Algorithmic Mechanisms for Reliable Master-Worker Internet-Based Computing under Communication Uncertainty

Evgenia Christoforou

Dept. of Computer Science,
University of Cyprus

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

Motivation
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, Amazons EC2
  - Volunteer Master-Worker computing: @home projects
    - e.g., SETI@home, AIDS@home, Folding@home, PrimeNet

Motivation
SETI@home

- Search for Extra Terrestrial Inteligence [1]
- Internet-based public volunteer computing project
  - Employs the BOINC software platform
  - Hosted by the Space Sciences Laboratory, at the University of California, Berkeley, USA
- Purpose: analyze radio (telescopic) signals, searching for signs of extra terrestrial intelligence
- How to use it:
  - Register your PC
  - Downloads the SETI data analyzer (screensaver mode)
  - When PC is idling, it starts analyzing data
  - When done, sends results, gets new data chunk to analyze

[1] http://setiathome.berkeley.edu/

Motivation
SETI@home

As reported in November 2009

- 278,832 active CPUs (out of a total of 2.4 million) in 234 countries
- 769 TFLOPs

⇓

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

Problem: Great potentials of Internet-based computing limited by untrustworthy platforms components

Motivation
SETI@home

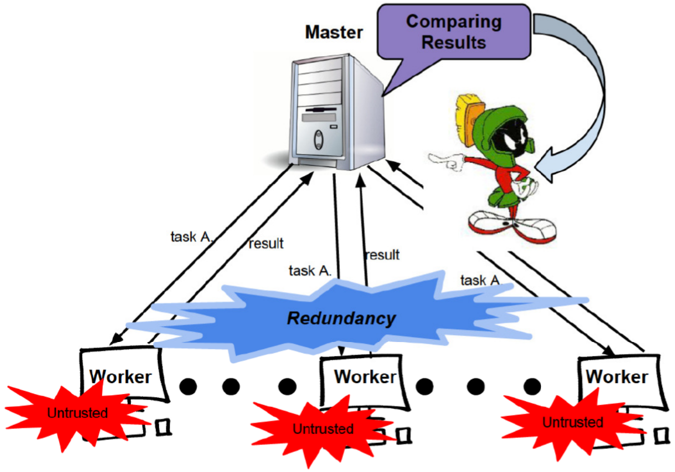

Motivation
Amazon's Mechanical Turk

- Master and worker humans
- Master processor
  - Has a problem to solve
  - Hires worker processors through the platform to compute it
- Worker processors
  - Contribute time in exchange to economic rewards

Redundant Task-Allocation

Two different approaches:

- "Classical" distributed computing (pre-defined worker behavior)
  [Sarmenta 02;Fernández et al 12; Konwar et al 06]
  - malicious workers always report incorrect result
    (sw/hw errors, Byzantine, etc.)
  - altruistic workers always compute and truthfully report result
    (the "correct" nodes)

  Malicious-tolerant voting protocols are designed
- Game-theoretic (no pre-defined worker behavior)
  [Yurkewych et al 05; Fernández Anta et al 08]
  - rational workers act selfishly maximizing own benefit

  Incentives are provided to induce a desired behavior
- BUT realistically, the three types of workers may coexist!
  [Fernández Anta et al 10]

Prior Work

In Fernandez, Georgiou and Mosteiro 10 an Internet-based master-worker framework was considered

- Consider all worker types
  - malicious: always report incorrect result
  - altruistic: always compute and report correct result
  - rational: selfishly (in a game-theoretic sense) choose to be honest or cheat
- Combine the two approaches
- **Goal:** Accept correct task with high probability, while maximizing master's utility
- A reliable network was considered

Communication Issues

- Communication uncertainty
  - Messages exchanged may get lost or arrive late
  - Around 5% of the workers are available more than 80% of the time
    Half of the workers are available less than 40% of the time
    [Heien, Anderson and Hagihara 09]
  - Long computational length incurred by a task
    [Kondo et al. 07]
- Probability the master does not receive a reply from a worker
- Allowing workers to abstain from the computation (low network reliability)

Contributions

- Develop and analyze a realistic game-theoretic mechanism
  - Unreliable communication, workers unavailability
- Designed two algorithmic mechanisms
  - Provides, when necessary, incentives to rational workers to act correctly so that
    - Master obtains correct task result (whp)
    - Despite malicious workers actions and network unreliability

  Both mechanisms are useful in different situations
- Analyzed the mechanism under two existing Internet-based Master-Worker settings
  - SETI-like volunteer computing systems
  - Profit-seeking Internet-based computational service
  - Provides clear tradeoffs between reliability, cost and network unreliability

Publications

- In Proceeding of NCA 2011

  Evgenia Christoforou, Antonio Fernandez Anta, Chryssis Georgiou and Miguel Mosteiro, Algorithmic Mechanisms for Internet Supercomputing under Unreliable Communication, in Proc. of the 10th IEEE International Symposium on Network Computing and Applications (NCA 2011), Cambridge, MA, 2011.
- Brief Announcement in Proceeding of DISC 2011

  Evgenia Christoforou, Antonio Fernandez Anta, Chryssis Georgiou and Miguel Mosteiro, Brief Announcement: Algorithmic Mechanisms for Internet-Based Computing under Unreliable Communication, in Proc. of the 25th International Symposium on Distributed Computing (DISC 2011), Rome, Italy, 2011.

Background

Definition

"A game consists of a set of players, a set of moves (or strategies) available to those players, and a specification of payoffs for each combination of strategies." [Wikipedia]

- Game Theory:
  - Players (processors) act on their self-interest
  - Rational behavior:
    seek to increase own utility choosing strategy according to payoffs
  - Protocol is given as a game
  - Design objective is to achieve equilibrium among players

Definition

Nash Equilibrium (NE): players do not increase their expected utility by changing strategy, if other players do not change [Nash 50]

Background

- Algorithmic Mechanism Design [Nisan, Ronen 01]

  Games designed to provide incentives s.t. players act "correctly"
  - Behave well: reward
  - Otherwise: penalize

  The design objective is to induce a desired behavior (e.g. unique NE)

Our approach

In this work: combine all

- Communication:
  - Unreliable network, workers may not reply
- Workers' types:
  - malicious: always report incorrect result
  - altruistic: always compute and report correct result
  - rational: selfishly choose to be honest, cheat or abstain
  - Unknown type of workers → Bayesian game [Harsanyi 67]
  - Known probability distribution over types
    p_ρ: Rational | p_μ: Malicious | p_α: Altruistic
    s.t. $p_\mu + p_\alpha + p_\rho = 1$
- Classical distributed computing approach:
  - Design malice/altruism-aware voting protocols
- Game-theoretic approach:
  - Computations modeled as strategic games

General Framework

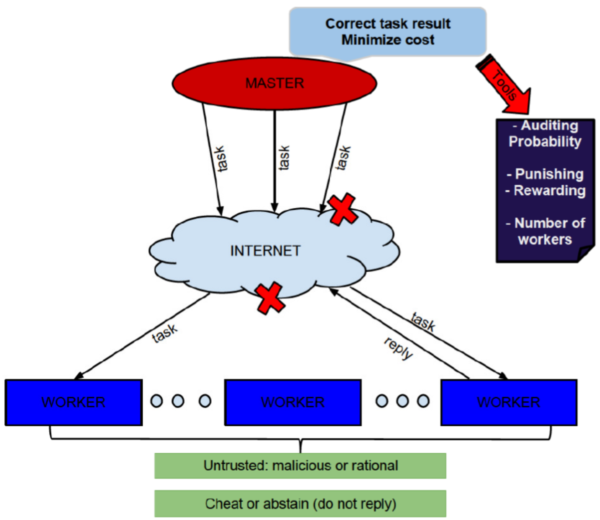

Framework

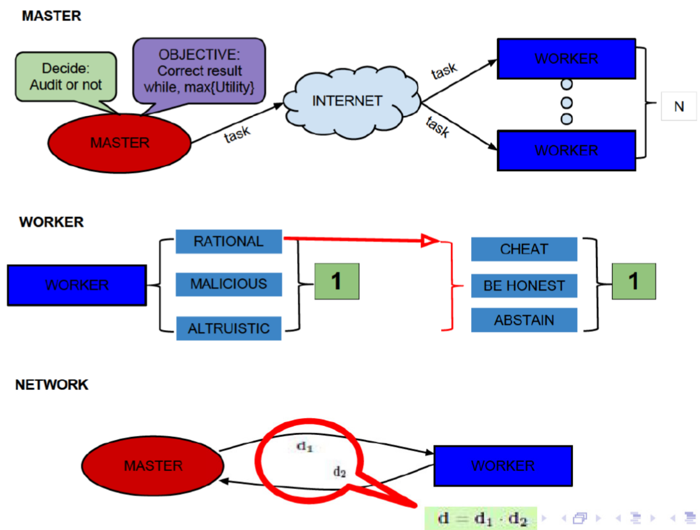

Time-based Protocol

- Master assigns a task to n workers
- Waits time T for replies
- Upon expire of time T the Master audits the responses with probability $p_{\mathcal{A}}$
- If master audits
  - rewards honest workers and
  - penalizes the cheaters
- If master does not audit
  - Accepts value returned by majority of workers
  - Rewards/penalizes according to a reward model

If by time T no replies are received, then the Master does nothing and incurs cost $MC_{\mathcal{S}}$

$\mathcal{R}_{\mathrm{m}}$	the master rewards the majority only
$\mathcal{R}_{\mathrm{a}}$	the master rewards all workers whose reply was received
$\mathcal{R}_{\emptyset}$	the master rewards no worker

Reply-based Protocol

- Master assigns a task to n workers
- If at least k replies are received then the Master audits the responses with probability $p_{\mathcal{A}}$
- If master audits
  - rewards honest workers and
  - penalizes the cheaters
- If master does not audit
  - Accepts value returned by majority of workers
  - Rewards/penalizes according to a reward model

If less than k replies are received, then the Master does nothing and incurs cost $MC_{\mathcal{S}}$

$\mathcal{R}_{\mathrm{m}}$	the master rewards the majority only
$\mathcal{R}_{\mathrm{a}}$	the master rewards all workers whose reply was received
$\mathcal{R}_{\emptyset}$	the master rewards no worker

Note: reward models may be fixed exogenously or chosen by the master

Estimating k

- For a given worker type distribution, the choice of n workers and d, even if all rational workers choose not to reply, the master receives at least

$$\boldsymbol{E} = nd(p_\alpha + p_\mu)$$

replies in expectation

- Using Chernoff bounds it follows that the master receives at least

$$k = \boldsymbol{E} - \sqrt{2\boldsymbol{E}\ln(1/\zeta)}$$

replies with probability at least $1-\zeta$ for $0<\zeta<1$ and large enough n (e.g., $\zeta = 1/n$)

Why two protocols?

- Master may have knowledge (e.g. statistics) for only one of two settings
  - Uses the mechanism designed for that setting
- Time-based mechanism, more likely to use auditing
- Reply-based mechanism may not receive enough replies
- Consequently
  - Time-based mechanism preferred when auditing cost low
  - Reply-based mechanism preferred when auditing cost high and small $MC_{\mathcal{S}}$

Payoff Parameters

$WP_{\mathcal{C}}$	worker's punishment for being caught cheating
$WC_{\mathcal{T}}$	worker's cost for computing the task
$WB_{\mathcal{Y}}$	worker's benefit from master's acceptance
$MP_{\mathcal{W}}$	master's punishment for accepting a wrong answer
$MC_{\mathcal{Y}}$	master's cost for accepting the worker's answer
$MC_{\mathcal{A}}$	master's cost for auditing worker's answers
$MC_{\mathcal{S}}$	master's cost for not getting any reply
$MB_{\mathcal{R}}$	master's benefit from accepting the right answer

Note: it is possible that $WB_{\mathcal{Y}} \neq MC_{\mathcal{Y}}$

Master's Goals

- Obtain the correct task result with a parameterized probability: $P_{succ} \geq 1 - \varepsilon$
- Then increase its utility (benefit): U_M
- Master's protocol to choose $p_{\mathcal{A}}$
  - Depending on the type distribution, the master might or might not rely on rational workers
  - The master must choose the auditing probability in such a way, to "force", when needed, the rational workers to act correctly ($p_C = 0$)
- We computed the equilibrium conditions under general payoffs values and system parameters

Conditions for mixed-strategy NE (MSNE)

Definition

For a finite game, a mixed strategy profile σ^* is a MSNE iff, for each player i

$$U_i(s_i, \sigma_{-i}) = U_i(s_i', \sigma_{-i}), \forall s_i, s_i' \in supp(\sigma_i)$$
$$U_i(s_i, \sigma_{-i}) \geq U_i(s_i', \sigma_{-i}), \forall s_i, s_i' : s_i \in supp(\sigma_i), s_i' \notin supp(\sigma_i)$$

[Osborne 2003]

s_i : strategy of player i in strategy profile s

σ_i : probability distribution over pure strategies of player i in σ

$U_i(s_i, \sigma_{-i})$: expected utility of player i using strategy s_i in σ

$supp(\sigma_i)$: set of positive-probability strategies in σ

Strategic payoffs

		$\mathcal{R}_{\mathrm{m}}$	$\mathcal{R}_{\mathrm{a}}$	$\mathcal{R}_{\emptyset}$
$\boldsymbol{w_{\mathcal{C}}}$	$w_{\mathcal{C}}^{\mathcal{AR}}$	$-WP_{\mathcal{C}}$	$-WP_{\mathcal{C}}$	$-WP_{\mathcal{C}}$
	$w_{\mathcal{C}}^{\mathcal{CR}}$	$WB_{\mathcal{Y}}$	$WB_{\mathcal{Y}}$	0
	$w_{\mathcal{C}}^{\mathcal{HR}}$	0	$WB_{\mathcal{Y}}$	0
	$w_{\mathcal{C}}^{\mathcal{X}\overline{\mathcal{R}}}$	0	0	0
$\boldsymbol{w_{\mathcal{H}}}$	$w_{\mathcal{H}}^{\mathcal{AR}}$	$WB_{\mathcal{Y}} - WC_{\mathcal{T}}$	$WB_{\mathcal{Y}} - WC_{\mathcal{T}}$	$WB_{\mathcal{Y}} - WC_{\mathcal{T}}$
	$w_{\mathcal{H}}^{\mathcal{CR}}$	$-WC_{\mathcal{T}}$	$WB_{\mathcal{Y}} - WC_{\mathcal{T}}$	$-WC_{\mathcal{T}}$
	$w_{\mathcal{H}}^{\mathcal{HR}}$	$WB_{\mathcal{Y}} - WC_{\mathcal{T}}$	$WB_{\mathcal{Y}} - WC_{\mathcal{T}}$	$-WC_{\mathcal{T}}$
	$w_{\mathcal{H}}^{\mathcal{X}\overline{\mathcal{R}}}$	$-WC_{\mathcal{T}}$	$-WC_{\mathcal{T}}$	$-WC_{\mathcal{T}}$
$\boldsymbol{w_{\mathcal{N}}}$	$w_{\mathcal{N}}^{\mathcal{XX}}$	0	0	0

Conditions for mixed-strategy NE (MSNE)

Desired condition for enforcing a unique NE at $p_{\mathcal{C}} = 0$ and $p_{\mathcal{N}} = 0$

$$\Delta U_{\mathcal{HC}} = \boldsymbol{\pi}_{\mathcal{H}} \cdot \boldsymbol{w}_{\mathcal{H}} - \boldsymbol{\pi}_{\mathcal{C}} \cdot \boldsymbol{w}_{\mathcal{C}} \geq 0$$
$$\Delta U_{\mathcal{HN}} = \boldsymbol{\pi}_{\mathcal{H}} \cdot \boldsymbol{w}_{\mathcal{H}} - \boldsymbol{\pi}_{\mathcal{N}} \cdot \boldsymbol{w}_{\mathcal{N}} \geq 0$$

$\Delta U_{S_1 S_2}$: difference on the expected utilities of a rational worker when choosing strategy S_1 over strategy S_2

$\boldsymbol{w}_X$: vector corresponding to different payoffs received by the given worker for each event when choosing strategy X

$\boldsymbol{\pi}_X$: vector corresponding to possibility that of the events occurring when the given worker choses strategy X

Analysis and Notations

Pr(worker cheats|worker replies): $q = \frac{p_\mu + p_\rho p_{\mathcal{C}}}{1 - p_\rho p_{\mathcal{N}}}$

Pr(worker does not cheat|worker replies): $\overline{q} = \frac{p_\alpha + p_\rho p_{\mathcal{H}}}{1 - p_\rho p_{\mathcal{N}}} = 1 - q$

Pr(reply received): $r = d(1 - p_\rho p_{\mathcal{N}})$

Pr(reply not received): $\overline{r} = 1 - r$

Then, $r(q + \overline{q}) + \overline{r} = 1$.

Pr(i out of n replies received): $r_i = \binom{n}{i} r^i \overline{r}^{n-i}$

Pr(majority honest | i replies received): h_i

Pr(majority cheats | i replies received): c_i

Equilibria Conditions

Guaranteeing : $P_{succ} \geq 1 - \varepsilon$ While maximizing U_M

Pr(master obtains correct answer):

$$P_{succ} = \sum_{i=k}^{n} r_i \left(p_{\mathcal{A}} + (1 - p_{\mathcal{A}}) h_i\right)$$

E(utility of master):

$$\textbf{master's utility } U_M = -\sum_{i=0}^{k-1} r_i MC_{\mathcal{S}} + \sum_{i=k}^{n} r_i (p_{\mathcal{A}} \alpha_i + (1 - p_{\mathcal{A}}) \beta_i)$$

where,

$$\alpha_i = MB_{\mathcal{R}} - MC_{\mathcal{A}} - nd(p_\alpha + p_\rho p_{\mathcal{H}}) MC_{\mathcal{Y}}$$
$$\beta_i = MB_{\mathcal{R}} h_i - MP_{\mathcal{W}} c_i - MC_{\mathcal{Y}} \gamma_i$$

and where, $\gamma_i = 0$ for $\mathcal{R}_\emptyset$, $\gamma_i = i$ for $\mathcal{R}_a$, and for $\mathcal{R}_m$ is expected number of worker's majority (as calculated in the paper)

Mechanism Design

Master protocol to chose $p_{\mathcal{A}}$

- Free rationals (master does not rely on rational workers)
  - Case 1: probability of malicious workers p_μ very large, high $p_{\mathcal{A}}$

$$p_{\mathcal{A}} \leftarrow 1 - \varepsilon / \sum_{i=k}^{n} r_i c_i$$

  - Case 2: probability of altruistic workers p_α big

$$p_{\mathcal{A}} \leftarrow 0$$

  - Case 3: rationals probability of being honest $p_{\mathcal{H}}$ is 1, even if $p_{\mathcal{A}} = 0$

$$p_{\mathcal{A}} \leftarrow 0$$

- Guided rationals (force the behavior of rational workers)
  - Rationals enforced to reply correctly ($p_{\mathcal{C}} = 0$ and $p_{\mathcal{N}} = 0$)
  - $p_{\mathcal{A}}$ is set according to worker's equilibria conditions depending on the reward model

Mechanism Design

Master protocol to chose $p_{\mathcal{A}}$ - Guided rational

- For $\mathcal{R}_{\emptyset}$, $p_{\mathcal{A}} = \frac{WC_{\mathcal{T}}}{d_2\, WB_{\mathcal{Y}} \sum_{i=k-1}^{n-1} r'_i}$, (for $p_N = 1$)
- For $\mathcal{R}_{a}$,

$$p_{\mathcal{A}} = \frac{WC_{\mathcal{T}}}{d_2(WB_{\mathcal{Y}}+WP_{\mathcal{C}}) \sum_{i=k-1}^{n-1} r'_i}, \text{ (for } p_C = 1)$$

$$d_2\, WB_{\mathcal{Y}} \sum_{i=k-1}^{n-1} r'_i \geq WC_{\mathcal{T}}, \text{ (for } p_N = 1)$$

- For $\mathcal{R}_{m}$,

$$p_{\mathcal{A}} = \frac{WC_{\mathcal{T}}/d_2 - WB_{\mathcal{Y}} \sum_{i=k-1}^{n-1} r'_i(h'_i - c'_i)}{(WB_{\mathcal{Y}}+WP_{\mathcal{C}}) \sum_{i=k-1}^{n-1} r'_i - WB_{\mathcal{Y}} \sum_{i=k-1}^{n-1} r'_i(h'_i - c'_i)}, \text{ (for } p_C = 1)$$

$$p_{\mathcal{A}} = \frac{WC_{\mathcal{T}}/d_2 - WB_{\mathcal{Y}} \sum_{i=k-1}^{n-1} r'_i h'_i}{WB_{\mathcal{Y}} \sum_{i=k-1}^{n-1} r'_i - WB_{\mathcal{Y}} \sum_{i=k-1}^{n-1} r'_i h'_i}, \text{ (for } p_N = 1)$$

SETI-like Scenario

Volunteering Computing

- each worker
  - incurs in no cost to perform the task ($WC_{\mathcal{T}} = 0$)
  - obtains a benefit ($WB_{\mathcal{Y}} > WC_{\mathcal{T}} = 0$) (recognition, prestige)
- master
  - incurs in a (possibly small) cost to reward a worker ($MC_{\mathcal{Y}} > 0$) (advertise participation)
  - may audit results at a cost ($MC_{\mathcal{A}} > 0$)
  - obtains a benefit for correct result ($MB_{\mathcal{R}} > MC_{\mathcal{Y}}$)
  - suffers a cost for wrong result ($MP_{\mathcal{W}} > MC_{\mathcal{A}}$)
- $d > 0$, as it is considered in the analysis as well

SETI-like Scenario

Volunteering Computing

Instantiating the mechanism designed on these conditions the master can choose $p_{\mathcal{A}}$ and n so that U_M is maximized for $P_{succ} \geq 1 - \varepsilon$ for any given worker-type distribution, reward model, and set of payoff parameters in the SETI scenario.

$$U_M \approx -\sum_{i=0}^{k-1} r_i MC_{\mathcal{S}} + \sum_{i=k}^{n} r_i \max\{\alpha_i, \beta_i\}$$

where $p_{\mathcal{N}} = 0$ and α_i, β_i as in the general equation of U_M.

SETI-like Scenario

Reliable Network ($d = 1$)

- Plots illustrating trade-off between reliability and cost
- Parameters' value:
  - $MC_{\mathcal{A}} = 1$, normalizing parameter
  - $MP_{\mathcal{W}} = 100$
  - Different values, don't change qualitatively the results
- 3D plots : Graphical characterization of the master's utility
  - $p_\mu \in [0, 0.5]$ ($p_\mu < 0.1$ in empirical evaluations on SETI-like system, Einstein@home, Estrada, Taufer and Anderson 09.)
  - $MC_{\mathcal{Y}} \in [0, 0.1]$, small maintenance cost of contribution list

SETI-like Scenario
Reliable Network ($d = 1$)

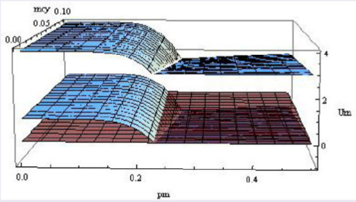

- $\mathcal{R}_\emptyset$, $n = 15$
- Upper plane $MB_\mathcal{R} = 4$, lower plane $MB_\mathcal{R} = 1$, red plane $U_M = 0$
- Master audits around $p_\mu = 0.2$

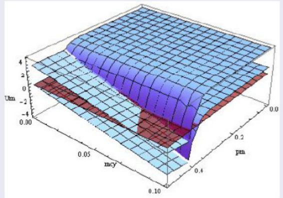

- $\mathcal{R}_\emptyset$, $n = 75$
- Upper plane $MB_\mathcal{R} = 4$, lower plane $MB_\mathcal{R} = 1$, red plane $U_M = 0$
- Master audits around $p_\mu = 0.4$

SETI-like Scenario
Unreliable Network ($d > 0$)

Time-based Mechanism

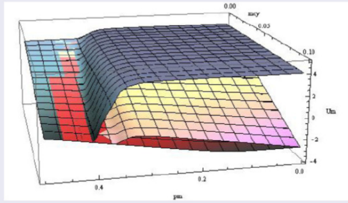

- $d = 0.9$, $n = 75$
- Upper plane $\mathcal{R}_\emptyset$, middle $\mathcal{R}_\mathrm{m}$ and lower plane $\mathcal{R}_\mathrm{a}$
- Master audits around $p_\mu = 0.35$

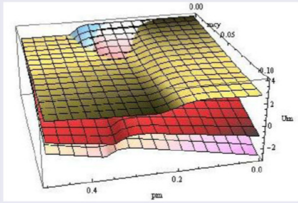

- Reward model $\mathcal{R}_\mathrm{m}$, $d = 0.9$
- Upper plane $n = 15$, middle $n = 55$, lower plane $n = 75$
- For $n = 15$, earlier change to auditing strategy

Contractor scenario

- each worker
  - receives payment for computing the task (not volunteers) ($S = WB_\mathcal{Y} = MC_\mathcal{Y}$)
  - incurs in a cost for computing ($WC_\mathcal{T} > 0$)
  - must have economic incentive ($U > 0$)
- master
  - pays each worker an amount ($MC_\mathcal{Y} > 0$)
  - receives a benefit (from consumers for the provided service) ($MB_\mathcal{R} > MC_\mathcal{Y}$)
  - may audit and has a cost for wrong result ($MP_\mathcal{W} > MC_\mathcal{A} > 0$)
- $d > 0$, as it is considered in the analysis as well

Contractor scenario

Instantiating the mechanism designed on these conditions the master can choose $p_\mathcal{A}$ and n so that U_M is maximized for $P_{succ} \geq 1 - \varepsilon$ for any given worker-type distribution, reward model, and set of payoff parameters in the Contractor scenario.

$$U_M = -\sum_{i=0}^{k-1} r_i MC_\mathcal{S} + \sum_{i=k}^{n} r_i \max\left\{\alpha_i, \beta_i + (\alpha_i - \beta_i)\frac{WC_\mathcal{T}}{d_2\, WB_\mathcal{Y} \sum_{i=k-1}^{n-1} r_i'}\right\}$$

α_i, β_i as in the general equation of U_M

Conclusions

- Reliable Master-Worker Internet-based Task computing under
  - Malicious, altruistic and rational workers
  - Communication uncertainty and worker unavailability
- When the network is considered reliable, conceptually we get Fernandez et al. algorithmic mechanism
- Mechanisms trade reliability (ε), cost (U_M) (and network unreliability)
- Realistic assumptions (unreliable network, worker abstain) :
  - Contractor scenario analysis opens the way for commercial Internet-based supercomputing where a company, given specific system parameters, could calculate its profit (if any) before agreeing into providing a proposed computational service

Ongoing and Future Work

- Consider task execution over multiple rounds over workers that their behavior changes over time
  - View the computations in the Master-Worker framework as Evolutionary Games
- Reinforcement learning
  - The Master uses knowledge gained in past rounds to
    - decrease of its probability of error in future rounds
    - increase its utility in future rounds
  - The workers use prior knowledge to increase their utility
- Worker reputation
  - Measure the workers reputation based on prior behavior and use it as an additional incentive for rational workers to act correctly.

Thank you!

Presentation available at:
http://www.cs.ucy.ac.cy/ric/dissemination.html

For further questions:
christoforou.evgenia@ucy.ac.cy