Delegating Responsibility in a Multiagent Process Management System

John Debenham
Faculty of IT
University of Technology, Sydney
NSW 2007, Australia
+61-2-9514-1837
debenham@it.uts.edu.au

ABSTRACT
In a multiagent process management system the distribution of work is achieved by negotiated delegation of responsibility for sub-processes by one agent to another. The responsibility delegation mechanism is based on a combination of estimates for subjective and objective payoff measures. This leads to estimates of the probability that one agent is a better choice than another. The probability of delegating responsibility to an agent is then expressed as a function of these probability estimates. This apparently convoluted probabilistic method is easy to compute and gives good results in process management applications even when successive payoff measurements are unpredictably varied.

Categories and Subject Descriptors
I.2.11 [Distributed Artificial Intelligence]: *Coherence and coordination, Multiagent systems.*

General Terms
Management, Measurement, Performance, Design, Reliability.

Keywords
Delegation, responsibility, trust

1. INTRODUCTION
The responsibility delegation mechanism described is based on a combination of subjective and objective payoff measures that give estimates of the expected relative value in delegating responsibility to one agent or another. This leads to estimates of the probability that one agent is a better choice than another. The probability of delegating responsibility to an agent is then expressed as a function of these probability estimates. This method defines one set of probabilities in terms of another. It is easy to compute and gives good results in process management applications even when successive payoff measurements are unpredictably varied.

The method has general application to multiagent negotiation in areas other than process management. For example in electronic business agents place a subjective value on other agents that recognises the value of those agents as business associates. The way in which this subjective value measure is combined with objective payoff measures enables an agent to express its preferences on whether to place its business with a valued associate or to chance an offer from a less-trusted agent that appears to present a more attractive deal.

AAMAS'03, July 14–18, 2003, Melbourne, Australia.
ACM 1-58113-683-8/03/0007.

2. DELEGATION
A *delegation strategy* is a mechanism for deciding who to give responsibility to for doing what. If agent X_0 wishes to delegate responsibility then: first X_0 announces a proposal to a focussed subset of n agents in its community $\{X_1,...,X_i,...,X_n\}$, second X_0 receives bids from these n agents, and third X_0 chooses an agent from this set. The strategies considered here achieve this indirectly by determining instead n probabilities $\{P_1,...,P_i,...,P_n\}$ where P_i is the probability that the i'th agent will be selected, and $S_i P_i = 1$. The choice of the agent to delegate to is then made with these probabilities. By expressing the delegation strategy in terms of probabilities, the agents have the flexibility to balance conflicting goals, such as achieving process quality and process efficiency.

$Pr(X_i \gg)$ denotes the *rank* of agent X_i. Rank is "the probability that agent X_i is the 'best' choice of agent, chosen from $\{X_1,...,X_i,...,X_n\}$, to delegate responsibility to". A delegation *strategy* is a set $\{P_1,...,P_i,...,P_n\}$ where $S_i P_i = 1$. The delegation strategies described here are determined by:

$$P_i = f(Pr(X_i \gg))$$

for some function f that preserves the constraint $S_i P_i = 1$.

The probabilities $Pr(X_i \gg)$ are calculated at the time at which the delegation is made. They are based on various estimates of future performance that are combined to give a single *expected payoff* vector for each agent $\underline{n}_i$. The payoff vector contains sufficient information to estimate the probability that one agent is expected to deliver higher payoff than another in some sense.

3. PAYOFF: { n_i }
There are five measures for agent X_0. Three are: *time*, *cost* and *likelihood of success* which are attached to all of its plans and sub-plans. The remaining two are *value* and a *delegate* parameters that are attached to other agents. *Time* is the total time taken to termination. *Cost* is the actual cost of the of resources allocated. For example, if an agent has a virtual document in its 'in-tray' then the time observation will be the total time that that document spent with that agent, and the cost may derived from the time that the agent—possibly with a human 'assistant'—actually spent working on that document. [Note: cost here does

not refer to costs incurred by the plan—this is considered in the eBusiness applications described in Sec. 7.] The *likelihood of success* is the probability that a plan will terminate successfully within its constraints. The *value* parameter is the value added to a process by a plan. Each agent represents the perceived subjective *value* of each other agent's work as a constant value for that agent.

The three measures *time*, *cost* and *likelihood of success* are recorded *every* time a plan or delegation is activated for a goal. This generates a large amount of data whose significance can reasonably be expected to degrade over time. Rather than record the raw data it is summarised using the geometric mean. Given a set of observations $\{ ob_i \}$ where ob_1 is the most recent observation:

$$\sum_{i=1}^{n} \alpha^{i-1} \_ ob_i \; / \; \sum_{i=1}^{n} \alpha^{i-1}$$

is the geometric mean where a is some constant, $0 < a < 1$. If the observations $\{ ob_i \}$ for some parameter p are drawn from a symmetrically distributed population then the geometric mean gives a point estimate of the mean of the population m_p:

$$\sum_{i=1}^{n} \alpha^{i-1} \_ | ob_i - \mu_p | \; / \; \sum_{i=1}^{n} \alpha^{i-1}$$

is a (geometric) estimate of $\sqrt{2/p}$ times the standard deviation of parameter p, s_p. Where the constant a is determined empirically.

We now assume that the parameters *time* and *cost* are normally distributed. This is "not unreasonable", and is highly desirable because the geometric means may be updated with the simple formulae:

$$\mu_{p_{new}} = (1 - \alpha) \_ ob_i + \alpha \_ \mu_{p_{old}}$$

$$\sigma_{p_{new}} = (1 - \alpha) \_ | ob_i - \mu_{p_{old}} | + \alpha \_ \sigma_{p_{old}}$$

with starting values $m_{p_{initial}}$ and $s_{p_{initial}}$. The likelihood of success observations are binary—ie "success" or "fail"—and so the *likelihood of success* parameter is binomially distributed, which is approximately normally distributed under the standard conditions.

Finally, consider measurements of the *delegate* parameter for each agent. This parameter is the pair: w_i^{in} is the amount of work delegated *to* agent i in a given discrete time period, and, w_i^{out} is the amount of work delegated *by* agent i in the same discrete time period:

$delegate_{new} = (1 - a) \_ w_i + a \_ delegate_{old}$. The two components of the *delegate* parameter are not normally distributed and the standard deviation is not estimated. The *delegate* and *value* estimates are associated with individuals. The *time*, *cost* and *likelihood of success* estimates are attached to plans and delegations.

4. RANK: { Pr(X$_i$ ») }

A bid consists of the five pairs of real numbers (Constraint, Delegate, Success, Cost, Time). The pair *Constraint* is an estimate of the earliest time that the agent could address the task—ie ignoring other non-urgent things to be done, and an estimate of the time that the agent would normally address the task if it "took its place in the queue".

The method described above estimates the probability Pr(A»B) that one agent, A, is a better choice than another, B. It may be extended to estimate the probability that one agent is a better choice than a number of other agents. For example, if there are three agents to choose from, A, B, and C, then for some t_A in the interval [0, 1]:

$$\Pr(A\,») = \Pr(A » B) \_ \Pr(A » C) + t_A \_ \left[\min[\Pr(A » B), \Pr(A » C)] - \Pr(A » B) \_ \Pr(A » C) \right]$$

To proceed assume that: $t_A = t_B = t_C = t$:

$$t = \frac{1 - d}{q - d} \quad \text{where:}$$

$$d = [\Pr(A»B)\_\Pr(A»C) + \Pr(B»C)\_\Pr(B»A) + \Pr(C»A)\_\Pr(C»B)]$$

$$q = \left[\min[\Pr(A»B), \Pr(A»C)] + \min[\Pr(B»C), \Pr(B»A)] + \min[\Pr(C»A), \Pr(C»B)] \right]$$

5. STRATEGY: { P$_i$ }

Given a sub-process, an expectation of the payoff n_i as a result of choosing X_i from $\{X_1,...,X_i,...,X_n\}$ to take responsibility for it, and of the probability $\Pr(X_i\,»)$ that X_i is the best choice. A *delegation strategy* at any given time is a set $S = \{P_1,...,P_i,...,P_n\}$ where P_i is the probability of delegating responsibility at that time for a given task to agent X_i chosen from $\{X_1,...,X_i,...,X_n\}$.

If community culture is to choose the agent whose expected payoff is maximal then the delegation strategy *best* is:

$$P_i = \begin{cases} \frac{1}{m} & \textit{if } X_i \text{ is such that } \Pr(X_i\,») \text{ is maximal} \\ 0 & \textit{otherwise} \end{cases}$$

The strategy *best* attempts to maximise expected payoff. Another strategy *prob* also favours high payoff but gives all agents a chance, sooner or later, and is defined by $P_i = \Pr(X_i\,»)$. The strategies *best* and *prob* have the feature of 'rewarding' quality work (ie. high payoff) with more work.

6. REFERENCES

[1]. Debenham, JK. Managing e-Market Negotiation in Context with a Multiagent System. in proceedings Twenty First International Conference on Knowledge Based Systems and Applied Artificial Intelligence, ES'2002: Applications and Innovations in Expert Systems X, Cambridge UK, December 2002.

[2]. Dellarocas, C. Contractual Agent Societies: Negotiated shared context and social control in open multi-agent systems. Workshop on Norms and Institutions in Multi-Agent Systems, 4th International Conference on Multi-Agent Systems (Agents-2000), Barcelona, Spain, June 2000

[3] Jennings, N.R., Faratin, P., Norman, T.J., O'Brien, P. and Odgers, B. Autonomous Agents for Business Process Management. Int. Journal of Applied Artificial Intelligence 14 (2) 145-189 (2000).

[4]. Kumar, A., van der Aalst, WMP, and H.M.W. Verbeek. Dynamic Work Distribution in Workflow Management Systems: How to Balance Quality and Performance? Journal of Management Information Systems, 18(3):157-193, 2002.

[5]. Wooldridge, M. Multiagent Systems. Wiley, 2002.

aamas

PROCEEDINGS OF

THE SECOND INTERNATIONAL JOINT CONFERENCE ON AUTONOMOUS AGENTS AND MULTIAGENT SYSTEMS

July 14-18, 2003
Melbourne, Australia

Editors:

Jeffrey S. Rosenschein
Tuomas Sandholm
Michael Wooldridge
Makoto Yokoo

acm PRESS

Sponsored by:

ACM / SIGART

The International Foundation for Multiagent Systems (IFMAS)

The International Workshop on Agent Theories, Architectures, and Languages (ATAL)

The International Conference on Autonomous Agents (AGENTS)

The Association for Computing Machinery
1515 Broadway
New York, New York 10036

Notice to Past Authors of ACM-Published Articles

ACM intends to create a complete electronic archive of all articles and/or other material previously published by ACM. If you have written a work that has been previously published by ACM in any journal or conference proceedings prior to 1978, or any SIG Newsletter at any time, and you do NOT want this work to appear in the ACM Digital Library, please inform permissions@acm.org, stating the title of the work, the author(s), and where and when published.

ISBN: 1-58113-683-8

Additional copies may be ordered prepaid from:

ACM Order Department
PO Box 11405
New York, NY 10286-1405

Phone: 1-800-342-6626
(US and Canada)
+1-212-626-0500
(all other countries)
Fax: +1-212-944-1318
E-mail: acmhelp@acm.org

ACM Order Number 607033
Printed in the USA

Foreword

The AAMAS conference series was initiated in 2002, with the goal of providing a single, high profile, internationally respected and recognised forum for research in the theory and practice of autonomous agents and multiagent systems. The first AAMAS conference (AAMAS-2002, Bologna, Italy) attracted a remarkable number of submissions and nearly 700 delegates, firmly establishing it as *the* major event in the academic history of agent systems to date. We expect that the 2003 conference, in the attractive and cosmopolitan setting of Melbourne, Australia, will build on the successes and strengths of the 2002 conference, and will confirm AAMAS as a key event on the international computing research calendar.

AAMAS-03 received 466 submissions, from 30 countries across the globe. The 33 members of the senior program committee recruited 178 program committee members to handle the reviewing process. Each paper was reviewed by at least three program committee members, with some submissions selected for publication as full papers, and some selected for presentation as posters. The acceptance rate for full papers was 24.7%: low enough to ensure high quality, yet high enough to include a variety of topics and perspectives. The acceptance rate for full papers and posters together was 56.9%. Full papers were accepted from 19 countries, and posters from 21 countries. This proceedings volume includes full papers, and poster summaries.

The AAMAS conference is a merger of three highly successful related events:

- The International Conference on Autonomous Agents (AGENTS);
- The International Conference on Multi-Agent Systems (ICMAS); and
- The International Workshop on Agent Theories, Architectures, and Languages (ATAL).

We trust that these proceedings will do justice to the rich scientific and technological heritage of these three founding organisations, and we hope that you will enjoy reading the proceedings as much as we enjoyed preparing them.

Acknowledgments

We would like to take this opportunity to thank everyone involved with the organisation of AAMAS-03. First, we would like to thank Mike Luck for his efforts as finance chair. Managing the finances of an event as large and complex as AAMAS is no small feat: it requires the ongoing scrutiny of a complex budget, and involves continually making difficult decisions, in the knowledge that mistakes will have serious repercussions for the event. Mike put a phenomenal amount of energy into handling the finances, and as a result, not only made our lives much simpler, but also contributed to the success of the conference in no small way.

While we are on the subject of "where would we be without them", we have to acknowledge the incredible work of the local organisation team, headed by Lin Padgham and Liz Sonenberg. We could not have asked for a more helpful and proactive ground crew. We should single out Iyad Rahwan for special mention: Iyad designed the WWW site, and did a job that compares favourably with the best work of professional web designers.

Simon Parsons and Franco Zambonelli did excellent jobs of handling workshops and tutorials respectively, while Andrea Omicini handled student scholarships, and Sandip Sen handled the doctoral mentoring program. Dave Shield at Liverpool handled the technical aspects of the conference web site, and in particular, spent a great deal of time editing and preparing the ConfMan scripts for the online submission and reviewing process. This was the first year that AAMAS tried out online submission, and we felt it worked smoothly – but all the more smoothly due to Dave's technical support.

We would also like to thank our other sponsors, without whose support AAAMAS-03 would have been a significantly less interesting event. The enthusiastic support of so many sponsors (listed in full elsewhere in this volume) is a good indicator of how seriously the world is taking agent technology. AAMAS-03 was able to offer over US$40,000 in travel support to students, and this was possible only because David Parkes, Monique Calisti, David Kinny, and Von Wun Soo put in so much hard work contacting possible sponsors.

Finally, we would like to extend our thanks to the steering committees of the three "owner" conferences, and in particular, Ed Durfee, Les Gasser, Victor Lesser, and Joerg Muller, who all gave us timely, useful, and pragmatic advice.

Jeffrey S. Rosenschein and Michael Wooldridge (General Chairs)
Tuomas Sandholm and Makoto Yokoo (Program Co-chairs)