Effect of Probabilistic Task Allocation Based on Statistical Analysis of Bid Values*

(Extended Abstract)

Toshiharu Sugawara
Computer Science and Engineering
Waseda University, Tokyo, Japan
sugawara@waseda.jp

Kensuke Fukuda
National Institute of Informatics/JST
Tokyo, Japan
kensuke@nii.ac.jp

Toshio Hirotsu
Faculty of Computer and Information Sciences
Hosei University, Tokyo, Japan
hirotsu@hosei.ac.jp

Satoshi Kurihara
Institute of Scientific and Industrial Research
Osaka University, Osaka, Japan
kurihara@ist.osaka-u.ac.jp

ABSTRACT

This paper presents the effect of adaptively introducing appropriate strategies into the award phase of the contract net protocol (CNP) in a massively multi-agent system (MMAS).

Categories and Subject Descriptors

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

General Terms

Experimentation

Keywords

Task Allocation, Massively Multi-Agent Systems, Coordination, Contract Net Protocol

1. INTRODUCTION

Although recent Internet technologies support advanced large-scale applications, such as e-commerce, grid computing, distributed computing and cloud computing, these applications still require flexible controls for complex *massively multi-agent systems* (MMAS) in order to (1) process large number of sophisticated requests in a timely way and (2) effectively use computer and information resources. In particular, proper allocation of tasks to agents is a key aspect to exploiting the capabilities of the entire system.

We first focus on task allocation using CNP because it is used in many applications. Because we are interested in the performance of MMAS, we assume that bids from contractors include values reflecting their abilities and workload, more precisely, the estimated time to complete the announced task.

*This work is supported by KAKENHI 20240015.

Cite as: Effect of Probabilistic Task Allocation Based on Statistical Analysis of Bid Values (Extended Abstract), Toshiharu Sugawara, Kensuke Fukuda, Toshio Hirotsu, Satoshi Kurihara, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1603-1604

The aim of our research is to investigate the overall performance of an MMAS when tasks are allocated using CNP and how it changes when a variety of manager-side controls are introduced in the announcement and award phases. We introduce the *probabilistic awardee selection strategy*, under which awardee is selected with certain probabilities based on bid values. We found that, by changing the award strategies according to the local workload, the overall performance can be considerably improved for a specific task consisting of a number of subtasks.[1] We extend it for more general task structures whose subtasks have a number of different costs and discuss *adaptive probabilistic awardee selection* in which the probabilistic award strategy and the conventional award strategy are selected alternatively according to the estimated local workloads. It can outperform the naive CNP under any workload for various types of tasks.

2. BRIEF MODEL OF AGENTS AND TASKS

Let $\mathcal{A} = \{1, \ldots, n\}$ be a set of agents and $\mathcal{T} = \{t_1, \ldots, t_l\}$ be a task which consists of a number of subtasks t_i (we assume $|\mathcal{T}| = l = 2$, hereafter). Agent i is expressed as a tuple, $(\alpha_i, L_i, S_i, Q_i)$, where $\alpha_i(\geq 0)$ is the agent's capability, L_i is the location of i, and Q_i is the queue where the agent's tasks are stored, waiting to be executed one by one. The maximum queue length is assumed to be 20. The set $S_i(\subset \mathcal{A})$ is i's scope, i.e., the set of agents that i knows. The *metric* between agents, $\delta(i,j)$, is based on their locations, L_i and L_j, and is used to define the communication time of messages between i and j.

Subtask t has an associated cost, $\gamma(t)$, which is the cost to complete it. Subtask t can be done by i in $\lceil \gamma(t)/\alpha_i \rceil$ time units. The time it takes to complete t is also called the *execution time* of t by i. $\mathcal{T}$ is completed when all its subtasks are completed.

In every unit time, $\mathcal{L}(\geq 0)$ tasks on average are generated according to a Poisson distribution and randomly assigned to different managers. The parameter $\mathcal{L}$ is called the *task load* and means $\mathcal{L}$ tasks per unit time, or simply $\mathcal{L}$ T/t.

For CNP, we define $\mathcal{M} = \{m_j\}(\subset \mathcal{A})$ as the set of managers who allocate tasks and $\mathcal{C} = \{c_k\}(\subset \mathcal{A})$ as the set of contractors who execute the allocated tasks. Let us assume that $|\mathcal{A}|$ is large (on the order of thousands); therefore, $|\mathcal{M}|$

and $|\mathcal{C}|$ are also large; Moreover, we shall assume that the agents are widely distributed, like servers on the Internet.

For the sake of efficiency, we used a modified CNP in which multiple bids and *regret* and *no-bid* messages are allowed. When manager m receives $\mathcal{T}$, it immediately initiates the modified CNP for each task $\tilde{t}(\in \mathcal{T})$. The contractor receiving the announcement of $\tilde{t}$ sends back bid message with a certain *bid value* containing the estimated required times for completing the task. Then, m selects a contractor, as the awardee, on the basis of an award strategy and sends the awardee a message with the announced task.

We assume that manager agents can observe, for each subtask t, the *completion time*, which is the elapsed time from the time the award message is sent, $s(t)$, to the time the message indicating that the subtask has been completed is received, $e(t)$. The completion time thus includes the communication time in both directions, the queue time, and the execution time. The completion time of $\mathcal{T}$ is defined as $\max_{t\in\mathcal{T}}(e(t)) - \min_{t\in\mathcal{T}}(s(t))$. The *overall performance*, which is the average of the completion times observed by all managers, is used as the system's performance measure.

3. USAGE OF PROBABILISTIC AWARD

We set $|\mathcal{C}| = 500$ and $|\mathcal{M}| = 10,000$ in our simulation. Other settings are identical to the one in [1].

We express the cost structure of subtasks by the superscript of $\mathcal{T}$, if necessary. For example, $\mathcal{T}^{25-5}$ consists of two subtasks, $\{t_1, t_2\}$ such that $\gamma(t_1) = 2500$ and $\gamma(t_2) = 500$. Contractor c_i is assigned different capabilities such that the values of $2500/\alpha_{c_i}$ $(c_i \in \mathcal{C})$ will be *uniformly distributed* over the range 20–100. We assume that manager agents can not do the tasks themselves.

A small number of high-capability agents that receive multiple awards will likely bear an excessive workload whenever many managers simultaneously announce numerous tasks. A simple award strategy to alleviate the burden of too many awards is to allocate some tasks to the non-best contractor by introducing a probability in the award phase.

Let $\{c_1, \ldots, c_p\}$ be contractors bid on the announced task. We denote the bid value from contractor c_i as b_{c_i}. In naive CNP, m selects the contractor who submitted the best bid (a smaller bid is better). The first award strategy selects the awardee according to the following probability:

$$\Pr(c_i) = \frac{1/(b_{c_i})^k}{\sum_{j=1}^{p} 1/(b_{c_j})^k}. \quad (1)$$

This *probabilistic awardee selection* strategy is denoted by PAS_k. k is a variable called the *fluctuation factor*, or simply the *f-factor*. The PAS_∞ corresponds to the naive CNP.

4. ADAPTIVE STRATEGY BASED ON BID STATISTICS

Assume that, for announced task $t \in \mathcal{T} = \{t_1, t_2\}$. manager m received bids whose values are $B_m(t) = \{b_1(t), b_2(t), \ldots\}$. Let the SD of $B_m(t)$ be denoted by $SD_m(t)$, and $D_m^{SD}(\mathcal{T})$ be $|SD_m(t_1) - SD_m(t_2)|$.

The algorithm for selecting the the f-factor is listed in Fig. 1. First, manager m calculates the SDs of bid values for each $t_i \in \mathcal{T}$ and the difference between these SDs (denoted by $D_m^{SD}(\mathcal{T})$). It also retains the maximum and minimum values of $D_m^{SD}(\mathcal{T})$ (denoted by *maxSDdiff*, and *minSDdiff*) that have been obtained thus far. It estimates the current task load using *maxSDdiff*, *minSDdiff*, and $D_m^{SD}(\mathcal{T})$. We call this award strategy *adaptive probabilistic awardee selection*, or *APAS*.

Initialize:
  $maxSDdiff = 0$, $minSDdiff = minMaxAv = \infty$.

for each $\mathcal{T} = \{t_1, t_2\}$
  Manager m announces tasks t_1 and t_2 to local contractors[1], and m calculates the average value, $Av_m(t_i)$, and the SD, $SD_m(t_i)$, of bid values for t_i.

  /* Then it calculates some statistical values. */
  $\overline{Av}_m(\mathcal{T}) \stackrel{\leftarrow}{=} \max\{Av_m(t_2), Av_m(t_2)\}$;
  $\overline{SD}_m(\mathcal{T}) \stackrel{\leftarrow}{=} \max\{SD_m(t_1), SD_m(t_2)\}$;
  $\underline{SD}_m(\mathcal{T}) \stackrel{\leftarrow}{=} \min\{SD_m(t_1), SD_m(t_2)\}$;
  $D_m^{SD}(\mathcal{T}) \stackrel{\leftarrow}{=} |SD_m(t_1) - SD_m(t_2)|$;
  $minMaxAv \stackrel{\leftarrow}{=} \min(minMaxAv, \overline{Av}_m(\mathcal{T}))$;

  /* If the system is not so busy, */
  if $(minMaxAv \times \alpha > \overline{Av}_m(\mathcal{T}))$\{ /* Condition (1) */
    $maxSDdiff \stackrel{\leftarrow}{=} \max(maxSDdiff, \overline{SD}_m(\mathcal{T}))$;
    $minSDdiff \stackrel{\leftarrow}{=} \min(minSDdiff, \underline{SD}_m(\mathcal{T}))$;
  \}

  /* Defining threshold values: */
  $Th = \varepsilon maxSDdiff + (\varepsilon - 1) minSDdiff$;
  /* where $0 < \varepsilon < 1$. */

  /* Then output PAS_k by following the rule: */
  if $(D_m^{SD}(\mathcal{T}) \geq Th)$ $k = \infty$;
  else $k = 3$;

Figure 1: Outline of the APAS strategy.

Parameter α and variable *minMaxAv* are referred to in order to determine whether *maxSDdiff* and *minSDdiff* should be revised. The constant ε in the figure is used to define the threshold Th to switch between award strategies. In our experiments, we chose $\varepsilon = 0.58$ and $\alpha = 1.5$ on the basis of the average D_m^{SD} and the SDs of the preliminary experiment (which is not shown in this paper due to page limitations). APAS is quite simple in which only PAS_3 or PAS_∞ is alternatively selected.

5. CONCLUSION

We proposed a probabilistic award strategy in CNP for a massively MAS to elicit the potential capabilities of all agents. In this strategy, a manager agent (a) announces subtasks, (b) statistically analyzes the bids for each of these, (c) estimates the current local task load, and (d) introduces an adaptive degree of fluctuation in the award phase. We experimentally demonstrated that this strategy provides considerably better performance than the naive CNP.

6. REFERENCES

[1] T. Sugawara, T. Hirotsu, S. Kurihara, and K. Fukuda. "controling contract net protocol by local observation for large-scale multi-agent systems". In *Cooperative Information Agents XII (CIA2008)*, pages 206–220. LNCS 5180, Springer, May 2008.