Reputation-aware Task Allocation for Human Trustees

[1]Han Yu, [1]Chunyan Miao, [1]Bo An, [1]Zhiqi Shen, and [1,2]Cyril Leung
[1]Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly (LILY),
Nanyang Technological University, 50 Nanyang Avenue, Singapore 639798
[2]Department of Electrical and Computer Engineering,
The University of British Columbia, 4102-2332 Main Mall, Vancouver, BC V6T1Z4, Canada
{han.yu, ascymiao, boan, zqshen, cleung}@ntu.edu.sg, cleung@ece.ubc.ca

ABSTRACT

Compared to automated entities, human trustees have two distinct characteristics: 1) they are *resource constrained* (with limited time and effort to serve requests), and 2) their utility is *not linearly related* to income. Existing research in reputation-aware task delegation did not consider these two issues together. This limits their effectiveness in human-agent collectives such as crowdsourcing systems. In this paper, we propose a distributed reputation-aware task allocation approach - RATA-NL - to address these issues simultaneously. It is designed to help an individual human trustee determine the optimal number of task requests to accept at each time step based on his situation to maximize his long term well-being. The resulting task allocation maximizes social welfare through efficient utilization of the collective capacity of the trustees, and provides provable performance guarantees. RATA-NL has been compared with five state-of-the-art approaches through extensive simulations based on human task delegation behavior abstracted from a user study involving over 100 trustees for eight weeks. The results demonstrated significant advantages of RATA-NL, especially under high workload conditions.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent agents, Multiagent systems*

General Terms

Algorithms, Human Factors

Keywords

Trust, reputation, decision support, well-being

1. INTRODUCTION

Reputation has been recognized as important mechanism for facilitating agent cooperation in an open environment. Traditionally, reputation based task delegation is studied in the context of multi-agent systems (MASs) where the agents involved are all regarded as software entities [10]. In recent years, systems that involve both software agents and human beings are starting to become widespread. For example, software agents may be deployed to assist buyers (*trusters*) in an e-commerce system to find sellers (*trustees*) who are qualified and sufficiently trustworthy to serve their requests. However, ultimately, the services are provided by the sellers who are human beings. This type of MASs can be referred to as *human-agent collectives* (HACs).

The need for agents to work together with human beings has prompted researchers to start looking for ways to infuse human factors into intelligent agents (e.g., in [5]). Over the years, human factors (e.g., cognitive processing [7]) have been explored by agent researchers. When a significant proportion of trustees in an HAC are human beings, reputation-aware task delegation (RATD) two significant new challenges. Firstly, compared to software agent, human being are more resource constrained. A software trustee can work around the clock serving trusters' requests. However, a human trustee has more limited cognitive and physical capacity, and cannot be expected to be always available as a software trustee. This leads to the second challenge. The utility of a software trustee can be measured simply by its income. However, a human trustee's utility includes not only the income aspect, but also how additional income impacts his quality of life. The mismatch between existing RATD approaches and these new challenges in HACs is reflected in reports on the widespread situation of sellers on the popular Chinese e-commerce site - Taobao.com - being chronically overworked, with some even died of exhaustion [16]. Although recent works such as [3, 13, 15] are starting to address the first challenge, they are still modeling trustees' utility as linear functions with respect to their income. In this paper, we propose a novel RATD approach to address both of these challenges simultaneously.

To address the second challenge, it is important to find a more realistic model of human trustees' utility function. In human factors research, *subjective well-being* (SWB) is commonly used as a holistic measure of people's quality of life [1]. SWB has been found to increase in a non-linear fashion with respect to increases in income [4, 9]. As a trustees' income increases, the rate of increase in this SWD starts to drop. Based on these observations, we propose a distributed Reputation-Aware Task Allocation approach for human trustees who has Non-linear utility functions (RATA-NL). By analyzing RATD from the perspective of queueing theory [8], RATA-NL helps each trustee agent determine the optimal number of new task requests to accept at each time step with a balanced consideration of his current workload, eagerness to work, expected income and task processing capacity.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

Through theoretical analysis, we proved the existence of a *lower bound* on the ratio between the social welfare (i.e., the collective well-being of the trustee agents in an MAS) and the theoretical optimal social welfare, and an *upper bound* on the waiting time for the trusters if all trustees in an MAS follow the RATA-NL approach.

To evaluate the effectiveness of RATA-NL, we have conducted a field study involving over 100 participants over an eight week period to collect data on how people make task delegation decisions in situations similar to a congestion game. To the best of our knowledge, this is the first large scale study of people's task delegation decision-making behavior under resource constraints. Based on the findings, we design a simulation test-bed to evaluate the performance of RATA-NL against five classic and state-of-the-art approaches. The results have shown that RATA-NL significantly outperforms existing approaches, especially under conditions where the workload level in an MAS is high.

2. RELATED WORK

As highlighted in [10] and [14], the trust and reputation research literature has a heavy focus on models that can produce accurate evaluations of the trustworthiness of agents. How to make interaction decisions based on these evaluations remains a relatively under explored area. The tasks received by trustees depend on how trusters make task delegation decisions based on the reputation information about the trustees. In general, trusters are assumed to be self-interested and select trustees based only on considerations about their own potential benefit.

Under such an assumption, two broad categories of self-interested RATD approaches exist: 1) the *deterministic* approach (e.g., [6, 11]), in which a truster agent always delegates tasks to the most reputable trustee agent it can find; and 2) the *probabilistic* approach (e.g., [12]), in which a truster agent delegates tasks to trustee agents with probability corresponding to their reputation. These approaches implicitly assume trustees suffer from no resource constraints and can handle addition requests without affecting the quality or timeliness of service.

In [3], the authors view RATD under the condition in which trustees are resource constrained and propose the Global Considerations (GC) approach to adjust the trustees' reputation based on their current workload. The GC approach is a probabilistic RATD approach where the probability of a trustee being selected by a truster is adjusted by the ratio between its servicing capacity and the number of newly accepted tasks whenever this ratio falls below 1.

In [13, 15], the authors formulate the RATD problem into a congestion game and propose the DRAFT and the SWORD approaches respectively to help allocate tasks to trustees so as to enhance the social welfare of a given MAS. The SWORD approach is a centralized approach designed to help trusters delegate tasks, whereas the DRAFT approach is a distributed approach to help trustees determine which requests to accept.

Nevertheless, these approaches are all designed for situations in which a trustee's utility function is linearly related to its income derived from serving truster's requests. Their models did not accommodate findings from human factors research indicating that to acquire additional income, human trustees need to expend additional effort which negatively impact their well-being.

3. PROBLEM DEFINITION

In an MAS where trustees are mostly human worker (e.g., crowdsourcing), the quality and timeliness of a trustee i in serving requests are affected by his competence and resources. His performance then determines his reputation standing in the MAS. His reputation, in turn, influences trusters' future decisions on how to delegate their requests. The confluence of these factors impacts i's workload and income, which affect his long term well-being. The need for trade-off between these considerations is well documented in human factors research.

The income of a trustee i derived from performing $\mu_i(t)$ number of tasks during time step t can be expressed as:

$$g(\mu_i(t)) = \sum_{j \in \mu_i(t)} u(j) \tag{1}$$

where $u(j) = R_j$ if task j is completed on time with quality acceptable to the truster. Otherwise, $u(j) = 0$. R_j represents the monetary reward received by i for successfully completing task j on time. Utility functions of this form is widely used by existing work in RATD such as [13] and [15]. In this paper, we enrich the RATD literature by modifying the utility function to reflect a metric most valued by human trustees according to human factors research - their *well-being*.

Through analyzing datasets published by multiple countries, various studies such as [2, 4] discovered that the marginal increase in people's SWB generally decreases as their income increases. Recent studies to quantify the relationship between SWB and income concludes that SWB can be approximated by a function of the form of $\ln(income)$ [9]. Based on these models, the SWB of i as a result of working for income during a unit time step t can be expressed as:

$$swb(g(\mu_i(t))) = \ln(1 + g(\mu_i(t))) \tag{2}$$

A "+1" term is included in the natural log function so that $swb(g(\mu_i(t)))$ evaluates to 0 in case $g(\mu_i(t)) = 0$. t normally represents a day when interpreted in the context of people working. We assume that trustees are specialized, and the tasks a trustee i is qualified to perform require similar effort from i. For example, in a crowdsourcing system, for a worker qualified to perform image labeling tasks, the effort required for him to label each image can be considered similar. Similarly, under these conditions, the value of R_j for different tasks can be regarded as the same.

The objective of this research is to design an approach to help an individual trustee i make decisions on how many of the new task requests should be accepted at any given time step t, $a_i(t)$, in order to optimize his long term well-being. Given that a large number of trustees will be involved in interactions over a potentially infinite time horizon in an uncertain environment, it is intractable (even impossible) to compute the optimal (equilibrium) strategy for all the trustees. Alternatively, this paper optimizes the well-being of each trustee indirectly by maximizing the social welfare of all trustees.[1] The *social welfare* (SW) is defined as the sum of individual trustees' SWB:

$$\bar{U} = \lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} \sum_{i=1}^{N} swb(g(\mu_i(t))) \tag{3}$$

[1]Experimental results show that a trustee has no incentive to choose a strategy other than the proposed one.

where N is the total number of trustees in the MAS. The optimization is subject to the capacity constraint of each trustee:

$$0 \leq \mu_i(t) \leq \mu_i^{max}, \forall i, t \quad (4)$$

where μ_i^{max} denotes the maximum number of tasks i can process per unit time. It is reasonable to assume that an agent of i can obtain the value of μ_i^{max} through observing i's past performance. The μ_i^{max} value can also be obtained by letting i provide an estimation.

Maximizing the sum of non-linear utility functions is a challenging problem, especially in a distributed fashion. To make the optimization problem more tractable, we view the MAS as a queueing system and analyze the optimization under the framework of reputation aware decision making. The queueing dynamics of a trustee i's pending task queue at time $t+1$, $q_i(t+1)$, can be expressed as:

$$q_i(t+1) = q_i(t) - \mu_i(t) + a_i(t) \quad (5)$$

4. THE RATA-NL APPROACH

The RATA-NL approach is designed for individual trustees to use and takes only local knowledge about the trustee as input. Intuitively, the task acceptance decision made by RATA-NL at any given point in time can be summarized as follow: "the more *eager to work* a trustee i is, the lighter his *current workload*, and the larger the *expected reward* for completing a task, the more new task requests i should accept". In this section, we present how such an intuition can be translated into an actionable task acceptance decision approach for trustees.

Let $\vec{q}(t) = (q_i(t))$ be a vector denoting the lengths of pending task queues of all trustees in an MAS. A commonly used metric for measuring the overall level of congestion in queueing system is the quadratic *Lyapunov function* [8] in the form of $L(\vec{q}(t)) = \frac{1}{2}\sum_{i=1}^{N}(q_i(t))^2$. Here, the coefficient $\frac{1}{2}$ is added to simplify the notations in subsequent analysis. We adopt this metric to measure the overall level of congestion in an MAS. Under the same framework, the stepwise change in the congestion level in the MAS can be measured by the *conditional Lyapunov drift* [8], $\triangle(\vec{q}(t))$, which is expressed as:

$$\triangle(\vec{q}(t)) = \sum_{i=1}^{N} \mathbb{E}\{L(\vec{q}(t+1)) - L(\vec{q}(t)) | \vec{q}(t)\} \quad (6)$$

where the conditional expectation is with respect to the one-step queueing dynamics given the current $\vec{q}(t)$. To ensure the decisions made by RATA-NL will not result in indefinite build-up of pending tasks in any $q_i(t) \in \vec{q}(t)$, $\triangle(\vec{q}(t))$ needs to be minimized.

By squaring both sides of Eq. (5), we have:

$$\begin{aligned} q_i^2(t+1) = q_i^2(t) + 2q_i(t)[a_i(t) - \mu_i(t)] + a_i^2(t) + \mu_i^2(t) \\ -2a_i(t)\mu_i(t) \end{aligned}$$

By re-arranging the terms and divide both sides by 2:

$$\begin{aligned} \frac{1}{2}q_i^2(t+1) - \frac{1}{2}q_i^2(t) &= \frac{1}{2}a_i^2(t) + \frac{1}{2}\mu_i^2(t) \\ &\quad + q_i(t)[a_i(t) - \mu_i(t)] - a_i(t)\mu_i(t) \\ &\leq \frac{1}{2}\mu_i^2(t) + \frac{1}{2}a_i^2(t) - q_i(t)[\mu_i(t) - a_i(t)] \end{aligned} \quad (7)$$

Due to the capacity constraints of the trustees, there exist a constant $\mu_i^{max} \geq \mu_i(t)$. To minimize the risk of i being overloaded with work, $a_i(t)$ values should be selected such that it is also less than or equal to μ_i^{max}. Thus, by taking expectation on both sides of Eq. (7), we have:

$$\triangle(\vec{q}(t)) \leq \sum_{i=1}^{N}[(\mu_i^{max})^2 - q_i(t)\mathbb{E}\{\mu_i(t) - a_i(t)|q_i(t)\}] \quad (8)$$

When determining $a_i(t)$, it is necessary to estimate its potential impact on i's well-being. The expected income from accepting $a_i(t)$ number of new task requests can be expressed in a similar way as in [13]:

$$\hat{g}(a_i(t)) = \mathbb{E}\{\sum_{j \in a_i(t)} u(j)\} = r_i(t)a_i(t)R_i \quad (9)$$

where R_i is the reward received by i, on average, for a task successfully completed on time. $r_i(t)$ is i's current reputation. $r_i(t)$ can be computed with any existing reputation model as long as the value produced by the model is within the range of $[0, 1]$. This enables $r_i(t)$ to be interpreted as the probability for i to complete a task on time with high quality at t. In many MAS-like system such as an e-commerce web site, the reputation of a seller (i.e., a trustee) is public knowledge. Thus, it is reasonable to assume that the trustee knows this value, too.

In order to ensure the stability of the pending task queues at all trustees (i.e., none of the pending task queues will grow in length indefinitely), the time averaged task acceptance rate must not exceed the time averaged task processing rate for all i [8]:

$$\lim_{T \longrightarrow \infty} \frac{1}{T}\sum_{t=0}^{T-1}\mu_i(t) \geq \lim_{T \longrightarrow \infty} \frac{1}{T}\sum_{t=0}^{T-1}a_i(t) \quad (10)$$

Each individual trustee i uses RATA-NL to determine the value of $a_i(t)$ at each time step to maximize i's SWB. At the same time, the collective decisions by trustees need minimize the general level of congestion in the MAS so that its long term operation can be sustained. From the perspective of a given MAS, these two considerations can be expressed as a *congestion-minus-well-being* expression

$$\Delta(\vec{q}(t)) - \sum_{i=1}^{N}\Psi_i\mathbb{E}\{swb(g(\mu_i(t)))|q_i(t)\} \quad (11)$$

which needs to be minimized. $\Psi_i > 0$ is a control parameter to be specified by the trustee indicating his eagerness to work. The larger the value of Ψ_i, the more eager to work i is. Since $swb(\cdot)$ is a concave and monotonically increasing function, based on Eq. (10), we have:

$$\mathbb{E}\{swb(g(\mu_i(t)))|q_i(t)\} \geq \mathbb{E}\{swb(\hat{g}(a_i(t)))|q_i(t)\} \quad (12)$$

Therefore, based on Eq. (8) and Eq. (12), Eq. (11) satisfies:

$$\begin{aligned} \Delta(\vec{q}(t)) - \sum_{i=1}^{N}\Psi_i\mathbb{E}\{swb(g(\mu_i(t)))|q_i(t)\} &\leq \Delta(\vec{q}(t)) \quad (13) \\ - \sum_{i=1}^{N}\Psi_i\mathbb{E}\{swb(\hat{g}(a_i(t)))|q_i(t)\} &\leq \delta \\ -\sum_{i=1}^{N}\Psi_i\mathbb{E}\{swb(\hat{g}(a_i(t)))|q_i(t)\} - \sum_{i=1}^{N}q_i(t)\mathbb{E}\{\mu_i(t) - a_i(t)|q_i(t)\} \end{aligned}$$

where N is the total number of trustees in the MAS. To simplify the notations, we define a constant $\delta = \sum_{i=1}^{N}(\mu_i^{max})^2$.

As $\mu_i(t)$ is not controlled by the RATA-NL approach, it can be excluded from the objective function. By isolating terms containing $a_i(t)$ on the right-hand side of Eq. (13), we have:

$$-\sum_{i=1}^{N}\mathbb{E}\{q_i(t)[\mu_i(t)-a_i(t)]+\Psi_i swb(\hat{g}(a_i(t)))|q_i(t)\} \quad (14)$$

Thus, in order to minimize Eq. (13), Eq. (14) needs to be minimized by selecting appropriate values of $a_i(t)$ at each time step. In this way, the optimization requires only the local knowledge about i, and a trustee i can individually determine the value of $a_i(t)$ to minimize his own objective function, $Obj_i(t)$:

Minimize:

$$Obj_i(t) = q_i(t)a_i(t) - \Psi_i \ln(1 + r_i(t)a_i(t)R_i) \quad (15)$$

Subject to:

$$0 \le a_i(t) \le \mu_i^{max} \quad (16)$$

$$0 \le a_i(t) \le \lambda_i(t), \forall t \quad (17)$$

where $\lambda_i(t)$ is the number of task requests for i at time step t. Its value depends on the decision-making process of the trusters which is influenced by the trustees' reputation values. The value of $a_i(t)$ can be solved by differentiating Eq. (15), which is a convex function, with respect to $a_i(t)$ and finding the critical point subject to Constraints (16) and (17):

$$\frac{\partial Obj_i(t)}{\partial a_i(t)} = q_i(t) - \frac{\Psi_i r_i(t)R_i}{1 + r_i(t)a_i(t)R_i} = 0 \quad (18)$$

$$a_i(t) = \min[\max[\lfloor \frac{\Psi_i}{q_i(t)} - \frac{1}{r_i(t)R_i} \rfloor, 0], \mu_i^{max}, \lambda_i(t)] \quad (19)$$

Eq. (19) can be interpreted as the following task acceptance policy: "the more *eager to work* a trustee i is (indicated by large Ψ_i values), the lighter his *current workload*, and the larger the *expected reward* for completing a task, the more new task requests i should accept subject to his physical limitation (Constraint (16)) and the actual number of task requests directed at him (Constraint (17))". Such a policy is rational for a human trustee and provides actionable guidance for a software agent of the trustee to compute the exact value of $a_i(t)$.

The RATA-NL approach is listed in Algorithm 1. It is designed for usage by individual trustees in a distributed fashion. In the case where not all incoming task requests are accepted by i, the RATA-NL approach informs the requesting trusters so that they can look for other alternatives. Throughout this process, no communication among trustees is required. The input for the variables required by RATA-NL can reasonably be assumed to be available with proper monitoring mechanisms in an MAS. A trustee only needs to provide a value for Ψ_i to RATA-NL following guidelines to be discussed in Section 5.

Algorithm 1 The RATA-NL Approach

Require: $q_i(t)$, $\lambda_i(t)$, $r_i(t)$, Ψ_i, R_i and μ_i^{max}.

Update $q_i(t)$ according to Eq. (5)
if $\lambda_i(t) > 0$ **then**
  if $r_i(t) == 0$ **then**
    $a_i(t) = 0$
  else
    if $q_i(t)$ is empty **then**
      $a_i(t) = \min[\mu_i^{max}, \lambda_i(t)]$
    else
      Determine $a_i(t)$ according to Eq. (19)
    end if
  end if
  Return unaccepted tasks to requesters
else
  $a_i(t) = 0$
end if

5. ANALYSIS

In this section, we analyze the impact on the social welfare of an MAS and the waiting time experienced by the trusters if RATA-NL were to be adopted by all trustees. Let $U^*(t)$ be the theoretical optimal social welfare produced by an MAS at t based on perfect foresight. Assume there are positive values Ψ, δ, and ϵ such that the *congestion-minus-utility* expression in Eq. (11) satisfies:

$$\Delta(\vec{q}(t)) - \Psi\sum_{i=1}^{N}\mathbb{E}\{swb(g(\mu_i(t)))|q_i(t)\}$$
$$\le \delta - \epsilon\sum_{i=1}^{N} q_i(t) - \Psi U^*(t)$$

where $\Psi = \frac{1}{N}\sum_{i=1}^{N}\Psi_i$. Taking expectations over the distribution of all $q_i(t)$ on both sides, we have:

$$\sum_{i=1}^{N}\mathbb{E}\{L(q_i(t+1)) - L(q_i(t))|q_i(t)\}$$
$$-\Psi\sum_{i=1}^{N}\mathbb{E}\{swb(\mu_i(t))|q_i(t)\}$$
$$\le \delta - \epsilon\sum_{i=1}^{N}\mathbb{E}\{q_i(t)\} - \Psi U^*(t)$$

which holds for all time steps. Summing both sides of the above expression over $t \in \{0, 1, ..., T-1\}$ yields:

$$\sum_{t=0}^{T-1}\sum_{i=1}^{N}\mathbb{E}\{L(q_i(t+1)) - L(q_i(t))\} - \Psi\sum_{t=0}^{T-1}\sum_{i=1}^{N}\mathbb{E}\{swb(\mu_i(t))\}$$
$$= \sum_{i=1}^{N}\mathbb{E}\{L(q_i(T)) - L(q_i(0))\} - \Psi\sum_{t=0}^{T-1}\sum_{i=1}^{N}\mathbb{E}\{swb(\mu_i(t))\}$$
$$\le T\delta - \epsilon\sum_{t=0}^{T-1}\sum_{i=1}^{N}\mathbb{E}\{q_i(t)\} - \Psi\sum_{t=0}^{T-1}U^*(t)$$

By re-arranging the terms in the above inequality, we have:

$$\epsilon\sum_{t=0}^{T-1}\sum_{i=1}^{N}\mathbb{E}\{q_i(t)\} \le T\delta + \Psi\sum_{t=0}^{T-1}\sum_{i=1}^{N}\mathbb{E}\{swb(\mu_i(t))\}$$
$$-\sum_{t=0}^{T-1}U^*(t) - \sum_{i=1}^{N}\mathbb{E}\{L(q_i(T))\} + \sum_{i=1}^{N}\mathbb{E}\{L(q_i(0))\}$$

Since $U^*(t) > 0$, $L(\cdot) \ge 0$ and $L(q_i(0)) = 0$, the above inequality can be simplified as:

$$\epsilon\sum_{t=0}^{T-1}\sum_{i=1}^{N}\mathbb{E}\{q_i(t)\} \le T\delta + \Psi\sum_{t=0}^{T-1}\sum_{i=1}^{N}\mathbb{E}\{swb(\mu_i(t))\}$$

Let $swb_{max} = \max_{[\forall n,t]} swb(\hat{g}(a_i(t)))$ be the largest observed per time step utility of any trustee in the MAS up to

$T-1$ such that $\Psi \sum_{t=0}^{T-1} \sum_{i=1}^{N} \mathbb{E}\{swb(\mu_i(t))\} \leq NT\Psi swb_{max}$. The above inequality can be written as:

$$\epsilon \sum_{t=0}^{T-1} \sum_{i=1}^{N} \mathbb{E}\{q_i(t)\} \leq T(\delta + N\Psi swb_{max})$$

By dividing both sides by $T\epsilon$, the *upper bound* on the time averaged task queue lengths for all trustees in an MAS is:

$$\limsup_{T \longrightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \sum_{i=1}^{N} \mathbb{E}\{q_i(t)\} \leq \frac{\delta + N\Psi swb_{max}}{\epsilon} \qquad (20)$$

Similarly, the *lower bound* on the time averaged social welfare produced by an MAS is:

$$\begin{aligned}
&\liminf_{T \longrightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \sum_{i=1}^{N} \mathbb{E}\{swb(g(\mu_i(t)))\} \\
&\geq \frac{1}{T} \sum_{t=0}^{T-1} U^*(t) - \frac{\delta}{\Psi} - \frac{\sum_{i=1}^{N} \mathbb{E}\{L(q_i(0))\}}{T\Psi} \\
&\quad + \frac{\epsilon \sum_{t=0}^{T-1} \sum_{i=1}^{N} \mathbb{E}\{q_i(t)\} + \sum_{i=1}^{N} \mathbb{E}\{L(q_i(T))\}}{T\Psi} \qquad (21) \\
&\geq \frac{1}{T} \sum_{t=0}^{T-1} U^*(t) - \frac{\delta}{\Psi}
\end{aligned}$$

Therefore, under the condition where the task allocation recommendations by RATA-NL to all trustees are fully complied with, the MAS can produce time averaged social welfare within $O(1/\Psi)$ of the optimal social welfare with average waiting time experienced by the trusters bounded by $O(\Psi)$.

6. EVALUATION

To evaluate the performance of RATA-NL, we design the experiments in two stages. Firstly, we conduct a user study to collect data reflecting people's task delegation decision making behavior both with and without considering trustees' reputation in a congestion game-like situation. Then, we build a simulation test-bed with agent behavior modeled based on the collected data and compare the performance of RATA-NL with five other state-of-the-art approaches.

6.1 Real World Task Allocation Decisions

The user study was conducted in conjunction with an eight week long group-based software engineering project in Beihang University, China from April to June 2013. In this round of data collection:

- 122 students are divided into 20 teams with an average team size of about 6 persons. Each team is required to produce a software system at the end of the project according to the course requirements.
- Each team proposes the tasks they need to complete in order produce the required software. Tasks are divided among the team members at the beginning of each week based on discussions within each team. A total of 733 tasks were proposed by all teams.
- Each student may be assigned 0, 1, or multiple tasks during each team meeting. Each student reports his estimations of the expected number of days needed to complete the task(s). All decisions are recorded.

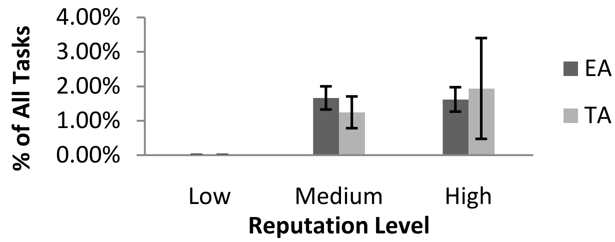

Figure 1: Real world task delegation decisions

- In the following team meeting, the quality of the completed tasks and the timeliness of their completion relative to the expected time are evaluated by the members of each team together. The performance of each team member is rated based on the quality of his work in a binary scale.

During the experiment, the participants adopted two different task allocation approaches: 1) *equality-based* approach (EA), in which team members are assigned tasks as evenly as possible without regard to their competence; and 2) *trust-based* approach (TA), in which team members are assigned tasks based on their past performance. Roughly half of the participants adopted each of the two approaches during the coursework. Figure 1 shows the distribution of workload measured by the average and the standard deviation of the percentage of all tasks delegated to individuals with respect to their reputation levels (*low* with reputation values within [0, 0.3]; *medium* with reputation values within (0.3, 0.7]; and *high* with reputation values within (0.7, 1]). The participants' reputation is computed after the experiment based on their performance using the *BRSEXT* method in [14]. Participants were not required to formally evaluate each other's reputation using computational methods during the study. Instead, those in the TA group relied only on their intuition to determine how trustworthy others are.

Let N_{EA} and N_{TA} denote the total number of tasks proposed in the EA and the TA groups respectively. In the experiment, none of the participants is of a low level of reputation. Under EA, participants with medium and high level of reputation are assigned an average of 1.65% of N_{EA} with standard deviations around 0.35% of N_{EA}. Under TA, participants with medium level of reputation are assigned an average of 1.24% of N_{TA} with standard deviations around 0.46% of N_{TA}. The numbers for those with high level of reputation are 1.93% and 1.46% respectively.

6.2 Experiment Design

In the simulation test-bed, we assume binary outcomes for task results (i.e., a task is considered successfully completed by a trustee agent if the quality of the result is satisfactory and the result is produced before its expected deadline; Otherwise, it is considered unsuccessful). We create six groups of 100 trustee agents each to compare the performance of six different RATD approaches. They are:

1. The Equality-based Approach (EA): this is an approach based on the patterns exhibited by participants in the EA Group in our dataset.
2. The Trust-based Approach (TA): this is also an approach based on the patterns exhibited by participants

in the TA Group in our dataset. Each truster agent uses its direct interaction experience with trustees in the past to evaluate their trustworthiness using the *BRSEXT* method in [14]. At each time step, a truster agent delegates tasks to the most trustworthy known trustee.

3. The Reputation-based Approach (RA): each truster agent then adopts the RATD approach in [12] in which the probability of a trustee agent being selected by a truster agent is directly related to its reputation standing among all trustee agents in the MAS.

4. The Global Considerations (GC) Approach: truster agents adapt the probability for each trustee agent being selected to serve task requests following the approach in [3].

5. The DRAFT Approach: trustee agents make request acceptance decisions following the approach in [13].

6. The RATA-NL Approach: trustee agents make request acceptance decisions following the proposed approach.

Truster agents in Approaches 4, 5, and 6 also adopt the RATD approach in [12].

Trustee agents behave following one of the listed patterns:

1. *Com*: competent trustees who produce satisfactory quality results with a 90% probability;

2. *MC*: moderately competent trustees who produce satisfactory quality results with a 70% probability;

3. *MI*: moderately incompetent trustees who produce satisfactory quality results with a 30% probability;

4. *Inc*: incompetent trustees who produce satisfactory quality results with a 10% probability.

The task processing capacities (μ_i^{max} values) of each type of trustee agents are set in such a way that more competent agents can process more tasks per time step than less competent ones. In this paper, when we refer to a trustee agent population as "x% competent", the exact composition of the population consists of $\frac{1}{2}x$% *Com* trustees, $\frac{1}{2}x$% *MC* trustees, $\frac{1}{2}(100-x)$% *MI* trustees, and $\frac{1}{2}(100-x)$% *Inc* trustees. In the experiments, the trustee agent population compositions are varied from 10% to 100% competent to simulate different trustee behavior patterns.

Another factor affecting the well-being of trustees is the generally level of workload in an MAS. As the workload is relative to the aggregate task processing capacity of a given trustee population, we define a formula to compute the *maximum throughput*, θ, of a trustee population per time step as $\theta = \sum_{i=1}^{N} c_i \mu_i^{max}$, where c_i denotes the competence value of a trustee i. The workload on a given trustee population is measured by a metric called *Load Factor* (LF) which is computed as $LF = \frac{N_{req}}{\theta}$, where N_{req} is the average number of task requests generated by trusters per time step. In the experiments, we vary the LF value from 25% to 200% to simulate different workload conditions. Under each configuration, the simulation is run for $T = 1000$ time steps. In all experiments, trustees process tasks at an average rate of $0.9\mu_i^{max}$ with a standard deviation of $0.1\mu_i^{max}$. On average, trusters expect a task to be completed within 2 time steps after it is allocated to a trustee.

6.3 Evaluation Metrics

In the experiments, we measure the performance of each approach using the following metrics:

1. *Social welfare to optimal social welfare ratio* (SW / Opt SW): The optimal social welfare (Opt SW), U^*, that can be produced by a trustee population per time step is expressed as $U^* = \sum_{i=1}^{\lfloor \min[1.0, LF]N \rfloor} \ln(1 + c_i \mu_i^{max} R_i)$. The N trustee agents are sorted in descending order of their c_i values. U^* can only be calculated in a controlled experimental environment where LF, μ_i^{max}, c_i and R_i can be definitively known. The SW/Opt SW ratio is calculated as $\bar{U}/U^*$.

2. *High quality rate* (HQR): this metric is computed as $\frac{N_q}{N_{acc}}$ where N_q denotes the total number of tasks completed with satisfactory quality, and N_{acc} denotes the total number of tasks accepted by the trustees over T time steps in a simulation.

3. *Timely completion rate* (TCR): this metric is computed as $\frac{N_t}{N_{acc}}$ where N_t denotes the total number of tasks completed before the expected deadlines over T time steps in a simulation.

The higher the values for these metrics, the better the performance of an approach.

In addition, we also measure how different RATD approaches may affect the truster agents' perceptions on the behavior of the trustees. It is important as it will affect the truster agents' subsequent task delegation decisions and, in turn, the trustee agents' well-being. Ideally, the reputation value of a trustee should only reflect its competence rather than performance variations caused by changing workloads which is not the trustee's own fault. We adopt commonly used metrics including *precision*, *recall*, *f-value*, and *mean absolute error* (MAE) to measure how accurately each approach classifies whether trustees are trustworthy against the ground truth. Precision=$\frac{t_p}{t_p+f_p}$, Recall=$\frac{t_p}{t_p+f_n}$, and f-value=$\frac{2\times Precision \times Recall}{Precision+Recall}$, where t_p (*true positive*) is the number of trustees correctly classified as competent (i.e., $t_p = \frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N} 1_{[r_i(t)>0.5|c_i>0.5]}$), f_p (*false positive*) is the number of trustees incorrectly classified as competent (i.e., $f_p = \frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N} 1_{[r_i(t)>0.5|c_i\leq 0.5]}$), and f_n (*false negative*) is the number of trustees incorrectly classified as incompetent (i.e., $f_n = \frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N} 1_{[r_i(t)\leq 0.5|c_i>0.5]}$). $1_{[condition]}$ equals to 1 if $[condition]$ is true, and 0 otherwise. MAE is defined as $\frac{1}{TN}\sum_{t=0}^{T-1}\sum_{i=1}^{N}|r_i(t) - c_i|$.

6.4 Results and Discussions

Figure 2 contains sub-figures showing the performance of the various approaches according to the evaluation metrics. Each data point in these figures represents the average value of the selected metric taken over 10 different trustee population configurations (10% to 100% competent) under a given load factor. As the results are based on simulations, the trends and relative performances of the approaches are more important than the exact numerical values.

As shown in Figure 2(a), there is almost no difference between the precision values achieved by various approaches. As the trusters in our simulations truthfully share their opinions about the trustees, the f_p values become very low given enough observations, resulting in generally high precision values. However, the differences in performance among

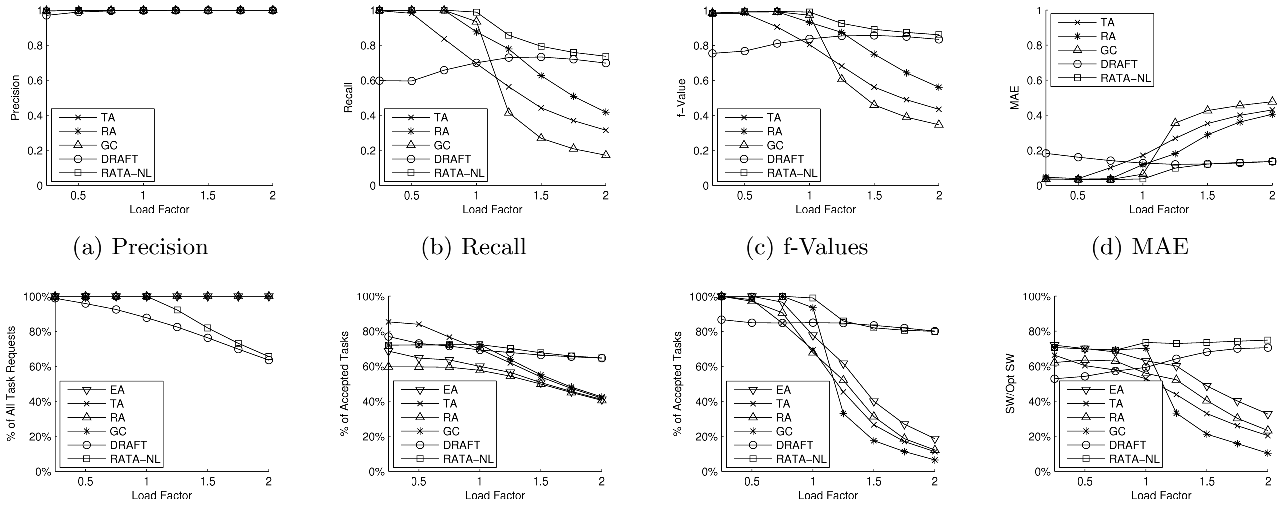

(a) Precision (b) Recall (c) f-Values (d) MAE

(e) N_{acc}/N_{req} (f) HQR (g) TCR (h) SW/Opt SW

Figure 2: Comparison of performance

various approaches measured by their recall values are more significant (Figure 2(b)). With $LF < 1.0$, under EA, TA, RA, GC, and RATA-NL, most unsuccessful tasks are caused by the competence of the trustees. Only a negligible percent of unsuccessful tasks are caused by failure to be completed on time due to trustees being overloaded. With $LF \geq 1.0$, the f_n values for EA, TA, RA, GC, and RATA-NL start to increase as more trustees suffer from reputation damage due to overloading, causing the recall values of these approaches to decrease. The stringent task allocation criteria used by DRAFT resulted in over concentration of tasks on reputable trustees under low workload conditions. As workload becomes higher, the performance of DRAFT improves to close to that of RATA-NL. These factors result in the relative performances of the approaches measured with their f-values (Figure 2(c)) and MAE values (Figure 2(d)). In terms of accurately reflecting the behavior of the trustees rather than the dynamics caused by inefficient task allocation decisions, RATA-NL significantly outperforms EA, TA, RA, and GC under high workload conditions, and significantly outperforms DRAFT under low workload conditions.

Figure 2(e) shows the ratio between the accepted tasks and all proposed tasks under different RATD approaches. Since EA, TA, RA, and GC do not provide mechanisms for trustees to reject incoming task requests, regardless of how the LF value changes, all proposed tasks are delegated to some trustees for processing. Under DRAFT, N_{acc} starts to drop even when workload is low ($LF < 1.0$). In the case of RATA-NL, all proposed tasks are allocated to trustees when $LF \leq 1.0$ and N_{acc} only starts to drop when $LF > 1.0$. When $LF > 1.0$, the task request arrival rates become larger than the task processing rates the trustee populations can effectively support. In this case, if the extra tasks are not dropped, they can cause delays and negatively affect the performance of the trustees as perceived by the trusters. Under such conditions, to protect the long term well-being of the trustees, it is advantageous to drop some requests.

Under $LF < 1.0$, the HQR achieved by TA beats other approaches (Figure 2(f)). However, as LF increases, the HQRs achieved by EA, TA, RA, and GC dropped significantly (for TA, as much as 45 percentage points). Whereas for DRAFT and RATA-NL, the HQRs remains relatively stable under changing workload conditions. A similar performance pattern can be observed from their TCRs (Figure 2(g)). Under this metric, the performance of RATA-NL matches that of DRAFT with $LF > 1.0$.

Figure 2(h) shows the SW/Opt SW ratios achieved by various approaches. With $LF \leq 1.0$, the differences in the performance of various approaches are not large (within 15%). The performance of RATA-NL, GC, and EA are almost the same, with RA and TA trailing not far behind. DRAFT delivered the worst performance under low LF conditions where the capacity of the trustees cannot be fully utilized. With $LF > 1.0$, the advantage of RATA-NL and DRAFT over other approaches become more significant with RATA-NL outperforming DRAFT by about 10%.

To study whether a self-interested trustee agent has incentive to follow an approach other than RATA-NL, we design another experiment. The truster agents in this experiment delegate tasks following the RA approach. A trustee agent population of 100 is generated to compete for tasks from the truster agents. The trustee agents are all of *Com* type. They are equally divided into four groups using GC, DRAFT, RATA-NL, and TRD (i.e., the traditional *accept-when-requested* approach) respectively. The load factor is varied from 25% to 200%, and the total utility of each individual trustee agent is recorded after each round of simulation which consists of 1000 time steps.

Figure 3 shows the highest utility of trustee agents using various approaches as a percentage of the lowest utility of trustee agents using RATA-NL under different load factors. It can be observed that the most well to do trustee agents under TRD, GC and DRAFT all achieved less than 100% of the utility of the least well to do trustee agent under RATA-NL under all workload conditions. On average, the least well to do trustee agents using RATA-NL achieved 48.17%, 65.03% and 21.49% more utility than the most well to do trustee agents using TRD, GC and DRAFT respectively. The results show that in a competitive environment, a rational trustee agent using RATA-NL has no incentive switching to another approach.

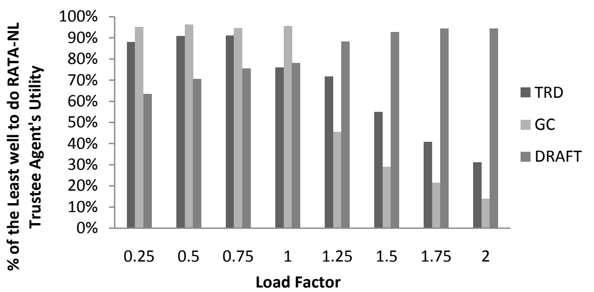

Figure 3: The *highest* utility of *Com* trustee agents using various approaches as a percentage of the *lowest* utility of those using RATA-NL

7. CONCLUSIONS

In this paper, we take a step towards bringing the discussion about reputation-aware decision-making into the domain of human-agent collectives. We propose the RATA-NL approach for human trustees, who are resource constrained and whose well-being has been found to be non-linearly related to their income. RATA-NL helps them make pragmatic and holistic task acceptance decisions at each time step so as to maximize the well-being of the trustee community, which implies that trustees can enjoy improved work-life balance. In addition, it also protects their reputation from being damaged by uncoordinated task allocation decisions by trusters. Theoretical analysis proves the existence of a lower bound on the social welfare of the trustees and an upper bound on the waiting time experienced by the trusters under the RATA-NL approach. Simulations infused with human task delegation decision patterns extracted from our field study involving over 120 participants demonstrate significant advantage of RATA-NL over five existing approaches, especially under high workload conditions.

As RATA-NL is designed for individual human trustees to use in a distributed fashion, it does not require trusters to modify their existing decision-making process. This makes it attractive for use in deployed systems such as e-commerce platforms to enhance user experience and overall system performance. In order to enable agents to operate in an MAS together with human beings, more data are needed to construct models of human performance and decision-making behaviors. In future research, we plan to set up a game based online platform for long term collection of such data to build a benchmark for this research topic.

8. ACKNOWLEDGMENTS

This research is supported in part by Interactive and Digital Media Programme Office, National Research Foundation hosted at Media Development Authority of Singapore (Grant No.: MDA/IDM/2012/8/8-2 VOL 01).

9. REFERENCES

[1] E. Diener. Subjective well-being. *Psychological Bulletin*, 95(3):542–575, 1984.

[2] E. Diener, E. Sandvik, L. Seidlitz, and M. Diener. The relationship between income and subjective well-being: Relative or absolute? *Social Indicators Research*, 28(3):195–223, 1993.

[3] A. Grubshtein, N. Gal-Oz, T. Grinshpoun, A. Meisels, and R. Zivan. Manipulating recommendation lists by global considerations. In *Proceedings of the 2nd International Conference on Agents and Artificial Intelligence (ICAART'10)*, pages 135–142, 2010.

[4] R. Inglehart, R. Foa, C. Peterson, and C. Welzel. Development, freedom, and rising happiness: A global perspective (1981-2007). *Perspectives on Psychological Science*, 3(4):264–285, 2008.

[5] G. A. Kaminka. Curing robot autism: A challenge. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'13)*, pages 801–804, 2013.

[6] S. Liu, J. Zhang, C. Miao, Y.-L. Theng, and A. C. Kot. iCLUB: an integrated clustering-based approach to improve the robustness of reputation systems. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11)*, pages 1151–1152, 2011.

[7] C. Miao, Q. Yang, H. Fang, and A. Goh. Fuzzy cognitive agents for personalized recommendation. In *Proceedings of the 3rd International Conference on Web Information Systems Engineering (WISE'02)*, pages 362–371, 2002.

[8] M. J. Neely. *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan and Claypool Publishers, 2010.

[9] D. W. Sacks, B. Stevenson, and J. Wolfers. Subjective well-being, income, economic development and growth. *The National Bureau of Economic Research Working Paper*, (16441):1–53, 2010.

[10] S. Sen. A comprehensive approach to trust management. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'13)*, pages 797–800, 2013.

[11] W. T. L. Teacy, J. Patel, N. R. Jennings, and M. Luck. Coping with inaccurate reputation sources: Experimental analysis of a probabilistic trust model. In *Proceedings of the 4th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'05)*, pages 997–1004, 2005.

[12] G. Vogiatzis, I. MacGillivray, and M. Chli. A probabilistic model for trust and reputation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)*, pages 225–232, 2010.

[13] H. Yu, C. Miao, B. An, C. Leung, and V. R. Lesser. A reputation management approach for resource constrained trustee agents. In *Proceedings of the 23rd International Joint Conference on Artifical Intelligence (IJCAI'13)*, pages 418–424, 2013.

[14] H. Yu, Z. Shen, C. Leung, C. Miao, and V. R. Lesser. A survey of multi-agent trust management systems. *IEEE Access*, 1(1):35–50, 2013.

[15] H. Yu, Z. Shen, C. Miao, and B. An. A reputation-aware decision-making approach for improving the efficiency of crowdsourcing systems. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'13)*, pages 1315–1316, 2013.

[16] D. Zhang. Taobao seller dies of exhaustion. *http://english.cntv.cn/program/china24/20121104/104649.shtml*, Dec 2012.