Algorithmic Mechanisms for Internet Supercomputing under Unreliable Communication

Evgenia Christoforou[1] Antonio Fernández Anta[2,3]
Chryssis Georgiou[1] Miguel A. Mosteiro[3,4]

[1]Dept. of Computer Science, University of Cyprus

[2]Institute IMDEA Networks

[3]LADyR, GSyC, Universidad Rey Juan Carlos

[4]Dept. of Computer Science, Rutgers University

NCA 2011

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

Motivation

- Internet emerges as a viable platform for supercomputing
  - @home systems, volunteering computing
    (e.g., SETI@home [Korpela et al 01])
  - P2P and Grid computing [Foster, Iamnitchi 03]

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

Motivation

Motivation

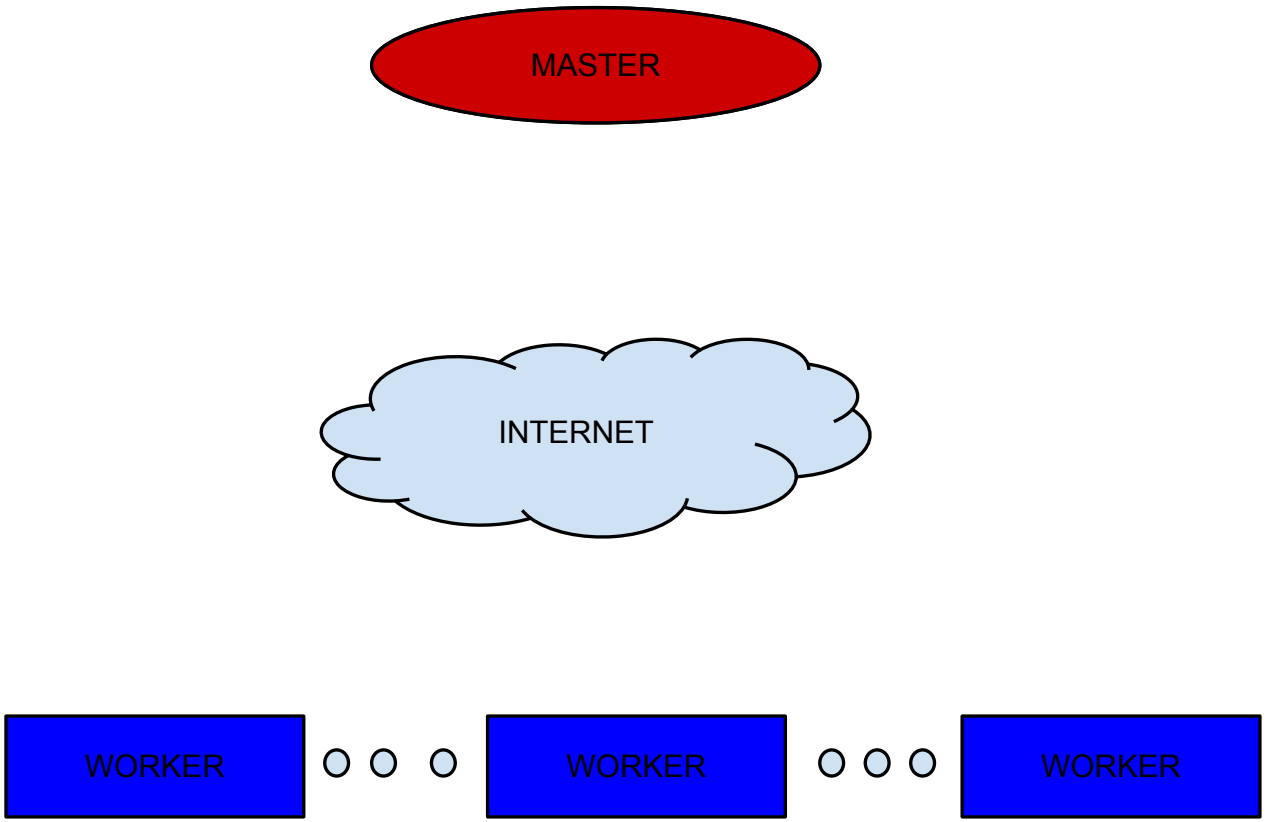

Motivation

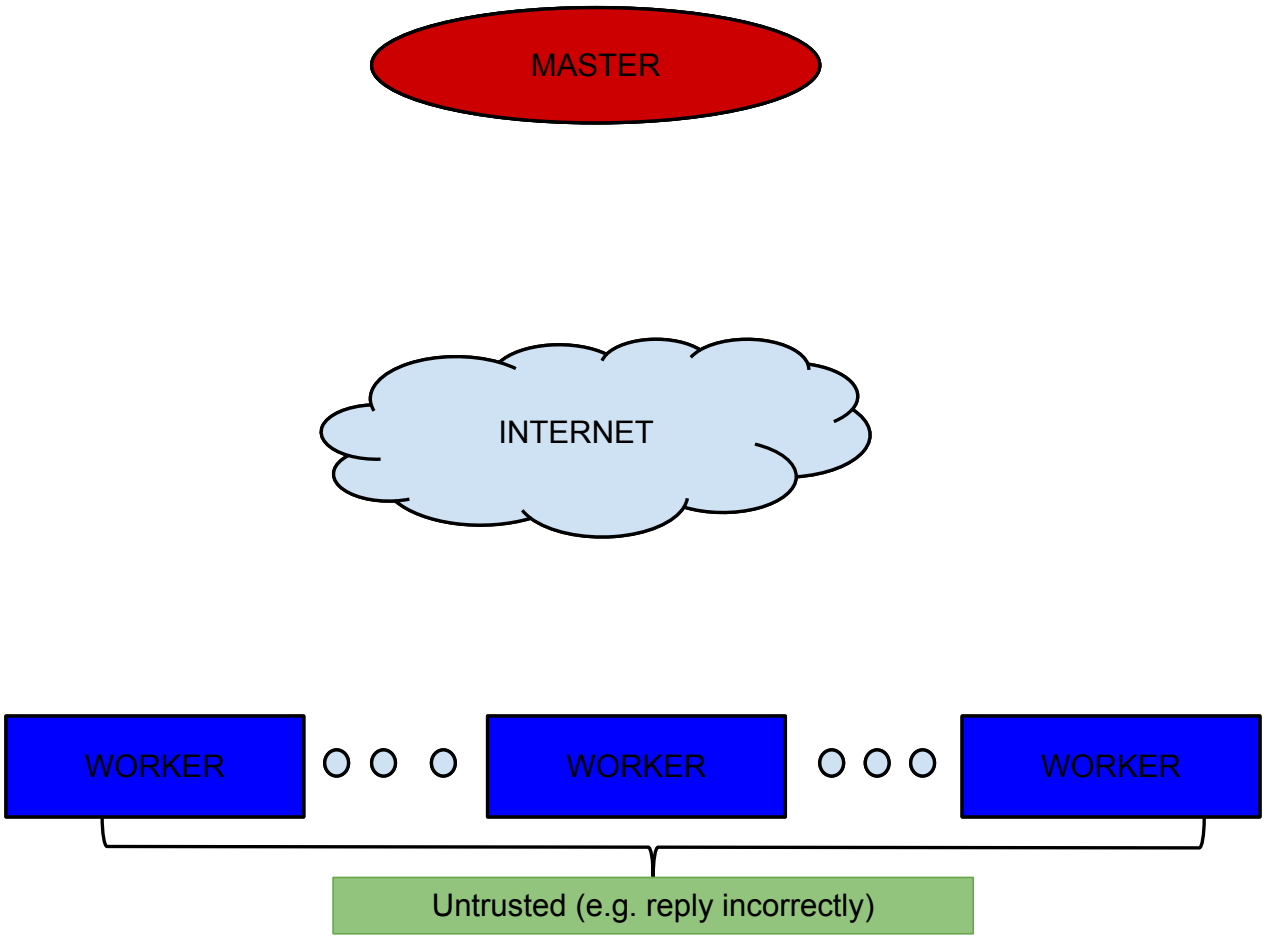
MASTER
INTERNET
WORKER
WORKER
WORKER
Untrusted (e.g. reply incorrectly)

Motivation

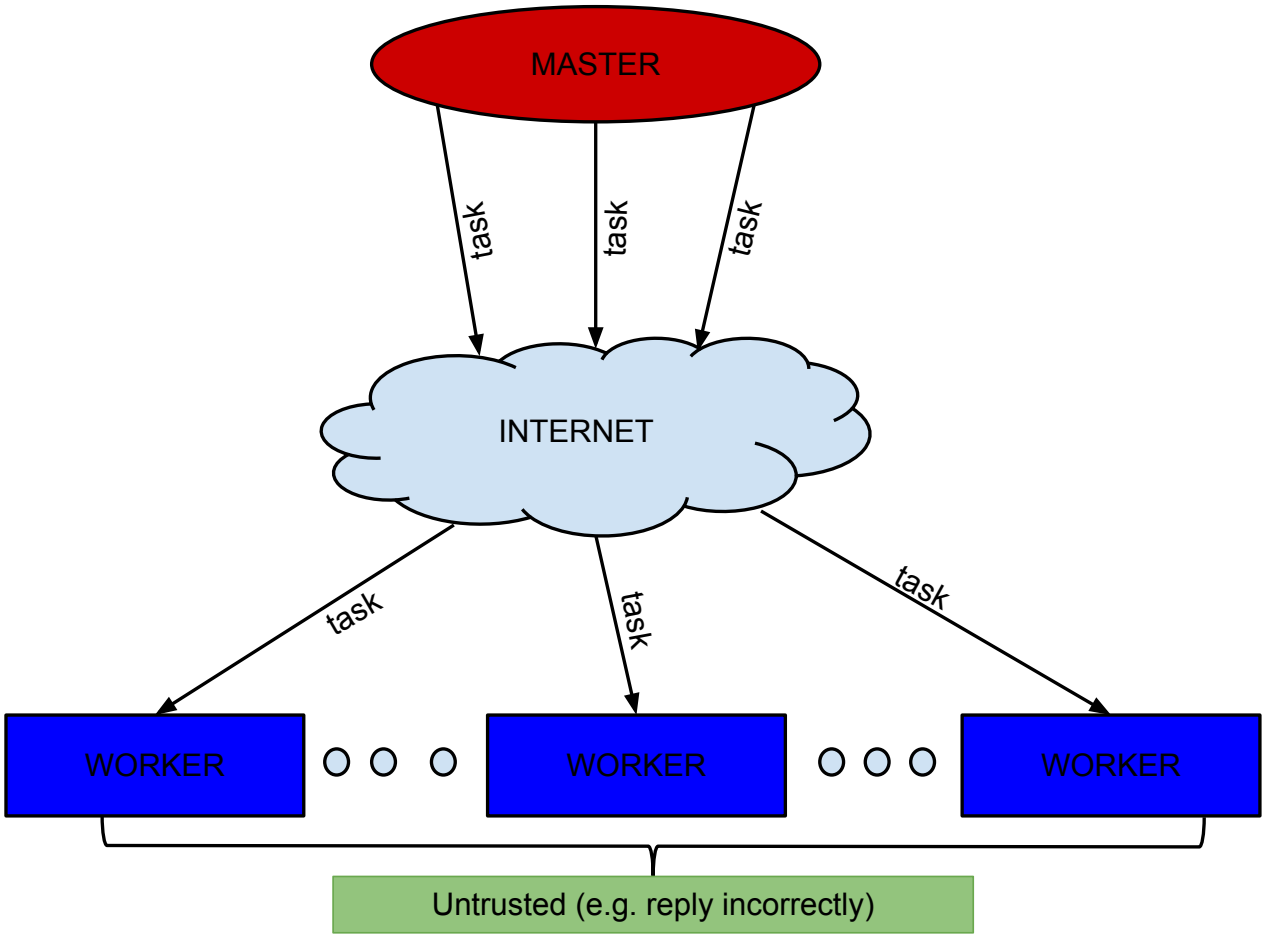
MASTER
task
task
task
INTERNET
task
task
task
WORKER
WORKER
WORKER
Untrusted (e.g. reply incorrectly)

Motivation

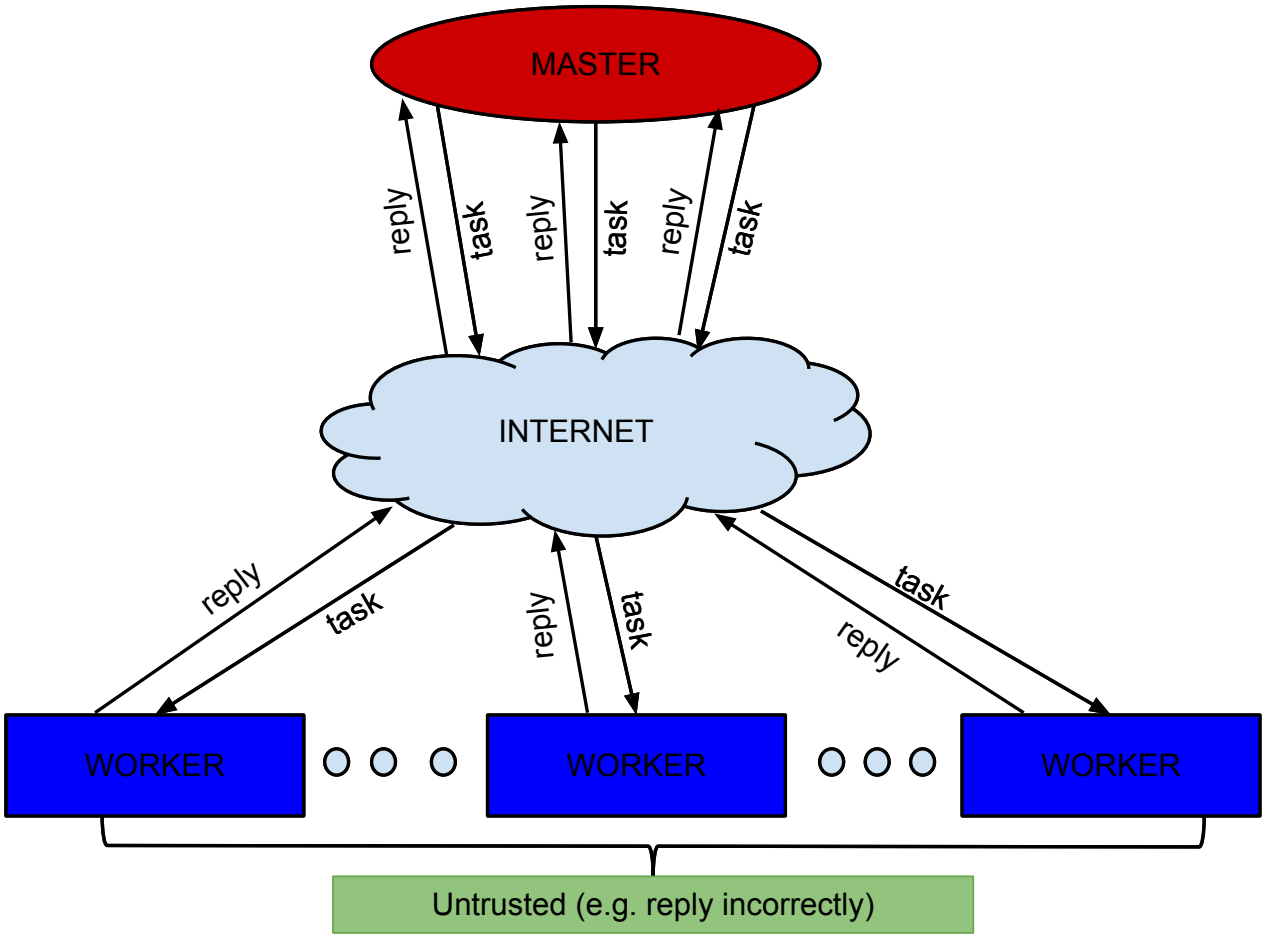
MASTER
reply
task
reply
task
reply
task
INTERNET
reply
task
reply
task
task
reply
WORKER
WORKER
WORKER
Untrusted (e.g. reply incorrectly)

Motivation

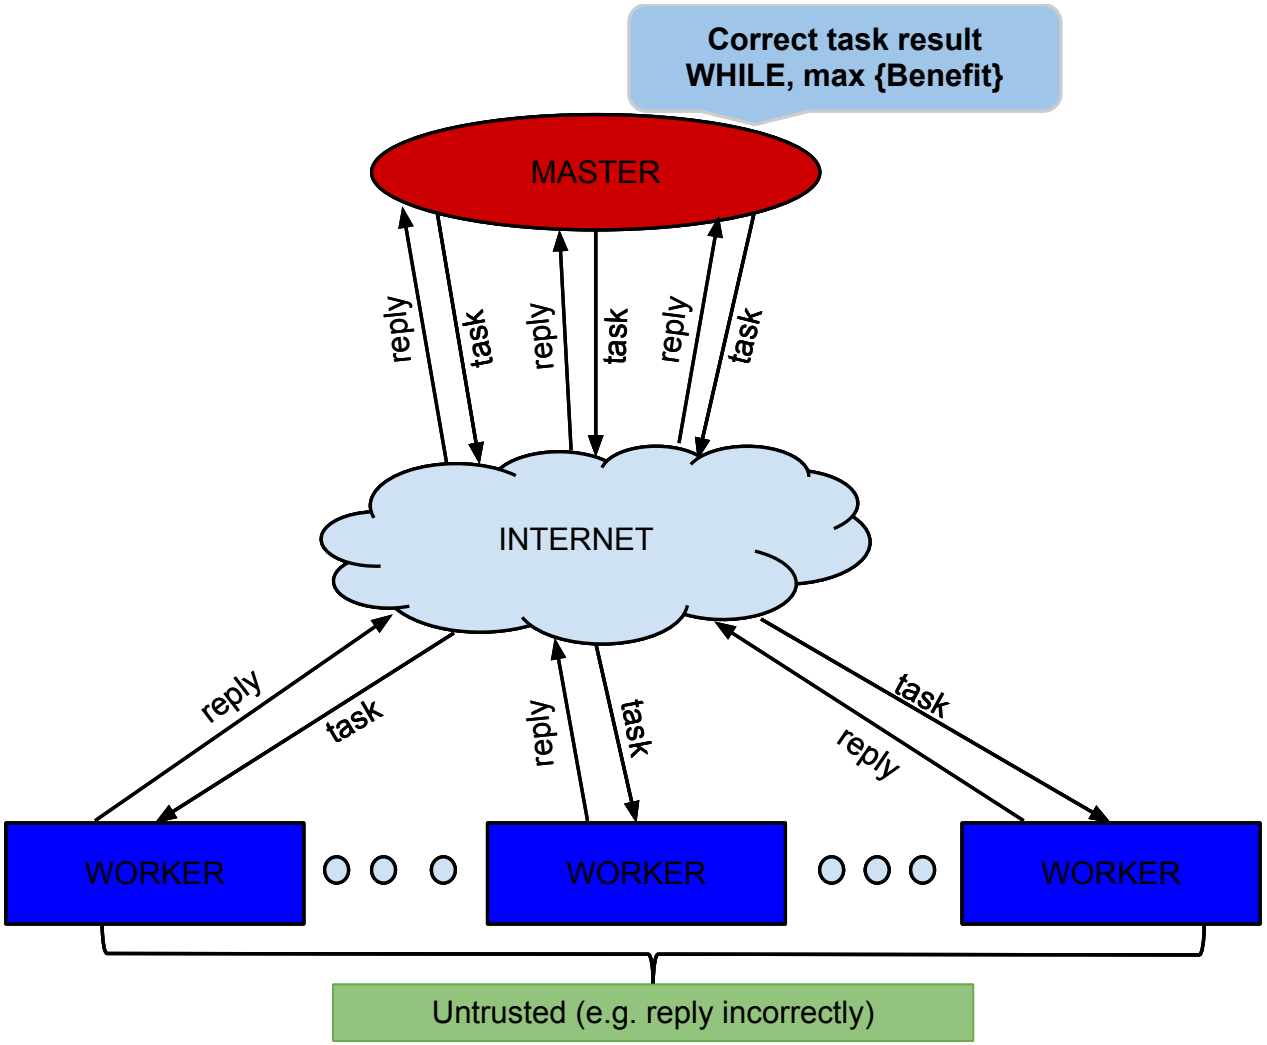
Correct task result
WHILE, max {Benefit}
MASTER
reply
task
reply
task
reply
task
INTERNET
reply
task
reply
task
task
reply
WORKER
WORKER
WORKER
Untrusted (e.g. reply incorrectly)

Motivation

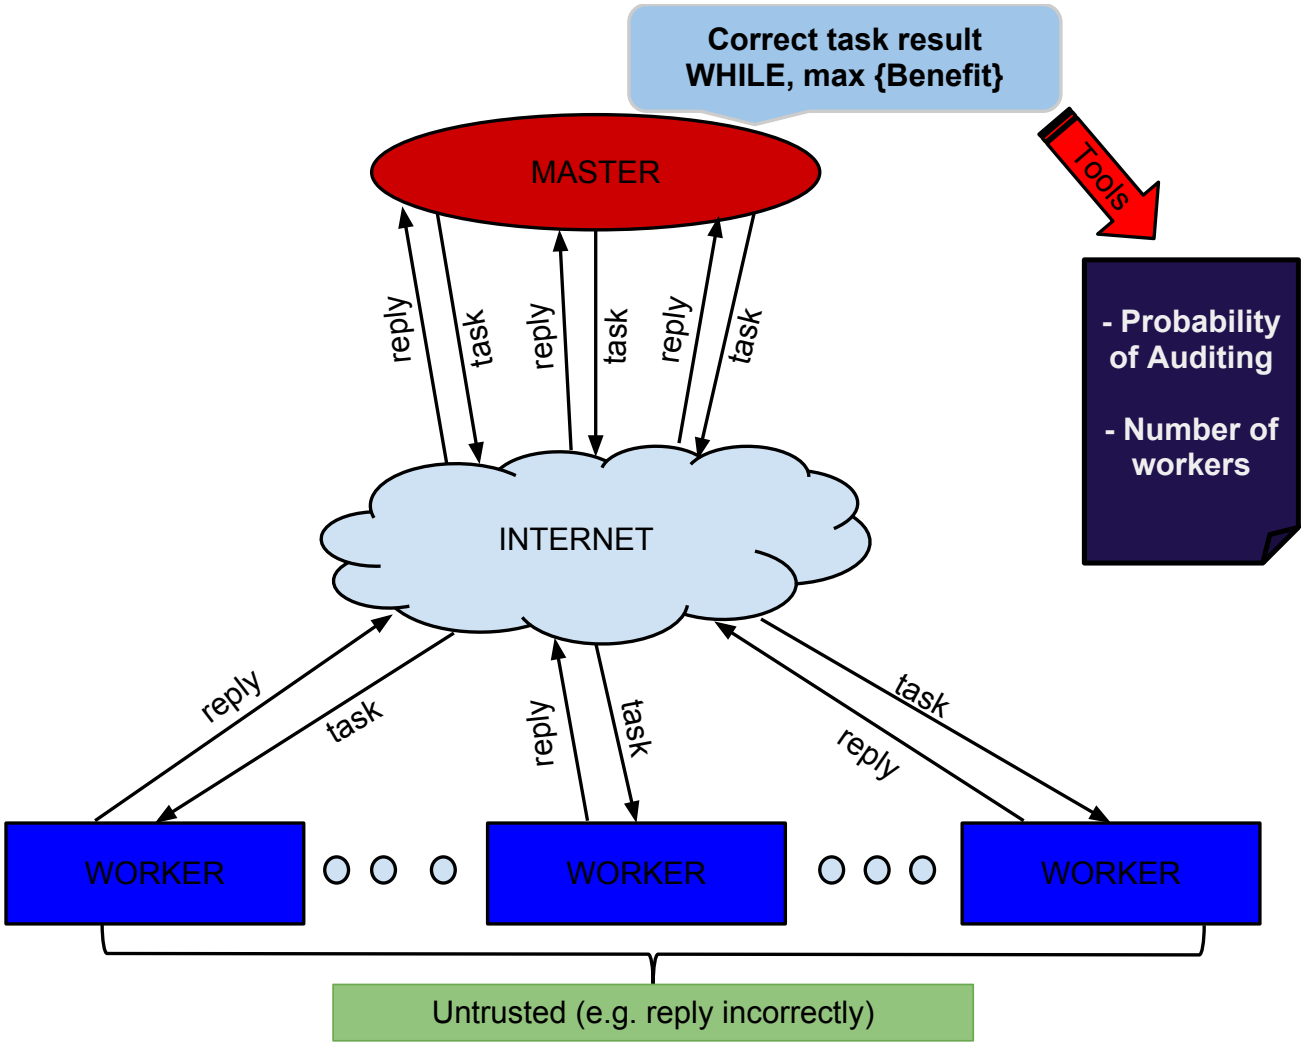
Correct task result
WHILE, max {Benefit}
MASTER
Tools
- Probability of Auditing
- Number of workers
reply
task
reply
task
reply
task
INTERNET
reply
task
reply
task
task
reply
WORKER
WORKER
WORKER
Untrusted (e.g. reply incorrectly)

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]

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

Background

Definition

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

- 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)

Prior Work

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

- Game-theoretic approach

Prior Work

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

- Game-theoretic approach
- Types of workers:
  - 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

Prior Work

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

- Game-theoretic approach
- Types of workers:
  - 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
- A reliable network was considered

Problem Statement

- Communication uncertainty
  - Messages exchanged may get lost or arrive late

Problem Statement

- Communication uncertainty
  - Messages exchanged may get lost or arrive late
- Possibility of workers not replying
  - 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 is incur by a task [Kondo et al. 07]

Problem Statement

- Communication uncertainty
  - Messages exchanged may get lost or arrive late
- Possibility of workers not replying
  - 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 is incur by a task [Kondo et al. 07]
- **Master's challenges**
  - **Provide incentives for workers to reply and reply truthfully**
  - **Ensure the above in the presence of low network reliability**

Contributions

- Develop and analyze two realistic game-theoretic mechanisms
  - Time-based mechanism
  - Reply-based mechanism

Contributions

- Develop and analyze two realistic game-theoretic mechanisms
  - Time-based mechanism
  - Reply-based mechanism
- Mechanisms provide the necessary incentives for rational workers to truthfully compute and return the task result, despite:
  - Malicious workers actions
  - Network unreliability

Contributions

- Develop and analyze two realistic game-theoretic mechanisms
  - Time-based mechanism
  - Reply-based mechanism
- Mechanisms provide the necessary incentives for rational workers to truthfully compute and return the task result, despite:
  - Malicious workers actions
  - Network unreliability
- Apply the mechanisms to two realistic settings:
  - SETI-like volunteer computing applications
  - Contractor-based applications (e.g. Amazons mechanical turk)

Framework

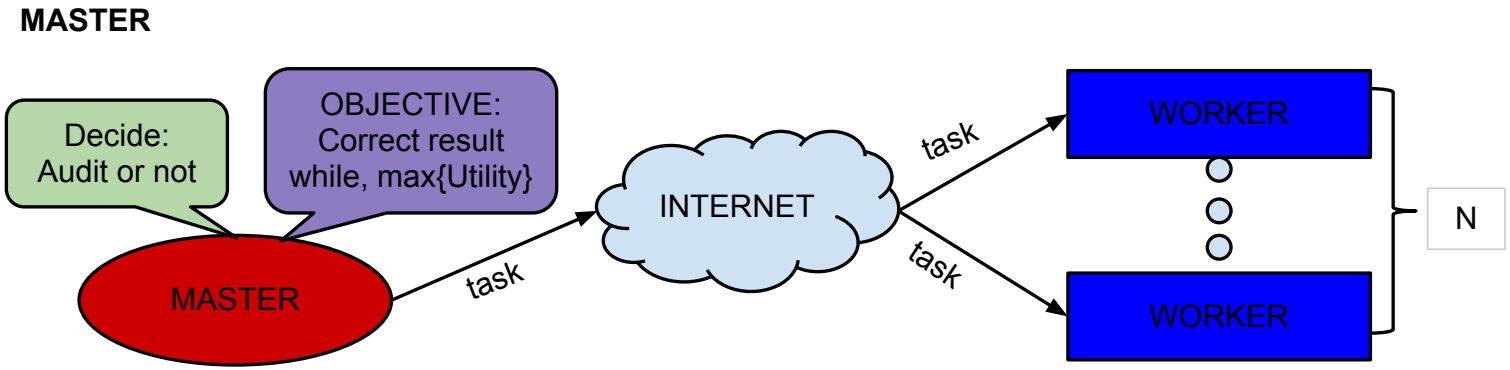
MASTER
Decide:
Audit or not
OBJECTIVE:
Correct result
while, max{Utility}
MASTER
task
INTERNET
task
task
WORKER
WORKER
N

Framework

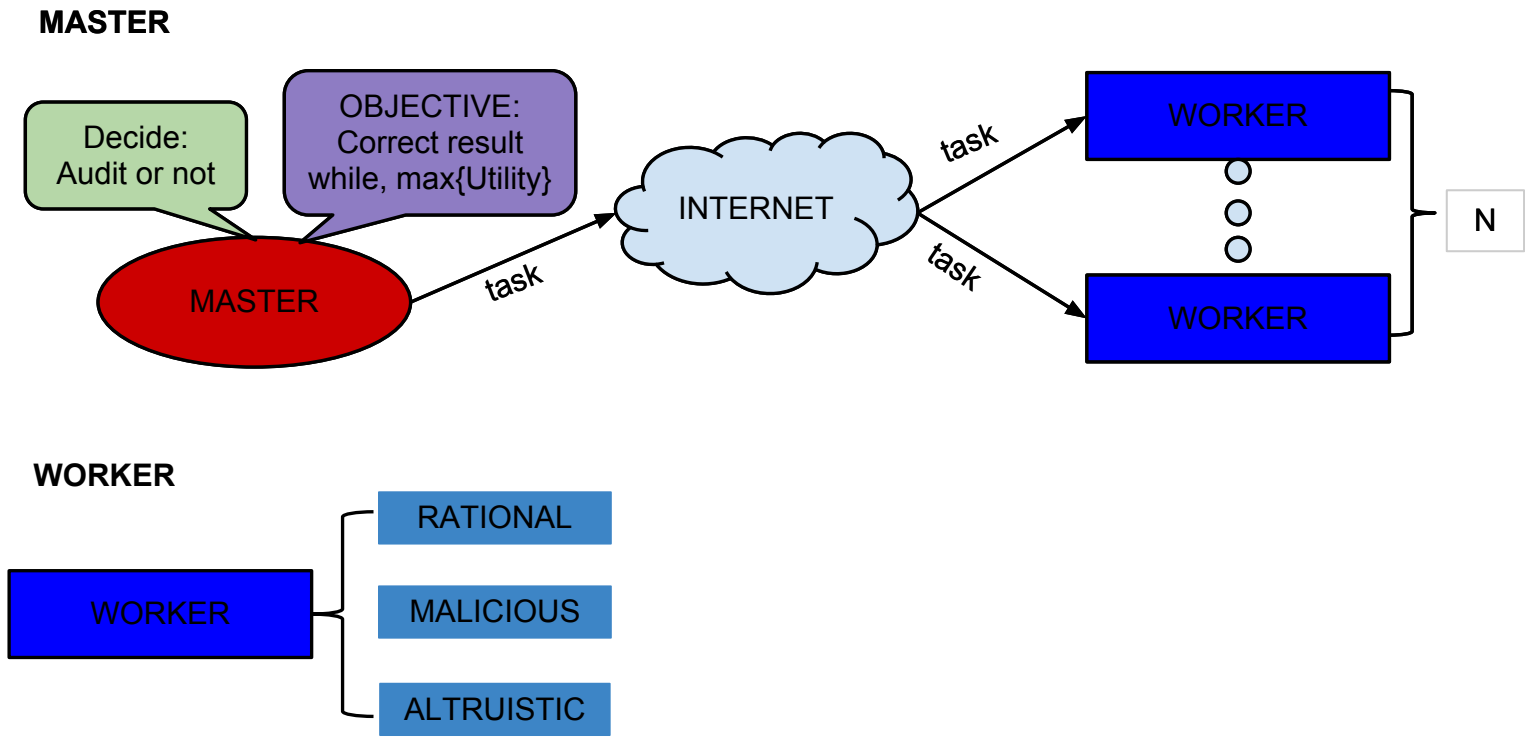
MASTER
Decide:
Audit or not
OBJECTIVE:
Correct result
while, max{Utility}
MASTER
task
INTERNET
task
task
WORKER
WORKER
N
WORKER
WORKER
RATIONAL
MALICIOUS
ALTRUISTIC

Framework

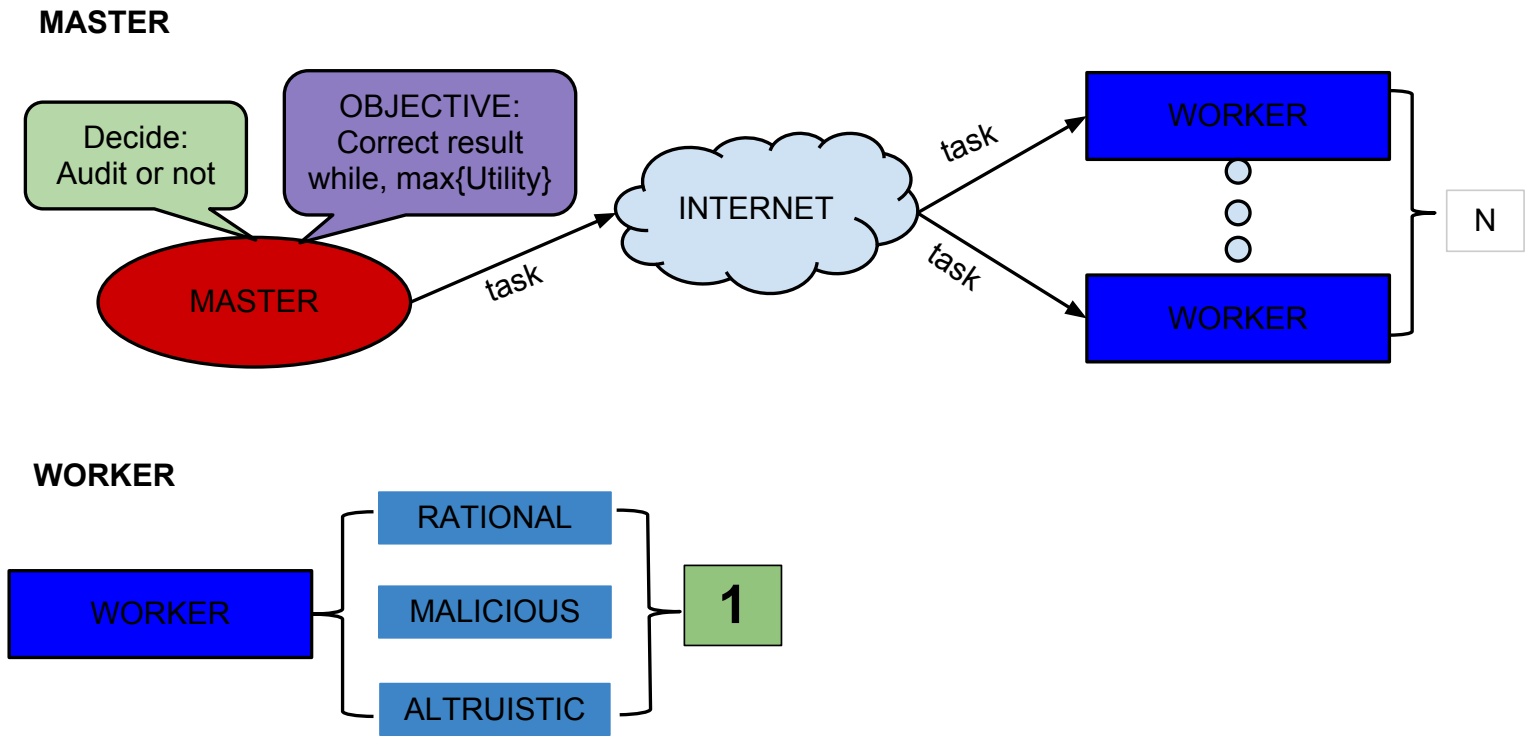
MASTER
Decide:
Audit or not
OBJECTIVE:
Correct result
while, max{Utility}
MASTER
task
INTERNET
task
task
WORKER
WORKER
N
WORKER
WORKER
RATIONAL
MALICIOUS
ALTRUISTIC
1

Framework

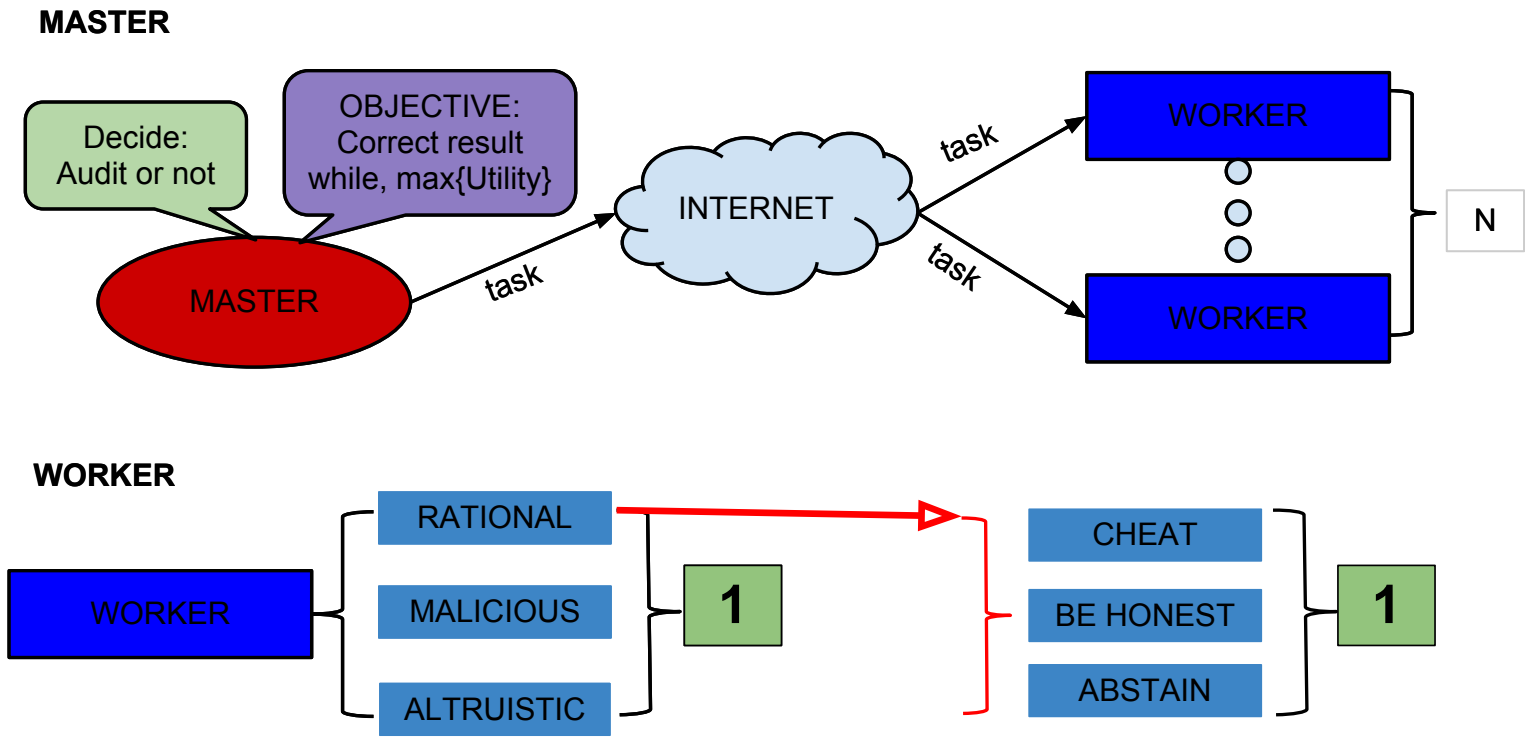
MASTER
Decide:
Audit or not
OBJECTIVE:
Correct result
while, max{Utility}
MASTER
task
INTERNET
task
task
WORKER
WORKER
N
WORKER
WORKER
RATIONAL
MALICIOUS
ALTRUISTIC
1
CHEAT
BE HONEST
ABSTAIN
1

Framework

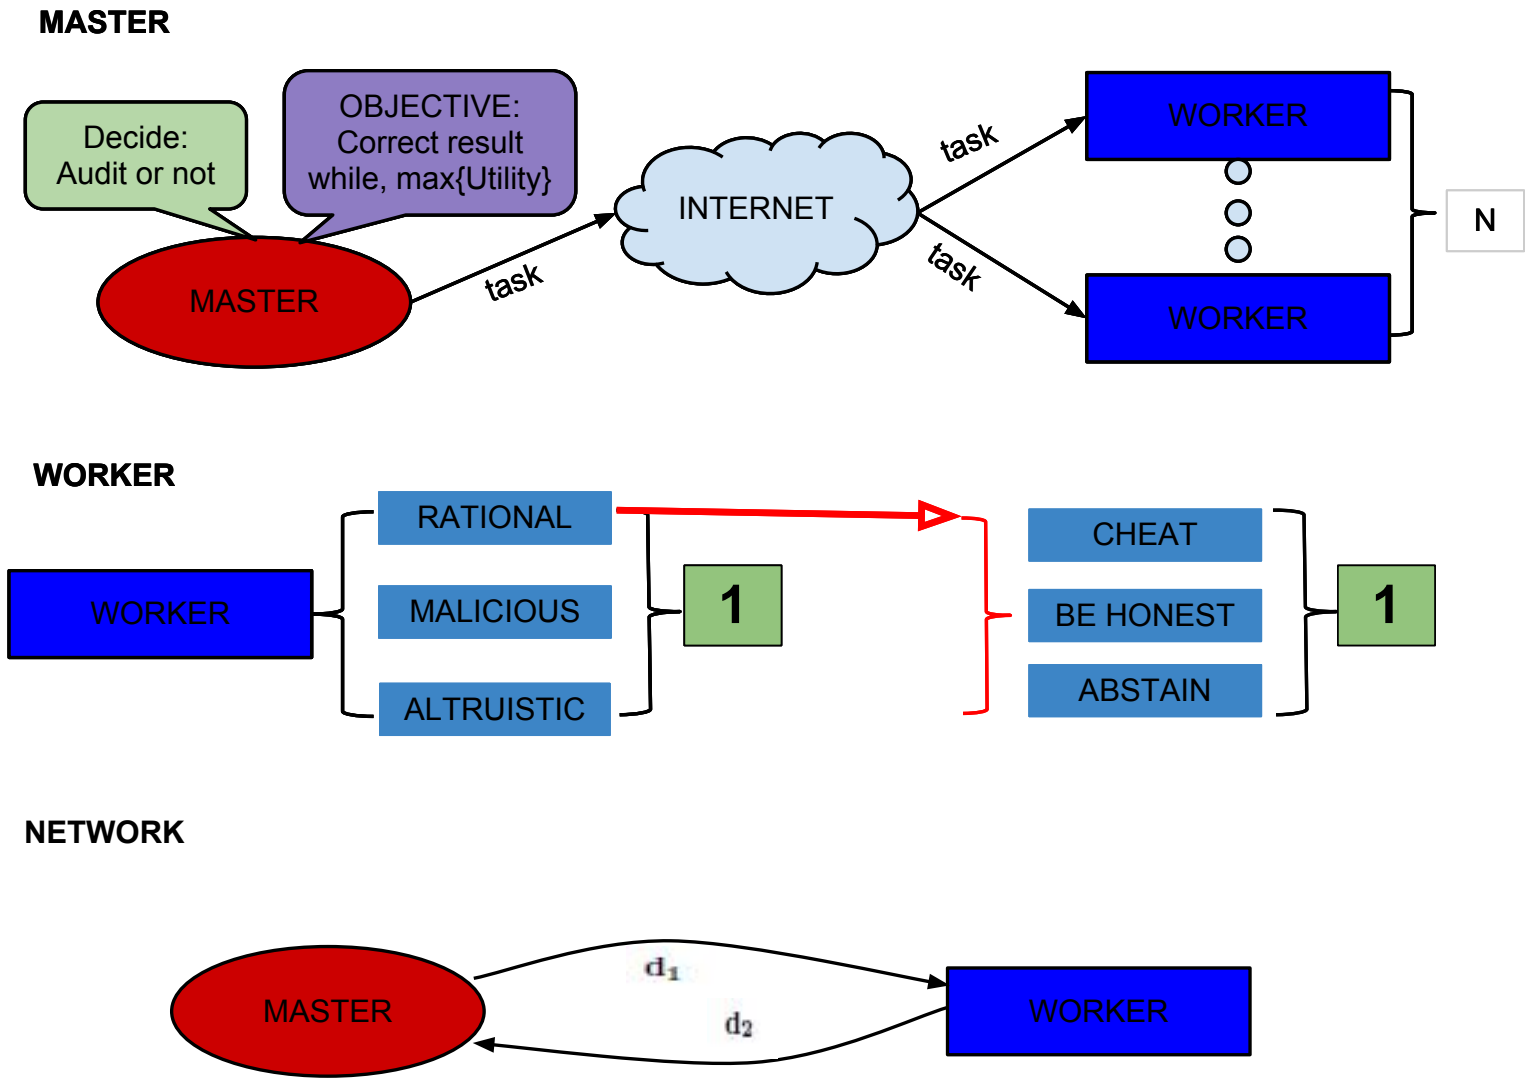
MASTER
Decide:
Audit or not
OBJECTIVE:
Correct result
while, max{Utility}
MASTER
task
INTERNET
task
task
WORKER
WORKER
N
WORKER
WORKER
RATIONAL
MALICIOUS
ALTRUISTIC
1
CHEAT
BE HONEST
ABSTAIN
1
NETWORK
MASTER
d_1
d_2
WORKER

Framework

MASTER

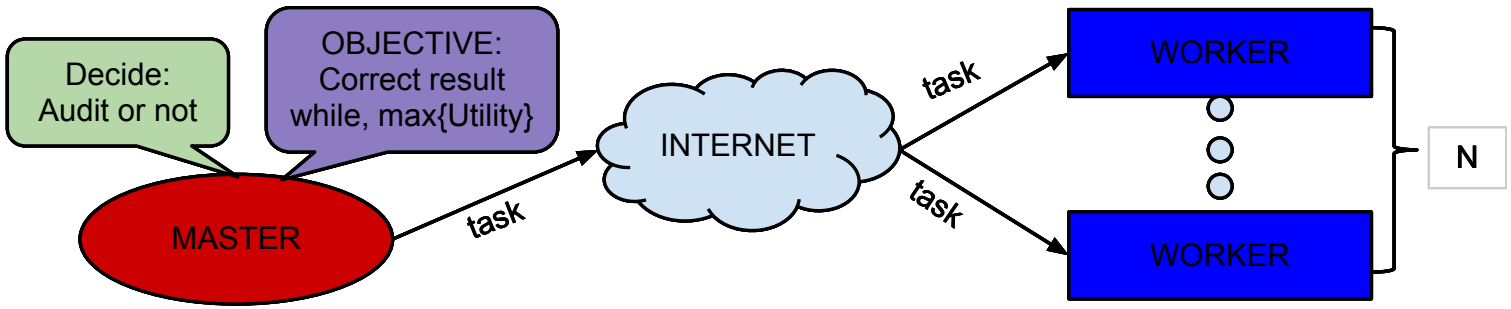

WORKER

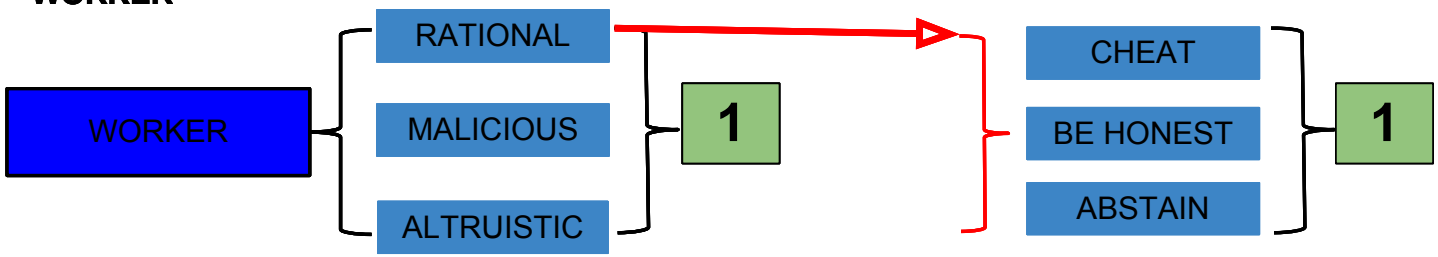

NETWORK

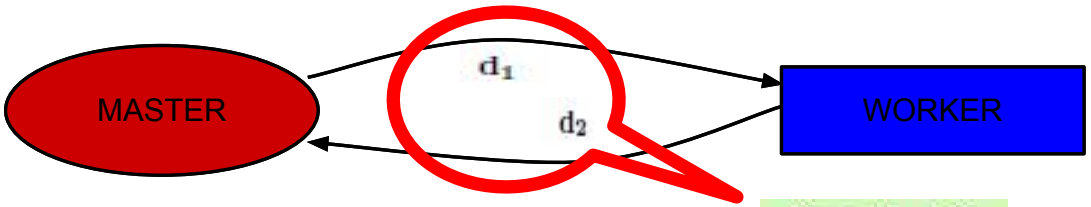

$$d = d_1 \cdot d_2$$

Framework
General Protocol

- Master assigns a task to n workers
- Rational worker cheats with probability $p_{\mathcal{C}}$ (seeking a NE)
- Master audits the responses with probability $p_{\mathcal{A}}$

Framework
General Protocol

- Master assigns a task to n workers
- Rational worker cheats with probability $p_{\mathcal{C}}$ (seeking a NE)
- Master audits the responses with probability $p_{\mathcal{A}}$
- If master audits (computes the task itself)
  - rewards honest workers and
  - penalizes the cheaters

Framework
General Protocol

- Master assigns a task to n workers
- Rational worker cheats with probability $p_{\mathcal{C}}$ (seeking a NE)
- Master audits the responses with probability $p_{\mathcal{A}}$
- If master audits (computes the task itself)
  - rewards honest workers and
  - penalizes the cheaters
- If master does not audit
  - Accepts value returned by majority of workers
  - Rewards/penalizes according to one of three models

Framework
General Protocol

- Master assigns a task to n workers
- Rational worker cheats with probability $p_{\mathcal{C}}$ (seeking a NE)
- Master audits the responses with probability $p_{\mathcal{A}}$
- If master audits (computes the task itself)
  - rewards honest workers and
  - penalizes the cheaters
- If master does not audit
  - Accepts value returned by majority of workers
  - Rewards/penalizes according to one of three models

$\mathcal{R}_{\mathrm{m}}$	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

Algorithms

- Time-based protocol
  - Master fixes a time T, once it is reached gathers all received replies
  - Ties are broken at random

Algorithms

- Time-based protocol
  - Master fixes a time T, once it is reached gathers all received replies
  - Ties are broken at random

- Reply-based protocol
  - Master fixes k, minimum estimated number of replies, by choosing n
  - If at least k replies are received, audit with $p_{\mathcal{A}}$
  - Else it does nothing, and incurs penalty $MC_{\mathcal{S}}$

Algorithms

- Time-based protocol
  - Master fixes a time T, once it is reached gathers all received replies
  - Ties are broken at random

- Reply-based protocol
  - Master fixes k, minimum estimated number of replies, by choosing n
  - If at least k replies are received, audit with $p_{\mathcal{A}}$
  - Else it does nothing, and incurs penalty $MC_{\mathcal{S}}$

- Note: Master based on statistics may have knowledge to only one of two settings

Equilibria Conditions

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

Pr(master obtains correct answer):

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

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):

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

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):

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

Expected utility of the worker when choosing to be honest over cheating and be honest over not replying

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

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

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

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

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

- 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

- 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.

SETI-like Scenario
Volunteering Computing

Time-based Mechanism

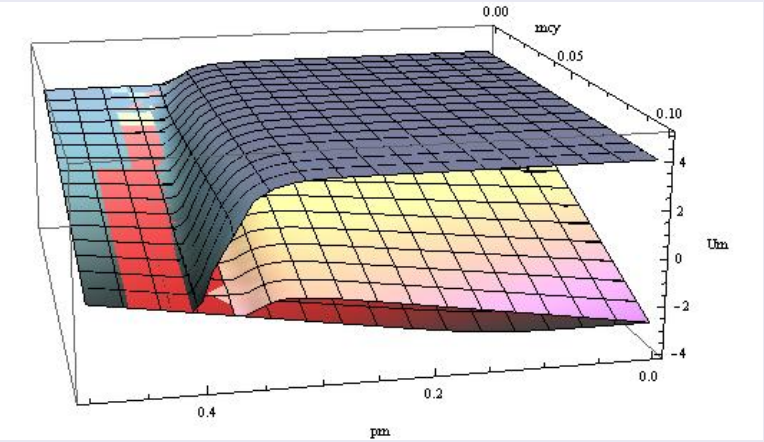

- $d = 0.9$, $n = 75$
- Upper plane $\mathcal{R}_{\emptyset}$, middle $\mathcal{R}_{\mathrm{m}}$ and lower plane $\mathcal{R}_{\mathrm{a}}$
- Master audits around $p_{\mu} = 0.35$

SETI-like Scenario
Volunteering Computing

Time-based Mechanism

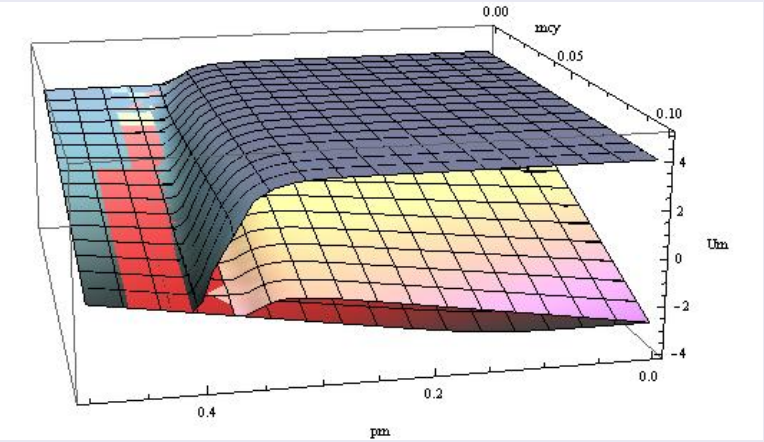

- $d = 0.9$, $n = 75$
- Upper plane $\mathcal{R}_{\emptyset}$, middle $\mathcal{R}_{\mathrm{m}}$ and lower plane $\mathcal{R}_{\mathrm{a}}$
- Master audits around $p_{\mu} = 0.35$

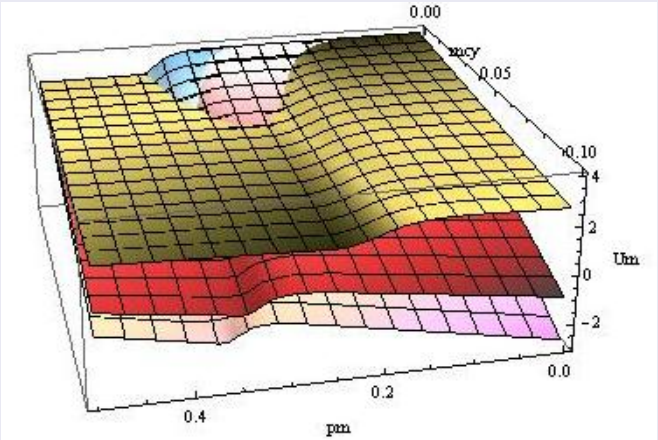

- Reward model $\mathcal{R}_{\mathrm{m}}$, $d = 0.9$
- Upper plane $n = 15$, middle $n = 55$, lower plane $n = 75$
- For $n = 15$, earlier change to auditing strategy

SETI-like Scenario
Volunteering Computing

Reply-based Mechanism

- $k \geq 1$
- Chernoff bounds for calculating k

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

with probability at least $1 - \zeta$, $0 < \zeta < 1$, where $\boldsymbol{E} = nd(p_\alpha + p_\mu)$

- $\zeta = 1/n$ (used in plot)

SETI-like Scenario
Volunteering Computing

Reply-based Mechanism

- $k \geq 1$
- Chernoff bounds for calculating k

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

with probability at least $1-\zeta$, $0 < \zeta < 1$, where $\boldsymbol{E} = nd(p_\alpha + p_\mu)$

- $\zeta = 1/n$ (used in plot)

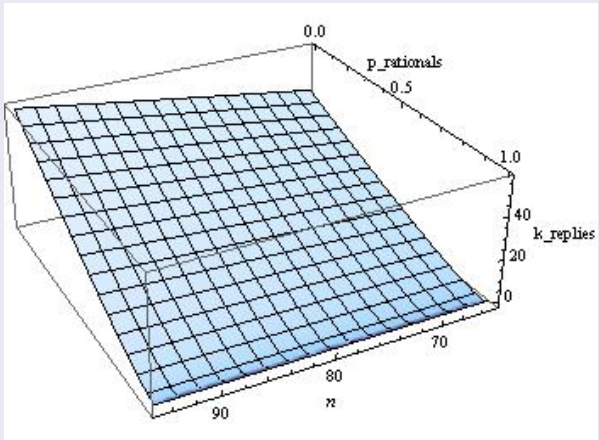

- $n \in [65, 95]$, $p_\rho \in [0, 1]$
- Appropriate value of n to get at least k replies
- p_ρ increase, k decrease

Conclusions and Future Work

- We present mechanisms for reliable computation in a volunteering computing environment
- Different types of workers
- Unreliable network
- **Future work :** We plan to explore systems with a continuous flow of tasks

Thank you

Appendix

Previous works considering master-worker computation problem:

- "Classical" distributed computing (pre-defined worker behavior)
  [Fernández et al 06; 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; Babaioff et al 06; Fernández Anta et al 08]
  - rational workers act selfishly maximizing own benefit

  Incentives are provided to induce a desired behavior

- In Fernandez, Georgiou and Mosteiro 10, the three types of workers coexist

Appendix

Framework / 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 $WB_{\mathcal{Y}} \neq MC_{\mathcal{Y}}$

Appendix

d_2 is the probability value that master achieves by

- Waiting T time, time-based mechanism
- Hiring n workers, reply-based mechanism

Why two protocols?

- Master may have knowledge to only one of two settings
  - For example based on statistics
  - 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 MC_S small

Time-based Mechanism

- Parameters values
  - $MC_{\mathcal{A}} = 1$, normalizing parameter
  - $MP_{\mathcal{W}} = 100$
  - $MC_{\mathcal{S}} = 10$
  - $MB_{\mathcal{R}} = 4$, large enough benefit given that $MC_{\mathcal{A}} = 1$
  - Different values, don't change qualitatively the results

Time-based Mechanism

- Parameters values
  - $MC_{\mathcal{A}} = 1$, normalizing parameter
  - $MP_{\mathcal{W}} = 100$
  - $MC_{\mathcal{S}} = 10$
  - $MB_{\mathcal{R}} = 4$, large enough benefit given that $MC_{\mathcal{A}} = 1$
  - Different values, don't change qualitatively the results

- 3D plots for U_M having
  - $p_\mu \in [0, 0.5]$ ($p_\mu < 0.1$ in empirical evaluations on SETI-like system, Einstein@home, Estrada et al.)
  - $MC_{\mathcal{Y}} \in [0, 0.1]$, small maintenance cost of contribution list