Designing Mechanisms for Reliable Internet-based Computing

Chryssis Georgiou
University of Cyprus, Cyprus

Antonio Fernandez   **Miguel Mosteiro**
Universidad Rey Juan Carlos, Spain

Trustworthy Network Computing workshop
IEEE NCA 2008
Cambridge, MA

July 12, 2008

Motivation (1)

- Demand for processing complex computational jobs
  - One-processor machines have limited computational resources
  - Powerful parallel machines are expensive
- Internet is emerging as an alternative platform for HPC
  - @home projects (e.g., SETI [Korpela Werthimer Anderson Cobb Lebofsky 01])
  - Volunteer computing
  - CPU scavenging
  - Convergence of P2P and Grid computing [Foster Iamnitchi 03]

Internet-based Computing or P2P Computing – P2PC

Motivation (2)

- Internet-based Computing
  - A machine acts as a server: Master
  - Distributes jobs, across the Internet, to client computers: Workers
  - Workers execute and report back the results
- Great potential but
  - Limited use due to *cheaters* [Kahney 01]
  - Cheater fabricates a bogus result and return it
- Possible solution: Redundant task allocation [Anderson 04, Fernandez et al 06, Konwar et al 06]
  - Master assigns same task to several workers and
  - Compares their returned results (voting)

Our Solution/Approach

- Consider Internet-based Computing from a game-theoretic point of view: Model computations as games
  - Master chooses whether to verify the returned result
  - Worker chooses whether to be honest or a cheater
- Design cost-sensitive mechanisms that incentive the workers to be honest
- Objective
  - Maximize the probability of the master for obtaining the correct result and
  - Maximize master benefit

Background (1)

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 in Distributed Computing [Halpern 08]
  - Internet routing [Koutsoupias Papadimitriou 99, Mavronicolas Spirakis 01]
  - Resource location and sharing [Halldorsson Halpern Li Mirrokni 04]
  - Containment of Viruses spreading [Moscibroda Schmid Wattenhofer 06]
  - Secret sharing [Halpern Teague 04]

Background (2)

- Traditional Distributed Computing
  - A priori behavior of processors: either good or bad

- Game Theory:
  - Processors (players) act on their self-interest
  - Rational [Golle Mironov 01] : seek to increase their utility
  - Protocol is given as a game, and the objective is to identify the *Nash equilibria* [Nash 50]

    NE: players don't increase their expected utility by choosing a different strategy, if other players don't change

Background (3)

- Algorithmic Mechanism Design [Nisan Ronen 01]
  - Games are designed to provide necessary incentives s.t. players act "correctly"
    - Behave well: Reward
    - Otherwise: Penalize
  - The design objective is to force a desired behavior (unique NE)

- Close connection between [Shneidman Parkes 03]
  - Rational players in Algorithmic Mechanism Design
  - Workers in realistic P2P (P2PC) systems

Framework (1)

- Master
  - Assigns a task to workers and collects responses
  - Can verify (audit) the values returned by the workers
    - Verification is cheaper that computing
    - The correct result might not be obtained

- Workers
  - Rational: seek to maximize their benefit
  - Honest: Computes the task and returns correct value
  - Cheater: Fabricates and returns a bogus result

Framework (2)

- We do not consider non-intentional errors produced by hardware or software problems
- Weak collusion (no Sybil attacks): Workers decide independently, but all cheaters collude in returning the same incorrect answer
- The probability of guessing the correct value (without computing the task) is negligible

General Protocol

- Master assigns a task to n workers
- Worker i cheats with probability $p_C^{(i)}$
- Master verifies the responses with probability p_V
- If master verifies
  - rewards honest workers and
  - penalizes the cheaters
- If master does not verify
  - Accepts value returned by majority of workers
  - Rewards majority (R_m), none (R_0) or all (R_a)
  - Does not penalize anyone ("In dubio pro reo")

Contributions (1)

- Identify a collection of realistic payoff parameters

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

Note that it is possible that

$$WB_{\mathcal{A}} \neq MC_{\mathcal{A}}$$

Contributions (2)

- Define the following games:
  - Between the master and a worker $1{:}1$
  - n games between the master and a worker $1{:}1^n$
  - Between n workers (master out of the game) $0{:}n$
  - Between the master and n workers $1{:}n$

  With the 3 reward models, we have 12 games in total!!

- Analysis of the 12 games under general payoff models
  - Characterize the conditions for a unique NE
  - Mechanisms that the master can run to trade cost and reliability

Contributions (3)

- Design mechanism for two specific realistic system scenarios
  - A system of volunteering computing like SETI
    - Best is single-worker allocation with game $(0{:}n,R_0)$
    - Always correct result, almost no verification, almost optimal master utility
  - Company that buys computing cycles from Internet computers and sells them to customers
    - No single optimal game
    - E.g., If only n chosen, best is single-worker allocation with game $(0{:}n,R_a)$ or $(0{:}n,R_0)$

GAMES

Game Definition

$W = \{1, 2, \ldots, n\}$	set of assigned workers
M	master processor
$\mathcal{S}_i$	set of pure strategies available to player i
$\{\mathcal{C}, \overline{\mathcal{C}}\}$	set of pure strategies of a worker
$\{\mathcal{V}, \overline{\mathcal{V}}\}$	set of pure strategies of the master
s	strategy profile (a mapping from players to pure strategies)
s_i	strategy used by player i in the strategy profile s
s_{-i}	strategy used by each player but i in the strategy profile s
$w_s^{(i)}$	payoff of worker i for the strategy profile s
m_s	payoff of the master for the strategy profile s
$p_{s_i}^{(i)}$	probability that worker i uses strategy s_i
p_{s_M}	probability that the master uses strategy s_M
σ	mixed strategy profile
σ_i	probability distribution over pure strategies used by player i in σ
σ_{-i}	probability distribution over pure strategies by all players but i in σ
$U_i(s_i, \sigma_{-i})$	expected utility of worker i with mixed strategy profile σ
$U_M(s_M, \sigma_{-M})$	expected utility of master with mixed strategy profile σ
$supp(\sigma_i)$	set of strategies of player i with probability > 0 in σ

Equilibrium Definition

- For a finite game, a mixed strategy profile σ^* is a mixed-strategy Nash Equilibrium (MSNE), iff, for each player *i*

$$U_i(s_i, \sigma^*_{-i}) = U_i(s_j, \sigma^*_{-i}), \forall s_i, s_j \in supp(\sigma^*_i),$$
$$U_i(s_i, \sigma^*_{-i}) \geq U_i(s_k, \sigma^*_{-i}),$$
$$\forall s_i, s_k : s_i \in supp(\sigma^*_i), s_k \notin supp(\sigma^*_i).$$

A strategy choice by each game participant s.t none has incentive to change it.

Methodology

- For each game in $\{1{:}1,\ 1{:}1^n,\ 0{:}n,\ 1{:}n\}$
  - Identify the parameter conditions for which there is a MSNE
    - Instantiate the two equations of the MSNE definition
    - Assume a general payoff model
  - From the above, obtain the conditions on the parameters (payoffs and probabilities) that make such a MSNE unique
  - Plug the specific reward models (R_m, R_0, R_a) on the conditions to obtain the trade-offs between cost and reliability

Game *1:1*, 1 master - 1 worker

Expected utility of the master in any equilibrium

$$U_M = p_{\mathcal{C}} p_{\mathcal{V}} m_{\mathcal{C}\mathcal{V}} + (1-p_{\mathcal{C}}) p_{\mathcal{V}} m_{\overline{\mathcal{C}}\mathcal{V}} + p_{\mathcal{C}} (1-p_{\mathcal{V}}) m_{\mathcal{C}\overline{\mathcal{V}}} + (1-p_{\mathcal{C}})(1-p_{\mathcal{V}}) m_{\overline{\mathcal{C}}\overline{\mathcal{V}}}$$

Expected utility of the worker in any equilibrium

$$U_W = p_{\mathcal{C}} p_{\mathcal{V}} w_{\mathcal{C}\mathcal{V}} + p_{\mathcal{C}} (1-p_{\mathcal{V}}) w_{\mathcal{C}\overline{\mathcal{V}}} + (1-p_{\mathcal{C}}) p_{\mathcal{V}} w_{\overline{\mathcal{C}}\mathcal{V}} + (1-p_{\mathcal{C}})(1-p_{\mathcal{V}}) w_{\overline{\mathcal{C}}\overline{\mathcal{V}}}$$

Probability of master accepting a wrong value:

$$\mathbf{P}_{wrong} = (1-p_{\mathcal{V}}) p_{\mathcal{C}}$$

Game *1:1*, 1 master - 1 worker

- Depending on the range of values that p_C and p_V take, we may have a MSNE or a pure NE
- Both p_C and p_V can take values either 0, 1, or in (0,1) (9 cases)
- Example

If $p_{\mathcal{C}} \in (0,1), p_{\mathcal{V}} \in (0,1)$, there is a a MSNE if, simultaneously,

$$U_M(\mathcal{V}, p_{\mathcal{C}}) = U_M(\overline{\mathcal{V}}, p_{\mathcal{C}}) \quad U_W(\mathcal{C}, p_{\mathcal{V}}) = U_W(\overline{\mathcal{C}}, p_{\mathcal{V}})$$

and hence

$$p_{\mathcal{C}} = \frac{m_{\overline{\mathcal{C}\mathcal{V}}} - m_{\overline{\mathcal{C}}\mathcal{V}}}{m_{\mathcal{C}\mathcal{V}} - m_{\overline{\mathcal{C}}\mathcal{V}} - m_{\mathcal{C}\overline{\mathcal{V}}} + m_{\overline{\mathcal{C}\mathcal{V}}}} \quad p_{\mathcal{V}} = \frac{w_{\overline{\mathcal{C}\mathcal{V}}} - w_{\mathcal{C}\overline{\mathcal{V}}}}{w_{\mathcal{C}\mathcal{V}} - w_{\mathcal{C}\overline{\mathcal{V}}} - w_{\overline{\mathcal{C}}\mathcal{V}} + w_{\overline{\mathcal{C}\mathcal{V}}}}$$

Reward Models

- R_m: Rewards only majority

$$
\begin{aligned}
m_{\mathcal{C}\mathcal{V}} &= -MC_{\mathcal{V}} & w_{\mathcal{C}\mathcal{V}} &= -WP_{\mathcal{C}} \\
m_{\bar{\mathcal{C}}\mathcal{V}} &= MB_{\mathcal{R}} - MC_{\mathcal{V}} - MC_{\mathcal{A}} & w_{\bar{\mathcal{C}}\mathcal{V}} &= WB_{\mathcal{A}} - WC_{\mathcal{T}} \\
m_{\mathcal{C}\bar{\mathcal{V}}} &= -MP_{\mathcal{W}} - MC_{\mathcal{A}} & w_{\mathcal{C}\bar{\mathcal{V}}} &= WB_{\mathcal{A}} \\
m_{\overline{\mathcal{C}\mathcal{V}}} &= MB_{\mathcal{R}} - MC_{\mathcal{A}} & w_{\overline{\mathcal{C}\mathcal{V}}} &= WB_{\mathcal{A}} - WC_{\mathcal{T}}
\end{aligned}
$$

- R_a: Rewards all workers (same as above)
- R_0: Does not reward any worker

$$
\begin{aligned}
m_{\mathcal{C}\mathcal{V}} &= -MC_{\mathcal{V}} & w_{\mathcal{C}\mathcal{V}} &= -WP_{\mathcal{C}} \\
m_{\bar{\mathcal{C}}\mathcal{V}} &= MB_{\mathcal{R}} - MC_{\mathcal{V}} - MC_{\mathcal{A}} & w_{\bar{\mathcal{C}}\mathcal{V}} &= WB_{\mathcal{A}} - WC_{\mathcal{T}} \\
m_{\mathcal{C}\bar{\mathcal{V}}} &= -MP_{\mathcal{W}} & w_{\mathcal{C}\bar{\mathcal{V}}} &= 0 \\
m_{\overline{\mathcal{C}\mathcal{V}}} &= MB_{\mathcal{R}} & w_{\overline{\mathcal{C}\mathcal{V}}} &= -WC_{\mathcal{T}}
\end{aligned}
$$

Game $1{:}1^n$, n Games 1:1

- Master runs n instances of Game $1{:}1$, one with each of the n workers
- Chooses to verify or not with prob p_V only once
- If no verification, rewards all or none like in the $1{:}1$ game
- Key difference now is that redundancy can be used to reduce the prob. of accepting a wrong value.

$$\mathbf{P}_{\mathcal{C}} = \sum_{\substack{(F,T)\in\mathcal{W}\\|F|>|T|}} \prod_{j\in F} p_{\mathcal{C}}^{(j)} \prod_{k\in T} (1 - p_{\mathcal{C}}^{(k)})$$

$$\mathbf{P}_{wrong} = (1 - p_{\mathcal{V}})\mathbf{P}_{\mathcal{C}}$$

Game $1{:}1^n$, R_m,R_a (Game $1{:}1$ for $n=1$)

Equilibrium $p_\mathcal{C}, p_\mathcal{V}$	Conditions	$\mathbf{P}_{wrong}$	U_M	U_W
$\frac{MC_\mathcal{V}}{MC_\mathcal{A}+MP_\mathcal{W}}, \frac{WC_\mathcal{T}}{WB_\mathcal{A}+WP_\mathcal{C}}$		$(1-p_\mathcal{V})\mathbf{P}_\mathcal{C}$	$p_\mathcal{V}((1-p_\mathcal{C}^n)MB_\mathcal{R}-$ $MC_\mathcal{V}-(1-p_\mathcal{C})nMC_\mathcal{A})+$ $(1-p_\mathcal{V})(MB_\mathcal{R}(1-\mathbf{P}_\mathcal{C})-$ $MP_\mathcal{W}\mathbf{P}_\mathcal{C}-nMC_\mathcal{A})$	$WB_\mathcal{A}-WC_\mathcal{T}$
$0,\ \frac{WC_\mathcal{T}}{WB_\mathcal{A}+WP_\mathcal{C}} \le p_\mathcal{V} < 1$ $0 < p_\mathcal{V}$	$MC_\mathcal{V}=0$	0	$MB_\mathcal{R}-nMC_\mathcal{A}$	$WB_\mathcal{A}-WC_\mathcal{T}$
$1,\ 0 < p_\mathcal{V} \le \frac{WC_\mathcal{T}}{WB_\mathcal{A}+WP_\mathcal{C}}$ $p_\mathcal{V} < 1$	$MC_\mathcal{V}=MP_\mathcal{W}+MC_\mathcal{A}$	$1-p_\mathcal{V}$	$-p_\mathcal{V}MC_\mathcal{V}-(1-p_\mathcal{V})(MP_\mathcal{W}+nMC_\mathcal{A})$	$(1-p_\mathcal{V})WB_\mathcal{A}-$ $p_\mathcal{V}WP_\mathcal{C}$
$0 \le p_\mathcal{C} \le \frac{MC_\mathcal{V}}{MC_\mathcal{A}+MP_\mathcal{W}}$, 0 $p_\mathcal{C} < 1$	$WC_\mathcal{T}=0$	$\mathbf{P}_\mathcal{C}$	$MB_\mathcal{R}(1-\mathbf{P}_\mathcal{C})-MP_\mathcal{W}\mathbf{P}_\mathcal{C}-nMC_\mathcal{A}$	$WB_\mathcal{A}$
$\frac{MC_\mathcal{V}}{MC_\mathcal{A}+MP_\mathcal{W}} \le p_\mathcal{C} < 1$, 1 $0 < p_\mathcal{C}$	$WC_\mathcal{T}=WB_\mathcal{A}+WP_\mathcal{C}$	0	$(1-\prod_{j\in W} p_\mathcal{C}^{(j)})MB_\mathcal{R}-MC_\mathcal{V}-$ $\sum_{(W_F,W_T)\in\mathcal{W}} \prod_{j\in W_F} p_\mathcal{C}^{(j)}\cdot$ $\prod_{k\in W_T}(1-p_\mathcal{C}^{(k)})\lvert W_T\rvert MC_\mathcal{A}$	$-WP_\mathcal{C}$
1, 1	$MC_\mathcal{V} \le MP_\mathcal{W}+MC_\mathcal{A}$ $WC_\mathcal{T} \ge WB_\mathcal{A}+WP_\mathcal{C}$	0	$-MC_\mathcal{V}$	$-WP_\mathcal{C}$
0, 1	$MC_\mathcal{V}=0$ $WC_\mathcal{T} \le WB_\mathcal{A}+WP_\mathcal{C}$	0	$MB_\mathcal{R}-nMC_\mathcal{A}$	$WB_\mathcal{A}-WC_\mathcal{T}$
1, 0	$MC_\mathcal{V} \ge MP_\mathcal{W}+MC_\mathcal{A}$	1	$-MP_\mathcal{W}-nMC_\mathcal{A}$	$WB_\mathcal{A}$

Game $1{:}1^n$, R_0 (Game $1{:}1$ for $n=1$)

Equilibrium $p_\mathcal{C}, p_\mathcal{V}$	Conditions	$\mathbf{P}_{wrong}$	U_M	U_W
$\frac{MC_\mathcal{V}+MC_\mathcal{A}}{MC_\mathcal{A}+MP_\mathcal{W}}, \frac{WC_\mathcal{T}}{WB_\mathcal{A}+WP_\mathcal{C}}$		$(1-p_\mathcal{V})\mathbf{P}_\mathcal{C}$	$p_\mathcal{V}((1-p_\mathcal{C}^n)MB_\mathcal{R} - MC_\mathcal{V} - (1-p_\mathcal{C})nMC_\mathcal{A}) + (1-p_\mathcal{V})(MB_\mathcal{R}(1-\mathbf{P}_\mathcal{C}) - MP_\mathcal{W}\mathbf{P}_\mathcal{C})$	$-p_\mathcal{V} WP_\mathcal{C}$
$0,\ \frac{WC_\mathcal{T}}{WB_\mathcal{A}+WP_\mathcal{C}} \le p_\mathcal{V} < 1$, $0 < p_\mathcal{V}$	$MC_\mathcal{A} = MC_\mathcal{V} = 0$	0	$MB_\mathcal{R}$	$p_\mathcal{V} WB_\mathcal{A} - WC_\mathcal{T}$
$1,\ 0 < p_\mathcal{V} \le \frac{WC_\mathcal{T}}{WB_\mathcal{A}+WP_\mathcal{C}}$, $p_\mathcal{V} < 1$	$MC_\mathcal{V} = MP_\mathcal{W}$	$1-p_\mathcal{V}$	$-MC_\mathcal{V}$	$-p_\mathcal{V} WP_\mathcal{C}$
$0 \le p_\mathcal{C} \le \frac{MC_\mathcal{V}+MC_\mathcal{A}}{MC_\mathcal{A}+MP_\mathcal{W}}$, $p_\mathcal{C} < 1$, 0	$WC_\mathcal{T} = 0$	$\mathbf{P}_\mathcal{C}$	$MB_\mathcal{R}(1-\mathbf{P}_\mathcal{C}) - MP_\mathcal{W}\mathbf{P}_\mathcal{C}$	0
$\frac{MC_\mathcal{V}+MC_\mathcal{A}}{MC_\mathcal{A}+MP_\mathcal{W}} \le p_\mathcal{C} < 1$, $0 < p_\mathcal{C}$, 1	$WC_\mathcal{T} = WB_\mathcal{A} + WP_\mathcal{C}$	0	$(1-\prod_{j\in W} p_\mathcal{C}^{(j)})MB_\mathcal{R} - MC_\mathcal{V} - \sum_{(W_F,W_T)\in\mathcal{W}} \prod_{j\in W_F} p_\mathcal{C}^{(j)} \cdot \prod_{k\in W_T}(1-p_\mathcal{C}^{(k)})\lvert W_T\rvert MC_\mathcal{A}$	$-WP_\mathcal{C}$
1, 1	$MC_\mathcal{V} \le MP_\mathcal{W}$, $WC_\mathcal{T} \ge WB_\mathcal{A} + WP_\mathcal{C}$	0	$-MC_\mathcal{V}$	$-WP_\mathcal{C}$
0, 1	$MC_\mathcal{V} = MC_\mathcal{A} = 0$, $WC_\mathcal{T} \le WB_\mathcal{A} + WP_\mathcal{C}$	0	$MB_\mathcal{R}$	$WB_\mathcal{A} - WC_\mathcal{T}$
1, 0	$MC_\mathcal{V} \ge MP_\mathcal{W}$	1	$-MP_\mathcal{W}$	0

Games $0{:}n$ and $1{:}n$

- In Game $0{:}n$ the master participates only indirectly (by fixing p_V)
  - Under our assumptions, we prove that there are only pure NE
  - We force the unique equilibrium in which no worker cheats
- Game $1{:}n$ is similar to $0{:}n$, but
  - The master is part of the game
  - This imposes additional restrictions on the NE
  - The unique NE in which no worker cheats requires $MC_V=0$

Mechanism Design

SETI-like Scenario

- Here we assume
  - $WB_A > WC_T = 0$ (CPU scavenging, worker's incentive)
  - $MB_R > MC_A > 0$ (master's incentive to compute)
  - $MP_W > MC_V > 0$ (master's incentive to verify)

- Under these constraints,
  - In both Games $1{:}1$ and $1{:}1^n$ one single MSNE remains
  - There is no unique NE for Game $1{:}n$

Results for SETI-like

(Game,Model)	Equilibrium $p_{\mathcal{C}}, p_{\mathcal{V}}$	$\mathbf{P}_{wrong}$	U_M	U_W
$(1{:}1,\mathcal{R}_\mathrm{m}),(1{:}1,\mathcal{R}_\mathrm{a})$	$0 \le p_{\mathcal{C}} \le \frac{MC_{\mathcal{V}}}{MC_{\mathcal{A}}+MP_{\mathcal{W}}}, p_{\mathcal{C}} < 1 \;, p_{\mathcal{V}} = 0$	$p_{\mathcal{C}}$	$MB_{\mathcal{R}} - p_{\mathcal{C}}(MB_{\mathcal{R}} + MP_{\mathcal{W}}) - MC_{\mathcal{A}}$	$WB_{\mathcal{A}}$
$(1{:}1,\mathcal{R}_\emptyset)$	$0 \le p_{\mathcal{C}} \le \frac{MC_{\mathcal{V}}+MC_{\mathcal{A}}}{MC_{\mathcal{A}}+MP_{\mathcal{W}}}, p_{\mathcal{C}} < 1 \;, p_{\mathcal{V}} = 0$	$p_{\mathcal{C}}$	$MB_{\mathcal{R}} - p_{\mathcal{C}}(MB_{\mathcal{R}} + MP_{\mathcal{W}})$	0
$(1{:}1^n,\mathcal{R}_\mathrm{m}),(1{:}1^n,\mathcal{R}_\mathrm{a})$	$0 \le p_{\mathcal{C}} \le \frac{MC_{\mathcal{V}}}{MC_{\mathcal{A}}+MP_{\mathcal{W}}}, p_{\mathcal{C}} < 1 \;, p_{\mathcal{V}} = 0$	$\mathbf{P}_{\mathcal{C}}$	$MB_{\mathcal{R}} - \mathbf{P}_{\mathcal{C}}(MB_{\mathcal{R}} + MP_{\mathcal{W}}) - nMC_{\mathcal{A}}$	$WB_{\mathcal{A}}$
$(1{:}1^n,\mathcal{R}_\emptyset)$	$0 \le p_{\mathcal{C}} \le \frac{MC_{\mathcal{V}}+MC_{\mathcal{A}}}{MC_{\mathcal{A}}+MP_{\mathcal{W}}}, p_{\mathcal{C}} < 1 \;, p_{\mathcal{V}} = 0$	$\mathbf{P}_{\mathcal{C}}$	$MB_{\mathcal{R}} - \mathbf{P}_{\mathcal{C}}(MB_{\mathcal{R}} + MP_{\mathcal{W}})$	0
$(0{:}n,\mathcal{R}_\mathrm{m})$	$p_{\mathcal{C}} = 0,\; \frac{WB_{\mathcal{A}}}{WP_{\mathcal{C}}+2WB_{\mathcal{A}}} < p_{\mathcal{V}} \le 1$	0	$MB_{\mathcal{R}} - p_{\mathcal{V}}MC_{\mathcal{V}} - nMC_{\mathcal{A}}$	$WB_{\mathcal{A}}$
$(0{:}n,\mathcal{R}_\mathrm{a})$	$p_{\mathcal{C}} = 0,\; 0 < p_{\mathcal{V}} \le 1$	0	$MB_{\mathcal{R}} - p_{\mathcal{V}}MC_{\mathcal{V}} - nMC_{\mathcal{A}}$	$WB_{\mathcal{A}}$
$(0{:}n,\mathcal{R}_\emptyset)$	$p_{\mathcal{C}} = 0,\; 0 < p_{\mathcal{V}} \le 1$	0	$MB_{\mathcal{R}} - p_{\mathcal{V}}(MC_{\mathcal{V}} + nMC_{\mathcal{A}})$	$p_{\mathcal{V}}WB_{\mathcal{A}}$

- The best choice is Game $0{:}n$ with reward model R_0 (no verification, no payment) and non-redundant allocation
- To obtain always the correct answer it is enough for the master to verify with arbitrarily small probability
- The master´s utility is arbitrarily close to optimal

Contractor Scenario

- A company buys computational power from Internet users and sells it to computation-needing consumers (commercial P2PC)
- Workers must have incentive to participate: $U_W > 0$
- Additionally
  - $WB_A = MC_A$ (worker's incentive)
  - $MB_R > MC_A$ (master's incentive to compute)
  - $WC_T > 0$ (computing cost)
  - $MP_W > MC_V > 0$ (master's incentive to verify)

Results for Contractor

(Game,Model)	Equilibrium $p_{\mathcal{C}}, p_{\mathcal{V}}$	$\mathbf{P}_{wrong}$	U_M	U_W
$(1{:}1, \mathcal{R}_{\mathrm{m}}), (1{:}1, \mathcal{R}_{\mathrm{a}})$	$\frac{MC_{\mathcal{V}}}{MC_{\mathcal{A}}+MP_{\mathcal{W}}}, \frac{WC_{\mathcal{T}}}{WB_{\mathcal{A}}+WP_{\mathcal{C}}}$	$(1-p_{\mathcal{V}})p_{\mathcal{C}}$	$MB_{\mathcal{R}} - p_{\mathcal{C}}(MB_{\mathcal{R}} + MP_{\mathcal{W}}) - MC_{\mathcal{A}}$	$WB_{\mathcal{A}} - WC_{\mathcal{T}}$
$(1{:}1, \mathcal{R}_{\emptyset})$	$\frac{MC_{\mathcal{V}}+MC_{\mathcal{A}}}{MC_{\mathcal{A}}+MP_{\mathcal{W}}}, \frac{WC_{\mathcal{T}}}{WB_{\mathcal{A}}+WP_{\mathcal{C}}}$	$(1-p_{\mathcal{V}})p_{\mathcal{C}}$	$MB_{\mathcal{R}} - p_{\mathcal{C}}(MB_{\mathcal{R}} + MP_{\mathcal{W}})$	$-p_{\mathcal{V}} WP_{\mathcal{C}}$
$(1{:}1^n, \mathcal{R}_{\mathrm{m}}), (1{:}1^n, \mathcal{R}_{\mathrm{a}})$	$\frac{MC_{\mathcal{V}}}{MC_{\mathcal{A}}+MP_{\mathcal{W}}}, \frac{WC_{\mathcal{T}}}{WB_{\mathcal{A}}+WP_{\mathcal{C}}}$	$(1-p_{\mathcal{V}})\mathbf{P}_{\mathcal{C}}$	$(p_{\mathcal{V}}(1-p_{\mathcal{C}}^n) + (1-p_{\mathcal{V}})(1-\mathbf{P}_{\mathcal{C}}))MB_{\mathcal{R}}$ $-p_{\mathcal{V}}MC_{\mathcal{V}} - (1-p_{\mathcal{V}})\mathbf{P}_{\mathcal{C}}MP_{\mathcal{W}}$ $-(1-p_{\mathcal{V}}p_{\mathcal{C}})nMC_{\mathcal{A}}$	$WB_{\mathcal{A}} - WC_{\mathcal{T}}$
$(1{:}1^n, \mathcal{R}_{\emptyset})$	$\frac{MC_{\mathcal{V}}+MC_{\mathcal{A}}}{MC_{\mathcal{A}}+MP_{\mathcal{W}}}, \frac{WC_{\mathcal{T}}}{WB_{\mathcal{A}}+WP_{\mathcal{C}}}$	$(1-p_{\mathcal{V}})\mathbf{P}_{\mathcal{C}}$	$(p_{\mathcal{V}}(1-p_{\mathcal{C}}^n) + (1-p_{\mathcal{V}})(1-\mathbf{P}_{\mathcal{C}}))MB_{\mathcal{R}}$ $-p_{\mathcal{V}}MC_{\mathcal{V}} - (1-p_{\mathcal{V}})\mathbf{P}_{\mathcal{C}}MP_{\mathcal{W}}$ $-p_{\mathcal{V}}(1-p_{\mathcal{C}})nMC_{\mathcal{A}}$	$-p_{\mathcal{V}} WP_{\mathcal{C}}$
$(0{:}n, \mathcal{R}_{\mathrm{m}})$	$0, \frac{WB_{\mathcal{A}}+WC_{\mathcal{T}}}{WP_{\mathcal{C}}+2WB_{\mathcal{A}}} < p_{\mathcal{V}} \leq 1$	0	$MB_{\mathcal{R}} - p_{\mathcal{V}}MC_{\mathcal{V}} - nMC_{\mathcal{A}}$	$WB_{\mathcal{A}} - WC_{\mathcal{T}}$
$(0{:}n, \mathcal{R}_{\mathrm{a}})$	$0, \frac{WC_{\mathcal{T}}}{WP_{\mathcal{C}}+WB_{\mathcal{A}}} < p_{\mathcal{V}} \leq 1$	0	$MB_{\mathcal{R}} - p_{\mathcal{V}}MC_{\mathcal{V}} - nMC_{\mathcal{A}}$	$WB_{\mathcal{A}} - WC_{\mathcal{T}}$
$(0{:}n, \mathcal{R}_{\emptyset})$	$0, \frac{WC_{\mathcal{T}}}{WP_{\mathcal{C}}+WB_{\mathcal{A}}} < p_{\mathcal{V}} \leq 1$	0	$MB_{\mathcal{R}} - p_{\mathcal{V}}(MC_{\mathcal{V}} + nMC_{\mathcal{A}})$	$p_{\mathcal{V}} WB_{\mathcal{A}} - WC_{\mathcal{T}}$

- To obtain always the correct answer, verification probability $p_V > WC_T/(WP_C+WB_A)$ can be large
- No single optimal game
  - Either Game $(0{:}n,R_a)$ or $(0{:}n,R_0)$ is the best if only $n(=1)$ or WP_C can be changed
  - If master can only change $WB_A = MC_A$, sometimes Game $(1{:}1^n,R_m)$ is the best

Summary

- We considered master-worker Internet-based computations from a game-theoretic point of view
- We defined the general model and cost-parameters
- Proposed and analyzed several games that the master can choose to play to achieve high reliability at low cost
- Designed appropriate mechanisms for two realistic scenarios
- No optimal game for every scenario or set of parameters.

Future Work

- Consider other realistic scenarios where our work can be applied
- Consider other forms of collusion (e.g., Sybil attacks) and investigate how the trade-offs are affected
- Remove the complete knowledge assumption (implicit in NE analyses)
- Study games with several rounds (reputation)
- Consider irrational, bounded-rational, and faulty players.

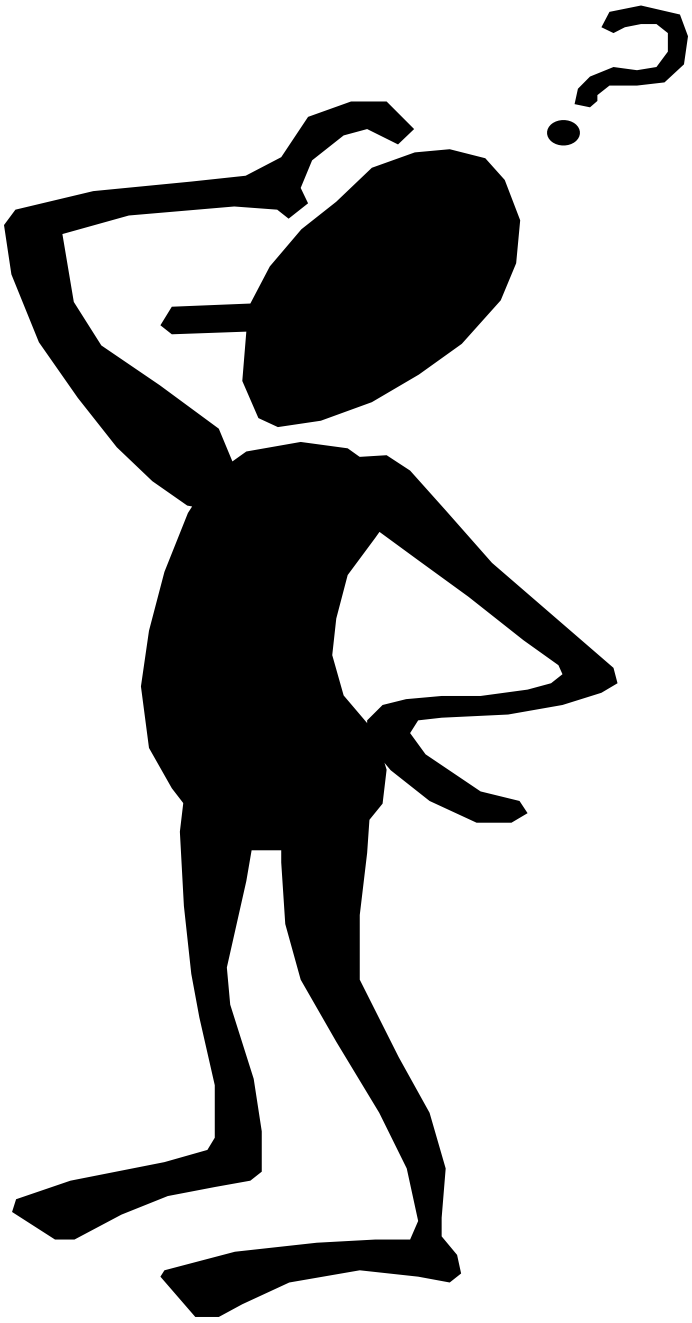

Thank you!

IAG vs Our Work

- Closer work to ours [Yurkewych Levine Rosenberg 05]
  - Master audits the results returned by rational workers with a tunable probability
  - Bounds for that probability are computed to provide incentives to workers to be honest in three scenarios
    - Redundant allocation with collusion
      - Cooperation among workers concealed from the master
    - Redundant allocation without collusion
    - Single-worker allocation
  - Their conclusion: Single-worker allocation is a cost-effective mechanism, especially in the presence of collusion

- Our model comprises a weaker type of collusion but
  - We study more algorithms and games
  - Consider a richer payoff model and probabilistic cheating
  - We have one-round protocols and
  - Show useful trade-offs between the benefit (cost) of the master and the probability of accepting a wrong result (reliability)

Prior/Related Work

- Internet Auditing Game [Yurkewych Levine Rosenberg 05]
  - Three master-worker scenarios
    - Redundant allocation with/without collusion
      (Cooperation among workers concealed from the master)
    - Single-worker allocation
  - Master (out of the game) audits the results or accepts majority
  - (Fixed) probability p_V of auditing
  - (Fixed) payments R for accepted results
  - (Fixed) penalty P for rejected results
  - Result: Bounds on R, P, and p_V to prevent cheating
  - Conclusions:
    - Lots of auditing ($p_V > 1/2$) with collusion
    - Redundancy only useful if no collusion
    - Single-worker allocation is cost-effective