Algorithmic Mechanisms for Reliable Master-Worker Internet-based Computing

Evgenia Christoforou, Antonio Fernández Anta, Chryssis Georgiou, Miguel A. Mosteiro

www.cs.ucy.ac.cy/ric

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, 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

- Search for ExtraTerrestrial Inteligence
- 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: http://setiathome.berkeley.edu/
  - Register your PC (or your Sony PS3!)
  - 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

SETI-like Internet-based Computing

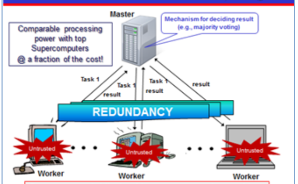

Great potential limited by untrustworthy entities

Redundant Task-Allocation

Two different approaches:

- "Classical" distributed computing: pre-defined worker behavior
  - Malicious workers fabricate and report a bogus result
  - Altruistic workers compute and truthfully report correct result

  Malicious-tolerant voting protocols are designed [Sarmenta 2002, Fernandez et al 2006, Konwar et al 2006]
- Game-theoretic: workers act upon their best interest
  - Workers are Rational, i.e., they act selfishly aiming to maximize their own benefit

  Incentives are provided to induce a desired behavior [Yurkewych et al 2005, Fernandez et al 2008]

BUT realistically, the three types of workers may coexist!

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)

Our Approach

- Communication:
  - Unreliable network, workers may not reply

Consider all worker types

- Types of workers:
  - **malicious:** always report incorrect result
  - **altruistic:** always compute and report correct result
  - **rational:** selfishly choose to be honest, cheater or abstain
- Combine the two approaches
  - Game-theoretic approach:
    - Computations modeled as strategic games
  - Classical distributed computing approach:
    - Design malice-aware voting protocols
- Objective: Reliable Internet-based Master-Worker Computing with *provable guarantees*

Framework

- Workers' types:
  - Unknown type of workers → Bayesian game [Harsanyi '67]
  - Known probability distribution over types

    p_ρ: Rational p_μ: Malicious p_α: Altruistic

    such that $p_\rho + p_\mu + p_\alpha = 1$
  - Rational has a probability to

    p_C: Cheat p_H: Honest p_N: Not reply

    such that $p_C + p_H + p_N = 1$ (if a worker decides not to reply, it does not perform the task)
- Network communication:
  - d1>0 : probability of any worker being available and receiving the task
  - d2>0 : the probability of the master receiving the worker's response (has the worker chosen to reply)
  - d=d1 · d2: probability of a round trip

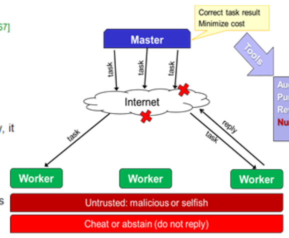

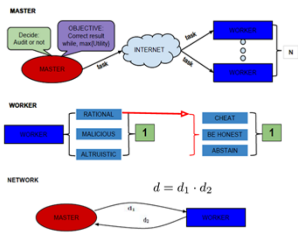

Algorithmic Mechanism (Master's Communication Protocol)

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_S

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_S

Master's Goals

Obtain the correct task result with a parameterized probability: $P_{succ} \geq 1-\varepsilon$

Then increase its utility (benefit): U_M

- Depending on the type distribution, the master might or might not rely on rational workers
- The master must choose the auditing probability $p_\mathcal{A}$ 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

Both Protocols are Useful

- Master may have knowledge (e.g., statistics) for only one of the two settings
  - Uses the protocol 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_S

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

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

replies in expectation.

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

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

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

Reward Models

$\mathcal{R}_m$	the master rewards the majority only
$\mathcal{R}_a$	the master rewards all workers
$\mathcal{R}_\emptyset$	the master does not reward any worker

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 a "sufficient" number of replies
$MB_\mathcal{R}$	master's benefit from accepting the right answer

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

Background

Terminology

A game consists of a set of players, a set of strategies available to those players, and a specification of payoffs (utilities) for each combination of strategies [wikipedia]

- Game Theory:
  - Players (processors) act on their self-interest
  - Rational [Golle Mironov 01] behavior: seek to increase their utility
  - Protocol is given as a game
  - Design objective is to achieve equilibrium among players

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

Algorithmic Mechanism Design

- Games are designed to provide necessary incentives such that rational players act "correctly"

  Reward/Punishing Schemes:
  - Behave well: Reward
  - Otherwise: Penalize
- The design objective is to induce a desired behavior (e.g., a unique NE) [Nisan Ronen 01]

Conditions for Mixed Strategy NE (MSNE)

For a finite game, a mixed strategy profile σ^* is a MSNE if and only if, for each player i: [Osborne 03]

$U_i(s_i,\sigma_{-i}) = U_i(s'_i,\sigma_{-i}), \forall s_i, s'_i \in supp(\sigma_i)$

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

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

Algorithmic Mechanism (Master's Protocol to Choose p_A)

Equilibria Conditions

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

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

$\Delta U_{S_1S_2}$: 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

Strategic Payoffs

		$\mathcal{R}_m$	$\mathcal{R}_a$	$\mathcal{R}_\emptyset$
$\boldsymbol{w}_\mathcal{C}$	$w_\mathcal{C}^{\mathcal{AR}}$	$-WP_\mathcal{C}$	$-WP_\mathcal{C}$	$-WP_\mathcal{C}$
	$w_\mathcal{C}^{\mathcal{CR}}$	$WB_\mathcal{Y}$	$WB_\mathcal{Y}$	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

Free rationals

- 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 $p_\mathcal{A} = 0$

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

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): $\widehat{q} = \frac{p_\alpha + p_\rho p_\mathcal{H}}{1 - p_\rho p_\mathcal{N}} = 1 - q$

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

Pr(reply not received from worker): $\widehat{r} = 1 - r$

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

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

Pr(majority honest | i replies received):

$h_i = \sum_{j=0}^{\lfloor i/2 \rfloor - 1} \binom{i}{j} q^j \widehat{q}^{i-j} + (1 + \lceil i/2 \rceil - \lfloor i/2 \rfloor)\frac{1}{2}\binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \widehat{q}^{\lceil i/2 \rceil}$.

Pr(majority cheats | i replies received):

$c_i = \sum_{j=\lceil i/2 \rceil + 1}^{i} \binom{i}{j} q^j \widehat{q}^{i-j} + (1 + \lceil i/2 \rceil - \lfloor i/2 \rfloor)\frac{1}{2}\binom{i}{\lceil i/2 \rceil} q^{\lceil i/2 \rceil} \widehat{q}^{\lfloor i/2 \rfloor}$

Guided rationals

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

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

- For $\mathcal{R}_a$,

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

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

- For $\mathcal{R}_m$,

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

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

$$r'_i = \binom{n-1}{i} r^i \widehat{r}^{n-1-i},$$

$$h'_i = \sum_{j=0}^{\lfloor i/2 \rfloor} \binom{i}{j} q^j \widehat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor)\frac{1}{2}\binom{i}{\lceil i/2 \rceil} q^{\lceil i/2 \rceil} \widehat{q}^{\lfloor i/2 \rfloor},$$

$$c'_i = \sum_{j=\lceil i/2 \rceil}^{i} \binom{i}{j} q^j \widehat{q}^{i-j} + (\lceil i/2 \rceil - \lfloor i/2 \rfloor)\frac{1}{2}\binom{i}{\lfloor i/2 \rfloor} q^{\lfloor i/2 \rfloor} \widehat{q}^{\lceil i/2 \rceil}$$

Putting the Mechanism into Action

Volunteering Computing (SETI-like)

- Each worker
  - Incurs in no cost to perform the task: $WC_\mathcal{T} = 0$
  - Obtains a benefit: $WB_\mathcal{Y} > 0$ (recognition, prestige – top contributors list)
- Master
  - Incurs in a (possibly small) cost to reward a worker (advertise participation): $MC_\mathcal{Y} > 0$
  - 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

Mechanism Instantiation

Instantiating the mechanism designed on these conditions the master can choose $p_\mathcal{A}$ and n so that U_M is maximized for $P_{wrong} \leq \varepsilon$ 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.

Plots

- 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 syste Einstein@home, Estrada, Taufer and Anderson 09.)
  - $MC_\mathcal{Y} \in [0, 0.1]$, small maintenance cost of contribution list

Examples (Reliable Network d=1)

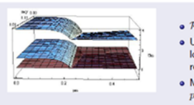

- $\mathcal{R}_\emptyset$, $n = 15$
- Upper plane $MB_\mathcal{R} = 4$, lower plane $MB_\mathcal{R} = 1$, red plane $U_M = 0$
- Master audits around $p_\mu = 0.2$

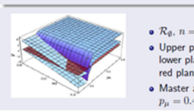

- $\mathcal{R}_\emptyset$, $n = 75$
- Upper plane $MB_\mathcal{R} = 4$, lower plane $MB_\mathcal{R} = 1$, red plane $U_M = 0$
- Master audits around $p_\mu = 0.4$

Examples (Unreliable Network 0<d<1)

Time-based Mechanism

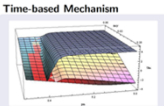

- $d = 0.9$, $n = 75$
- Upper plane $\mathcal{R}_\emptyset$, middle $\mathcal{R}_m$ and lower plane $\mathcal{R}_a$
- Master audits around $p_\mu = 0.35$

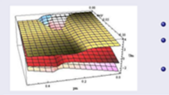

- Reward model $\mathcal{R}_m$, $d = 0.9$
- Upper plane $n = 15$, middle $n = 55$, lower plane $n = 75$
- For $n = 15$, earlier change to auditing strategy

ΔΕΣΜΗ 2009-2010

Ίδρυμα Προώθησης Έρευνας

ΚΥΠΡΙΑΚΗ ΔΗΜΟΚΡΑΤΙΑ

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