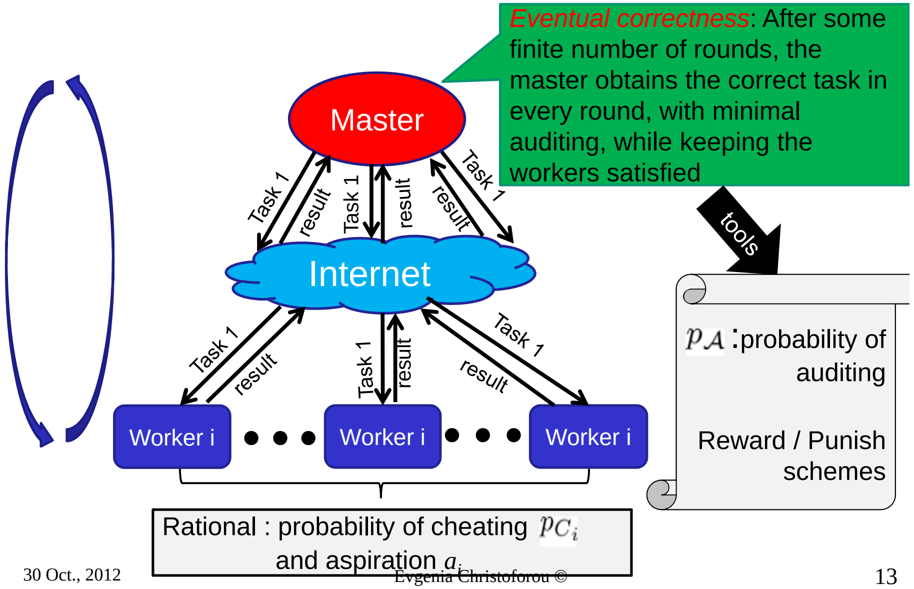

Payoffs

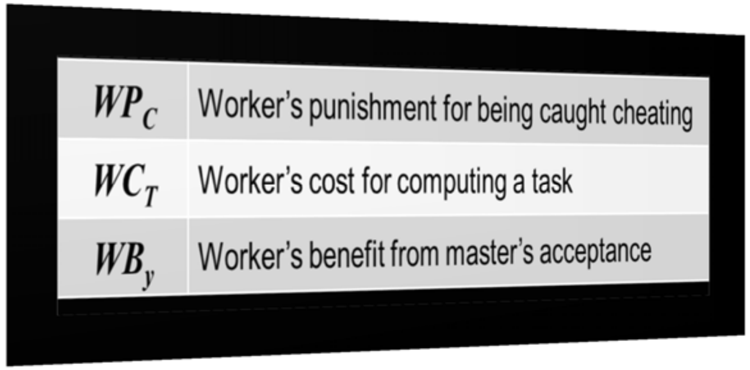

WP_C	Worker's punishment for being caught cheating
WC_T	Worker's cost for computing a task
WB_y	Worker's benefit from master's acceptance

EVOLUTIONARY MECHANISM
[RATIONAL WORKERS]

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

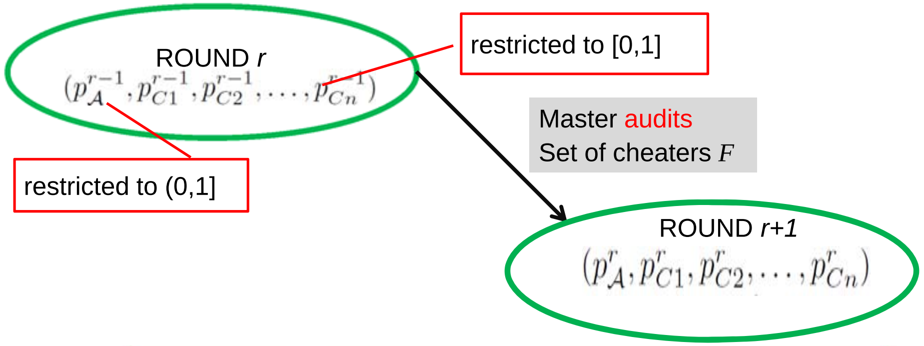

Master's audit prob. update: $p_{\mathcal{A}}^r = p_{\mathcal{A}}^{r-1} + \alpha_m(|F|/n - \tau)$

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

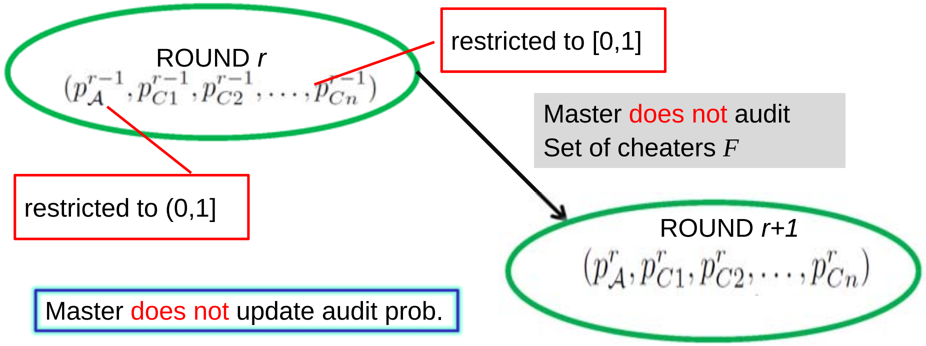

worker's cheating prob. update

	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 $\forall i \in Z$: $WB_{\mathcal{Y}} < a_i + WC_{\mathcal{T}}$ and $\forall i \in W : \alpha(a_i + WP_{\mathcal{C}}) < 1$, 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 $\forall i \in Z : WB_{\mathcal{Y}} \geq a_i + WC_{\mathcal{T}}$ 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 $p_{\mathcal{A}} > 0$ and for all $i \in W : \alpha(a_i + WP_{\mathcal{C}}) < 1$ 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* $\boxed{WB_{\mathcal{Y}} \geq a_i + WC_{\mathcal{T}}}$ *for all* $i \in Z' \subseteq Z$ *such that* $|Z'| = \lfloor n/2 \rfloor + 1$.

Examples of Convergence

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

$$\left(\alpha_w \cdot (WB_y - WC_T - \max_i\{a_i\}) \cdot \varepsilon\right)^{-1}$$

where

$$\varepsilon \in (0, 1 - (WC_T + \max_i\{a_i\})/WB_y).$$

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

$$dec = \min_i\{\alpha_w \cdot \min\{a_i,\ WB_y - WC_T - a_i\}\}, \text{and } \varepsilon \in (0,1)$$

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
  - $\alpha = \alpha_m = \alpha_w$, $\alpha \in \{0.1, 0.01\}$
  - $WB_{\mathcal{Y}} \in \{1, 2\}$ 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

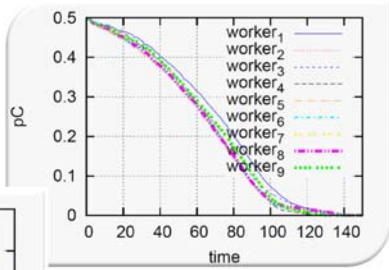

$\alpha = 0.01$ $WB_{\mathcal{Y}} = 1$ $WP_{\mathcal{C}} = 0$

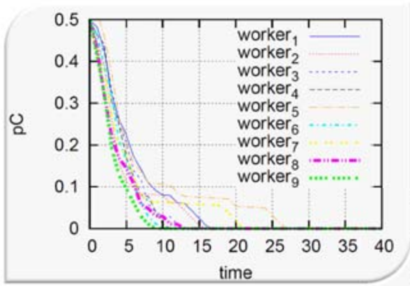

$\alpha = 0.1$ $WB_{\mathcal{Y}} = 2$ $WP_{\mathcal{C}} = 0$

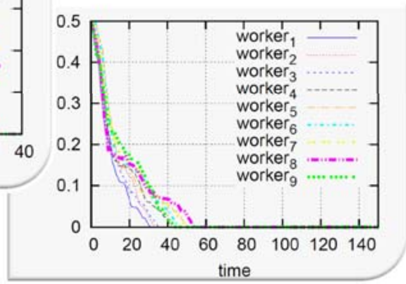

$\alpha = 0.1$ $WB_{\mathcal{Y}} = 1$ $WP_{\mathcal{C}} = 0$

Simulations

Cheating probability for the workers as a function of evolutionary rounds

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

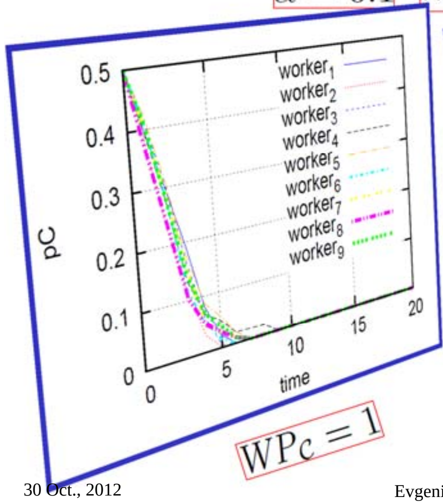

$WP_{\mathcal{C}} = 1$

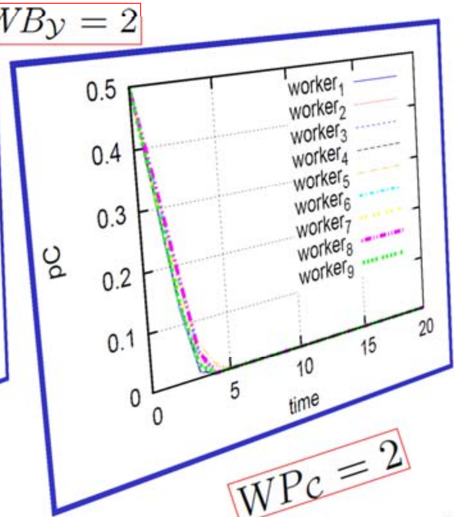

$WP_{\mathcal{C}} = 2$

Simulations

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

$\alpha = 0.1$ $a_i = 0.1$ $WC_{\mathcal{T}} = 0.1$ $WP_{\mathcal{C}} = 0$

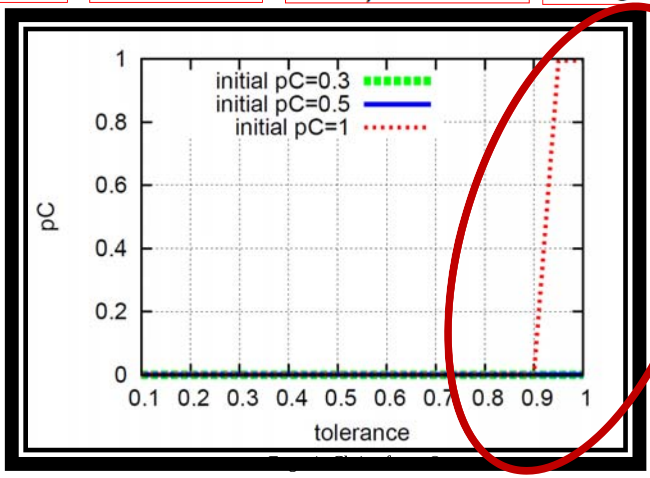

EVOLUTIONARY MECHANISM
[MALICIOUS WORKERS]

Workers with Predefined Behavior

- Workers can be:
  - Ratio…
  - Have…
    - ➢Ma…
    - ➢Altr…
- Master's … malicious …

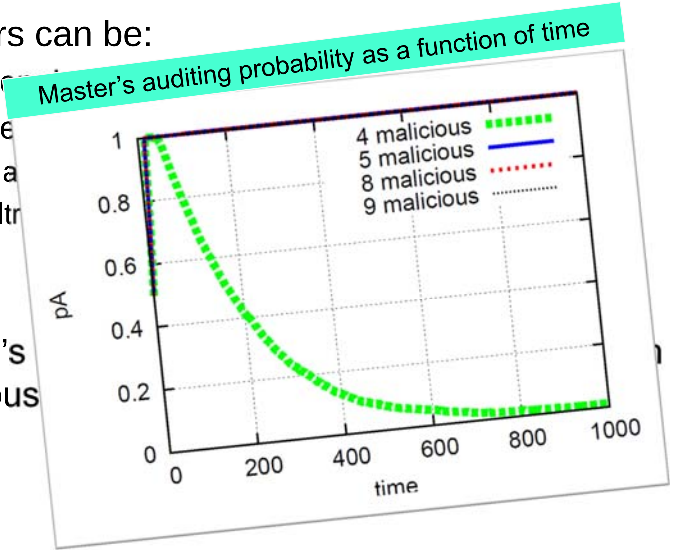

Background: Reputation

- 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
    - ➢bitTorrent: increase reputation by uploading => increase download speed (tit for tat)
    - ➢Gnutella (who to download from?)
- A reputation system can be:
  - Centralized [Sonnek et al. 2007]
  - Decentralized [Damiani et al. 2002]

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* : $correct\_audit_i(r)$
- Reputation types:
  - Type 1: $\rho_i = \frac{correct\_audit_i(r) + 1}{audit(r) + 2}$ [based on Sonnek et al. 2007]
  - Type 2: $\rho_i = \varepsilon^{audit(r) - correct\_audit_i(r)}$ where $\varepsilon \in (0,1)$

Master's Protocol

α_m: learning rate
τ: tolerance
W: set of workers
F: set of cheaters
W_j: workers in W with the same answer

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~~ workers in W_j where $j \in \{0,1\}$ and $\sum_{i \in W_j} \rho_i > \sum_{i \in W \setminus W_j} \rho_i$
Else
Update worker's *i* reputation

$$p_{\mathcal{A}} = p_{\mathcal{A}} + a_m \cdot \left(\frac{\sum_{i \in F} \rho_i}{\sum_{i \in W} \rho_i} - \tau \right)$$

Give appropriate payoff Π_i to each worker *i*

Simulations

- Choose parameters likely to be encountered:
  - 9 workers (e.g. SETI@home 3 workers)
  - initial p_{C_i} = {0.5, 1}
  - 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
  - $\alpha = \alpha_m = \alpha_w$, $\alpha = 0.1$
  - $WB_{\mathcal{Y}} = 1$ set as our normalizing
  - $WC_{\mathcal{T}} = 0.1$
  - $WP_{\mathcal{C}} = 0$

Simulations
Only Rational Workers

Auditing probability of the master as a function of time

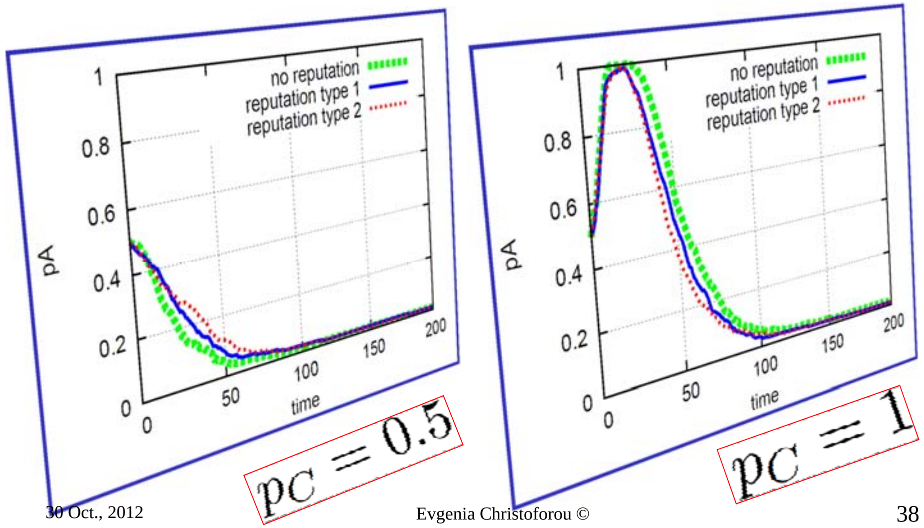

Simulations
Malicious and Rational Workers

Auditing probability of the master as a function of time
Extreme case where initial pC=1

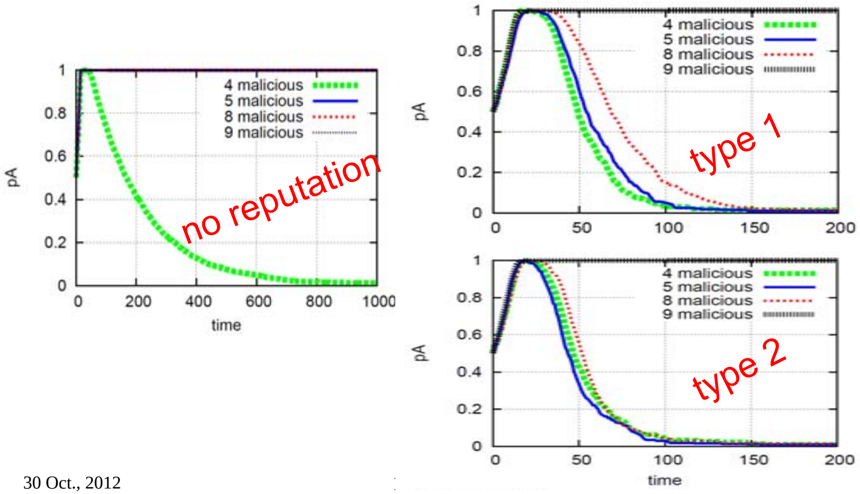

Simulations
Malicious and Altruistic Workers

Auditing probability of the master as a function of time

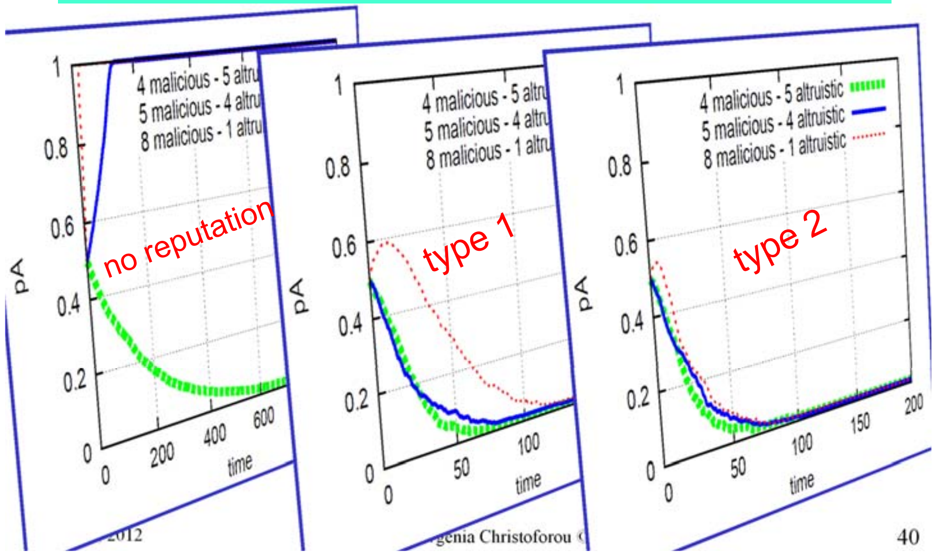

EVOLUTIONARY MECHANISM
[n CHOOSE N]

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

α_m : learning rate
τ : tolerance
W': set of n most reputable workers
F: cheaters in W'
$W'j$: workers in W' with the same answer

Set initial $p_{\mathcal{A}}$ (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 $p_{\mathcal{A}}$
    If the answers were *not* audited **then**
    Accept the value returned by workers in W'_j where $j \in \{0,1\}$ and $\sum_{i\in W'_j}\rho_i > \sum_{i\in W'\setminus W'_j}\rho_i$
    Else
    Update worker's i reputation, and

$$p_{\mathcal{A}} = p_{\mathcal{A}} + a_m \cdot \left(\frac{\sum_{i\in F'}\rho_i}{\sum_{i\in W'}\rho_i} - \tau\right)$$

  Give appropriate payoff Π_i to each worker i in W'
  If master audited **then** update the n most reputable workers

Simulations

4 malicious, 5 rational workers, reputation type 2

Workers cheating probability as a function of time

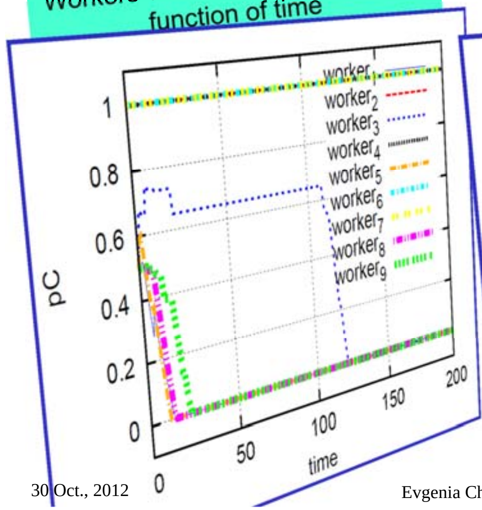

Workers reputation as a function of time

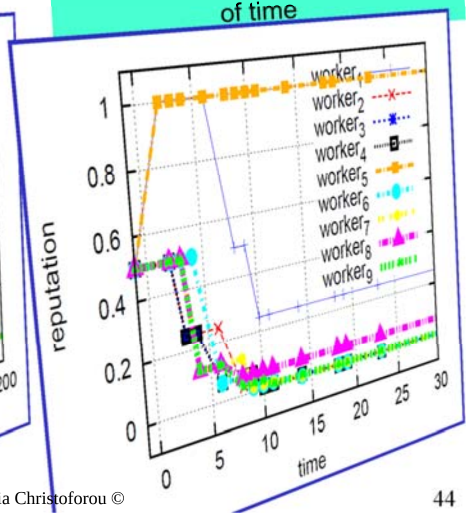

Dealing with communication uncertainty

- 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

CONCLUSIONS

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

Summary

- Use reputation mechanism to deal with malicious workers
  - Reputation not efficient in the presence of only rational workers
  - Reputation **type 2** more efficient than type 1
- Considering only replies from the most reputable workers
  - Deal with network's unreliability
  - Decreases master's cost

ΔΕΣΜΗ 2009-2010

Research Promotion Foundation

ΚΥΠΡΙΑΚΗ ΔΗΜΟΚΡΑΤΙΑ

ΤΠΕ/ΠΛΗΡΟ/0609(ΒΕ)/05

www.cs.ucy.ac.cy/ric

University of Cyprus

Universidad Rey Juan Carlos

University of Connecticut