Relaxed-criteria G-negotiation for Grid Resource Co-allocation (Position Paper)

KWANG MONG SIM
Department of Computer Science, Hong Kong Baptist University,
Kowloon Tong, KLN, Hong Kong. Email: *prof_sim_2002@yahoo.com*

Supporting resource ***co-allocation*** is essential for realizing the Grid vision because computational intensive applications may require more resources than a single computing machine can provide in one administrative domain. Given that the various stakeholders often have their own requirements and supply-and-demand patterns, successfully obtaining commitments through concurrent negotiations with multiple resource owners to simultaneously access several resources is a very challenging task. This position paper (i) suggests that a *relaxed-criteria G-negotiation mechanism* may be used to enhance the success rates of negotiation agents in Grid resource co-allocation, (ii) proposes a testbed for realizing the relaxed-criteria G-negotiation mechanism, (iii) discusses experimental designs for evaluating the proposed mechanism, and (iv) discusses previous results obtained from using a relaxed-criteria G-negotiation protocol for Grid resource allocation.
Additional Key Words and Phrases: Grid Resource co-allocation, Grid Resource Management, Automated Negotiation, Bargaining Protocol, Resource allocation, Software Agent

1. INTRODUCTION

A Grid resource management system should bolster *co-allocation* of computing resources (i.e., allocating to an application *multiple resources belonging to* possibly *different administrative domains*) [1, p.2]. Supporting resource co-allocation is essential for realizing the Grid vision because (i) computationally intensive applications may require more resources than a single computing machine can provide in one administrative domain [2, p1], and (ii) an application may require several types of computing capabilities from resource owners in other administrative domains. Providing effective Grid resource co-allocation mechanisms will have a significant impact on facilitating smooth operations of computational Grids executing computationally intensive applications on computing resources across multiple administrative domains. Both (i) establishing multiple contracts (or agreements) between applications and providers through negotiation, and (ii) coordinating simultaneous multiple resources utilization, are essential functionalities of a Grid resource co-allocation mechanism. However, to date, there are very few works on Grid resource negotiation [3–12] (see [13] for a survey), and among these few extant works on Grid resource negotiation, only *SNAP* (*Service Negotiation and Acquisition Protocol*) [3] has considered the issue of Grid resource co-allocation, but the protocols in [4–12] do not address the issue of coordinating resource utilization. Nevertheless, unlike [4–12] which considered strategies for optimizing utilities of Grid participants, *SNAP* [3] only searches for the solutions for satisfying simultaneous multiple resources requirements of Grid consumers, and does not focus on specifying the negotiation strategies nor the protocols that Grid participants should adopt for their negotiation activities.

Successfully obtaining contracts from multiple resource owners for simultaneously accessing several resources is a very challenging task given that the stakeholders often

Permission to make digital/hard copy of part of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date of appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.
© 2006 ACM 1529-3785/2004/0700-0024 $5.00

have different requirements and supply-and-demand patterns. Making allocation to multiple resources in a coordinated fashion across virtual organizational boundaries is not only an important problem in Grid resource management, but is also a very difficult problem. For instance, mapping application workflows consisting of interacting components that need to be executed in a certain partial order to Grid resources is an NP-complete problem [14]. One way of solving such a hard problem is to use heuristic approaches, and this position paper proposes novel heuristics that will guide G-negotiation agents in overlooking slight differences in their bargaining terms (e.g., acquiring a slightly more expensive resource when many resources are already occupied) in the hope of increasing the chance of acquiring all required resources and acquiring them more rapidly. Although still in its infancy, this work builds on the author's work on *relaxed-criteria G-negotiation* (section 2) and Czajkowski et al.'s work on *SNAP* [3] (section 3). The impetus of this work is proposing a relaxed-criteria negotiation mechanism for supporting Grid resource **co**-allocation by *bolstering multiple concurrent pairs of negotiations simultaneously* and *coordinating the concurrent negotiations* (section 4). In this position paper, a testbed for simulating resource planning, concurrent negotiation activities, and coordination of multiple resources utilization in a Grid-computing environment is proposed (section 4.1). Experimental designs for evaluating the relaxed-criteria negotiation mechanism are discussed in section 4.2. Section 5 concludes this position paper by discussing the suitability of adopting a relaxed-criteria G-negotiation protocol for enhancing agents' success rates in Grid resource co-allocation.

2. RELAXED-CRITERIA G-NEGOTIATION

The G-negotiation protocol used in [7–8] enhances the *alternating offers protocol* [15] by slightly relaxing the criteria for agents to reach a consensus. In the alternating offers protocol and also in most negotiation models (e.g., [10–12,16], only to name a few because of space limitation), a pair of negotiation agents (B_1, S_1) reaches an agreement when one agent proposes a deal that matches (or exceeds) what another agent asks for (see *R1*). *R1* is relaxed in [7–8] where a G-negotiation agent *also* accepts another agent's (counter-)proposal if it is *sufficiently close* to its own proposal following *R2*.

R1: An agreement is reached if two agents B_1 and S_1 propose deals b_1 and o_1, respectively, such that either (i) $U(b_1)\geq U(o_1)$ or (ii) $U(o_1)\geq U(b_1)$, where U is a utility function mapping b_1 and o_1 to [0,1].

R2: An agreement is reached if either (i) $\eta=U(o_1)-U(b_1)$, such that $\eta\to 0$ or (ii) $\eta=U(b_1)-U(o_1)$, such that $\eta\to 0$, where η is the amount of relaxation determined using a fuzzy decision controller (*FDC*).

In Sim's relaxed-criteria G-negotiation protocol [7–8], *market-driven G-negotiation agents* (see section 2.1) representing resource providers and consumers are programmed to *slightly* relax their bargaining criteria under intense pressure (e.g., when a consumer has a higher demand for resources) with the hope of enhancing their chance of successfully acquiring resources. A consumer agent and a provider agent are both designed with an *FDC*: *FDC-C* and *FDC-P*, respectively. Two sets of relaxation criteria (for consumers and providers, respectively) that are specific to Grid resource management are used as inputs to *FDC-C* and *FDC-P*, respectively.

Consumers' relaxation criteria: Two criteria that can influence a consumer agent's decision in the amount of relaxation of bargaining terms are (i) recent statistics in failing/succeeding in acquiring resources called *failure to success ratio* (fs_t), and (ii) demand for computing resources called *demand factor* (df_t). If a consumer agent is *less* successful in acquiring resources recently to execute its set of tasks, it will be under more pressure to slightly relax its bargaining criteria with the hope of completing a deal. Furthermore, if it has a *greater* demand for computing resources it is more likely to be under more pressure to slightly relax its bargaining criteria. Both fs_t and df_t are inputs to *FDC-C*, which a consumer agent uses to determine η (its amount of relaxation) [7–8].

Providers' relaxation criteria: Two criteria that can influence a provider agent's decision are: (i) the amount of the provider's resource(s) being utilized (i.e., the *utilization level* (ul_t)), and (ii) recent requests from consumers for resources (i.e., called the *request factor* (rf_t)). If more of its resources are currently being used to execute its own tasks or are already leased to other consumers, then a provider is less likely to slightly relax its bargaining terms. If there are fewer recent demands from consumers to lease its resources, a provider is more likely to slightly relax its bargaining criteria since it is under more pressure to trade its idle resources. Both ul_t and rf_t are inputs to *FDC-P*, which a provider agent uses to determine η [7–8].

Empirical results obtained in [7–8] show that by slightly relaxing their bargaining terms under intense negotiation pressure, both consumer and provider agents generally achieved higher success rates in negotiation (without sacrificing much of their average utilities).

2.1. Market-driven G-negotiation Agents

Using a *market-driven strategy* [17–23], *market-driven G-negotiation agents* (*MDGAs*) [4–6] make *adjustable amounts* of concession in different market situations by considering factors such as time, opportunity, and competition. Sim[17,23] has proven that agents adopting a market-driven strategy negotiate optimally by making *minimally sufficient concession* in different market conditions (see lemmas 4.1 and 4.2 in [17,pp.718–719]). That is, they avoid making excessive concessions in favorable markets and inadequate concessions in unfavorable markets. An *MDGA* determines the appropriate amount of concession for a given market situation using a combination of three negotiation functions: time (***T***), competition (***C***), and opportunity (***O***) functions.

T models the intuition that as time passes, an *MDGA* relaxes its proposal by making attempt(s) to narrow its difference(s) with other parties using: $T(t,\tau,\lambda) = 1 - (t/\tau)^{\lambda}$ where t is the current trading time (measured in rounds), τ is the deadline, and λ is an *MDA*'s time preference. Negotiators with different time preferences may have different concession rates with respect to time. With infinitely many values of λ, there are infinitely many possible strategies in making concession with respect to remaining trading time, but they are classified into *Conservative, Linear* and *Conciliatory strategies* corresponding to making (i) larger concessions in the early trading rounds and smaller concessions at the later stage, (ii) a *constant* rate of concession, and (iii) smaller concessions in early rounds and larger concessions in later rounds.

C determines the amount of competition of an *MDGA* by determining the probability that it is not being considered as the most preferred party. Since *MDGAs* are utility maximizing agents, an *MDGA* is more likely to reach a consensus if its proposal is ranked the *highest* by some other agent. Suppose an agent ***B*** has m–1 competitors $\{B_2,\ldots, B_m\}$ and n trading partners $\{S_1,\ldots, S_n\}$. The probability that ***B*** is *not* the most preferred trading partner of *any* S_j (where $S_j \in \{S_1,\ldots, S_n\}$) is $(m-1)/m$. Hence, the probability that ***B*** is *not* the most preferred partner of *all* $S_j \in \{S_1,\ldots, S_n\}$ is $[(m-1)/m]^n$. In general, the probability that ***B*** is considered the *most preferred* trading partner by at *least one* of

$S_j \in \{S_1, \ldots, S_n\}$ is: $C(m,n) = 1 - \left[(m-1)/m\right]^n$ where m and n are, respectively, the numbers of buyer agents (including ***B***) and seller agents at round t.

O determines the amount of concession based on (i) trading alternatives (i.e., outside options or number of trading partners) and (ii) differences in utilities generated by the proposal of an *MDGA* and the counter-proposal(s) of its trading partner(s). When determining opportunity, it was shown in [17,23] that if there is a large number of trading alternatives, the likelihood that an agent proposes a bid/offer that is potentially close to an *MDGA*'s offer/bid may be high. However, it would be difficult for the *MDGA* to reach a consensus if none of the so many options are viable (i.e., there are large differences between the proposal of the *MDGA* and the counter-proposals of all its trading partners). On this account, ***O*** determines the probability of reaching a consensus at its own term by determining its bargaining position based on trading alternatives, differences between its proposal and others, and considering the notion of *conflict probability* [24].

3. SERVICE NEGOTIATION AND ACQUISITION PROTOCOL

To coordinate the utilization of multiple resources owned by different administrative domains, advance reservation of resources that specifies the time and duration of a resource capacity is essential [25, p.3]. Unlike generic e-commerce negotiation where a buyer-seller pair negotiates for a product/service, perhaps in a single negotiation phase, a Grid application may need to be engaged in a multi-phase negotiation process with multiple resource owners, to reserve, acquire, coordinate, schedule, and potentially renegotiate resource access. Dealing with negotiation of resource co-allocation and advance resource reservations requires more sophisticated negotiation protocols. One such protocol is *SNAP* (Service Negotiation and Acquisition Protocol) [3,25]. In *SNAP*, Grid participants negotiate a *service-level agreement* (*SLA*) in which a resource provider establishes a contract with a consumer to provide some measurable capabilities or to perform a task. Establishing a single *SLA* across a set of (simultaneously required) resources that may be owned and operated by different providers is very difficult. *SNAP* defines a resource management model in which consumers (i) can submit tasks to be performed, and (ii) get promises of capability (commitment from the providers or servers), and bind (i) and (ii). In *SNAP*, *SLAs* are classified into: *Resource SLAs* (*RSLAs*), *Task SLAs* (*TSLAs*), and *Binding SLAs* (*BSLAs*). In an *RSLA*, consumers negotiate with resource providers for the rights to consume a resource without specifying how the resource will be utilized. For example, an advance resource reservation takes the form of an *RSLA*, and it characterizes a resource in terms of its abstract service capabilities. In a *TSLA*, clients negotiate with resource providers for the performance of an activity or a task. For example, a *TSLA* is created by submitting a job description to a queuing system and it characterizes a task in terms of its service steps and resource requirements. In a *BSLA*, consumers negotiate with resource providers for the application of a resource to a task. A *BSLA* associates a task defined by a *TSLA* to a *RSLA*. However, *SNAP* [3,25] only searches for the solutions for satisfying (multiple) resource requirements of Grid consumers, and coordinates multiple resources utilization, but does not specify the negotiation strategies to be adopted or the basic negotiation protocol for Grid applications to send their terms of a desired *SLA* to a target resource provider. Section 4 proposes a *SNAP*-like coordination protocol with the specifications of agents' negotiation activities following a relaxed-criteria G-negotiation protocol.

4. RELAXED-CRITERIA NEGOTIATION AND RESOURCE CO-ALLOCATION

Complete details of the design of a negotiation mechanism for Grid resource co-allocation have not been fully formulated at this stage. One of the possible approaches for constructing such a mechanism is to incorporate into a *SNAP*-like coordination protocol (section 3) the detailed specifications of the negotiation activities between consumers and providers by taking into consideration the issues of enhancing negotiation success rates by relaxing bargaining criteria, optimizing utility, and modeling market dynamics (section 2).

With the view of adopting *MDGAs* (section 2.1) to act as intermediaries between Grid consumers and resource providers, a possible relaxed-criteria negotiation protocol (based on [7–8]) for specifying the negotiation activities among consumer agents and provider agents is given as follows:

- Negotiation proceeds in a series of rounds.
- Adopting Rubinstein's *alternating offers protocol* [15, p.100], a pair of consumer and provider agents negotiates by making proposals in alternate rounds.
- Multiple consumer-provider agent pairs can negotiate deals simultaneously.
- When an agent makes a proposal, it proposes a deal from their space of possible deals (e.g., consisting of the most desirable price, the least desirable (reserve) price, and those prices in between). Typically an agent proposes its most preferred deal initially [26].
- If no agreement is reached, negotiation proceeds to the next round. At every round, an agent determines its amount of concession using the ***T***, ***O***, and ***C*** functions [17–23] (see section 2.1).
- Negotiation between two agents terminates (i) when an agreement is reached, or (ii) with a conflict when one of the bargaining agents' deadline is reached.
- Agents follow one of the following rules for reaching an agreement:
  (i) An agreement is reached when an agent's trading party's offer matches or exceeds what it asked for (see *R1* in section 2).
  (ii) An agreement is reached if the offer is sufficiently close (albeit, it does not totally match the agent's negotiation terms). (see *R2* in section 2). However, consumer agents and provider agents will be designed with ***<u>different</u>*** *FDCs* since they will use ***<u>different</u>*** sets of relaxation criteria (from that of [8–9]) and ***<u>different</u>*** sets of fuzzy rules (from that of [7–8]) to determine the amount of relaxation.

Although a protocol for coordinating concurrent Grid resource negotiation has not been fully devised at this stage, echoing [13], this position paper suggests that a *SNAP*–like coordination protocol may be used. Using a *SNAP*–like coordination protocol, a consumer may achieve advance resource reservation and coordinate simultaneous access to multiple resources following a 4-state resource planning that is similar to *SNAP* shown in Fig. 1. The four states in resource planning consist of: S_0, S_1, S_2, and S_3. Note that in Fig. 1, a solid arrow represents a request (or action) by a consumer, and a dashed arrow represents an action or internal behavior of a resource provider. In S_0, *SLAs* have not been created or have been resolved by termination or cancellation of the *SLAs*. In S_1, both *RSLAs* and *TSLAs* have been agreed upon, but they are not matched with each other. There are 2 possible movements from S_0: (i) S_0 to S_1 (dashed arrow), and (ii) S_0 to S_0 (solid curly arrow). S_0 to S_0 represents the transition in which a consumer is still waiting to establish both *RSLAs* and *TSLAs* through negotiation activities between consumers and providers following the relaxed-criteria negotiation protocol described above. A movement from S_0 to S_1 (solid arrow) indicates that a client has successfully negotiated with resource providers to establish both *RSLAs* and *TSLAs*. There are 3 possible movements from S_1: (i) S_1 to S_0 (dashed arrow), (ii) S_1 to S_1 (solid curly arrow), and (iii) S_1 to S_2 (solid arrow). S_1 to S_0 represents the transition in which *SLAs* have been either cancelled by a resource provider or a client, or expired. S_1 to S_1 represents the transition in which a consumer is still waiting to establish the *BSLAs* (even though it has already established both *RSLAs* and *TSLAs*). S_1 to S_2 represents the transition of a client that has successfully negotiated with resource providers for the application of resources to tasks (i.e., successfully establishing *BSLAs*). In S_2, the *TSLA* is matched with the *RSLA*, and this binding represents a dependent *BSLA* to resolve the task. There are 3 possible movements from S_2: (i) S_2 to S_1 (dashed arrow), (ii) S_2 to S_2 (solid curly arrow) and (ii) S_2 to S_3 (dashed arrow). S_2 to S_1 represents the transition in which a resource provider moves the control back to the prior state because some fault has occurred and the task cannot be scheduled. S_2 to S_2 represents the transition that even though a *BSLA* has been established, a client is waiting for the task to be scheduled. S_2 to S_3 represents the scheduling of resources by a resource provider to satisfy a *TSLA*. In S_3, although resources are actively being utilized to support a task, they can still be controlled and changed (e.g., moving back to S_2 from S_3). Whereas the movement from S_3 to S_3 represents the transition of task execution (a client's task is being executed and it is waiting for the task to complete execution), S_3 to S_2 represents either task completion or faults in the execution so that the resource provider moves the control back to the prior state.

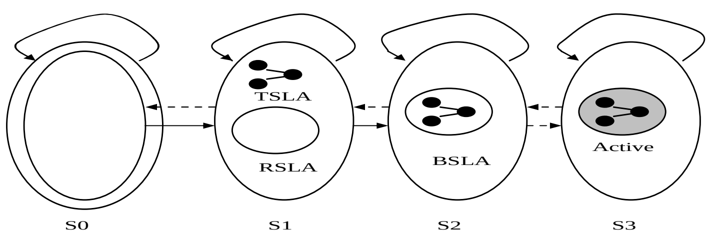

Fig. 1 State transitions of a *SNAP-like* coordination Protocol [3,25]

4.1. Possible Components in a Preliminary Testbed

While a testbed for simulating Grid resource co-allocation using a relaxed-criteria negotiation mechanism is still under construction, this position paper envisions that it will consist of (i) a set of *Grid resources*, (ii) a set of *resource consumers*, (iii) a set of *MDGAs* (*provider agents* and *consumer agents*) acting as an intermediary between resource providers and consumers, (iv) a *repository of resource information*, and (v) a Grid resource record agent (*GRRA*).

Grid resource providers/consumers: A Grid resource provider may possess a series of computing resources. Each computing resource can be formed by one or more processing elements, and each processing element can have different speeds measured in terms of MIPS (millions of instructions per second). Each resource consumer can have one or more jobs. Each job is characterized by a job length measured in MI (millions of instructions), length of input and output data, execution start and end times, as well as the originator of the job. In particular, each job may require multiple computing resources (owned by providers in different administrative domains) for its executions.

Provider agents and *consumer agents*: Provider agents act as resource brokers between computing resources and consumer agents which perform the function of acquiring multiple computing resources and coordinating job executions of consumers. Rather than submitting to the resources, jobs are submitted to provider agents, which schedule and coordinate the job activities in these resources. Consumer agents negotiate with provider agents to establish contracts for provisioning computing capabilities to perform tasks. While negotiation activities between consumer and provider agents are anticipated to be carried out following the specifications of the relaxed-criteria negotiation protocol, the coordination of the utilization of resource for job execution is expected to follow the *SNAP*-like coordination protocol.

Resource information: The resource information dictionary is a repository of information about the computing resources registered in the Grid. The *GRRA* updates the resource information dictionary when a resource joins/leaves the Grid.

4.2. Experimental Design

Following [7–8,18], two types of agents are anticipated to be used in future experiments: (i) Type 1 agents are either consumer agents or provider agents that are programmed to follow *only R1* when reaching an agreement (see section 2), and (ii) Type 2 agents are either consumer agents or provider agents that are programmed to follow *both R1* and *R2*, i.e., slightly relax their bargaining terms in the face of intense pressure (e.g., urgent need to acquire a resource or facing fast approaching deadlines). Whereas it is expected that Type 1 agents' negotiation activities will follow an alternating offers protocol (*AO*-protocol), Type 2 agents' will follow Sim's relaxed-criteria negotiation protocol (*RC*-protocol) [7–8,18] (section 2). For resource planning and coordinating multiple resources utilization, it is anticipated that both Type 1 and Type 2 agents will follow the *SNAP*-like coordination protocol described above. In summary, while Type 1 agents will follow a *SNAP-AO*-protocol, Type 2 agents will follow a *SNAP-RC*-protocol.

5. DISCUSSION AND CONCLUSION

This position paper attempts to answer the following questions:

(i) What relaxation heuristics could be used by G-negotiation agents in Grid resource co-allocation?

(ii) What are some of the possible classes of heuristics for determining the appropriate amount of relaxation to achieve both optimal utilities and optimal success rates?

(iii) Explain how preliminary empirical results in [7–8] show the suitability of adopting a relaxed-criteria G-negotiation protocol for enhancing the success rates of G-negotiation agents in Grid resource co-allocation?

Relaxation heuristics: It is anticipated that if a relaxed-criteria bargaining protocol [7–8] is applied to Grid resource co-allocation involving coordination of multiple concurrent negotiations, ***different*** *sets of heuristics and fuzzy rules should be adopted* for the G-negotiation agents. For instance, in addition to factors such as recent demand for resources and recent statistics of failing/succeeding in acquiring resources, perhaps, another factor that needs to be considered is the control/data dependencies among interacting components of a Grid application. It is noted in [14, p.2] that some important Grid applications fall into the category of *workflow applications* consisting of collections of several interacting components that need to be executed in certain partial order for successful execution of the application in its entirety. On this account, this position paper suggests that one of the research questions that needs to be answered is "What relaxation heuristics should a G-negotiation agent be programmed to follow when it is negotiating for multiple resources for a Grid application that has multiple interacting components with control/data dependencies". For example, should a G-negotiation agent be programmed so that it would be more likely to overlook small differences when bidding for a resource to execute a sub-task on which the control and executions of (many) other sub-tasks depend. Answering questions such as this, finding the appropriate set of relaxation heuristics, and devising the fuzzy rules for guiding G-negotiation agents when relaxing negotiation terms, are among the list of future agenda for this work.

Relaxation vs. utility optimization: Although relaxing bargaining terms slightly (at the expense of achieving slightly lower utility) may be desirable to enhance the success rates of acquiring computing resources (section 2), the problem of determining the appropriate amount of relaxation to achieve both optimal utilities and optimal success rates under different market conditions (e.g., given different resource alternatives and demands) and constraints (e.g., given different deadlines) remains open. This issue was raised by the author in a recent survey paper [13]. Whereas a previous work of the author on relaxed-criteria G-negotiation for Grid resource allocation [7–8] adopts *FDC-C* and *FDC-P* and two corresponding sets of fuzzy rules for guiding consumer and provider agents in relaxing bargaining terms, a future agenda of this research is searching for an appropriate approach for tuning the two sets of fuzzy rules to improve and perhaps optimize the performance of both consumer and provider agents. Even though this work is still in its infancy, this position paper suggests that there are two possible classes of heuristics: (i) *utility-optimizing* and (ii) *deal-optimizing*. A utility-optimizing (respectively, deal-optimizing) heuristics is characterized by making smaller (respectively, larger) amounts of relaxation, and may be more appropriate to be applied when a G-negotiation has longer (respectively, shorter) deadlines in acquiring resources. For instance, an agent that is urgently in need of acquiring computing resources rapidly may adopt a deal-optimizing heuristic to increase its chance of successfully obtaining the necessary resources before its deadline. Agents with more relaxed deadline constraints may adopt a utility-optimizing heuristic to increase its chance of acquiring less expensive resources.

Preliminary results: A series of stochastic simulations was carried out using the testbed in [7–8] to compare the performance of agents following (i) the *AO*-protocol and (ii) the *RC*-protocol for Grid resource negotiation *without* considering the issue of resource planning and coordination of multiple resources utilization. Four input parameters were used: (i) the capacity of the Grid, (ii) market density (corresponding to the arrival rate of consumers joining the Grid), (iii) deadlines for consumers to complete their set of tasks, and (iv) available budgets of consumers. Whereas the arrival rate of consumers follows a Poisson distribution, consumers' deadlines and budgets are randomly generated. Preliminary empirical results in [7–8] showed that agents following the *RC*-protocol (Type 2 agents) (i) are generally more successful in reaching agreements, and (ii) reached agreements with fewer negotiation rounds, than agents that followed the *AC*-protocol (Type 1 agents). Selected (preliminary) results from [7–8] are summarized in Fig. 2 in the appendix. The results in Fig. 2 show that, in general, agents following the *RC*-protocol (i.e., "with relaxation") achieved higher success rates and required fewer negotiation rounds to complete deals than agents that followed the *AC*-protocol (i.e., "without relaxation"). However, comparing the performance of Type 1 agents following the *SNAP-AO*-protocol and Type 2 agents following the *SNAP-RC*-protocol is a much more challenging task. Building on previous empirical results in [7–8] on using the *RC*-protocol for Grid resource allocation, a future agenda is to demonstrate the effectiveness of using the *SNAP-RC*-protocol for Grid resource co-allocation by comparing the performance of Type 1 agents following the *SNAP-AO*-protocol and Type 2 agents following the *SNAP-RC*-protocol.

Contribution: The contribution of this position paper include: (i) suggesting a possible mechanism for enhancing success rates of G-negotiation agents for Grid resource co-allocation, (ii) discussing what relaxation heuristics could be used and the possible classes of heuristics that may optimize agents' success rates and utilities, (iii) suggesting some of the possible high-level components that one may envision in a preliminary testbed for demonstrating the ideas of relaxed-criteria G-negotiation for Grid resource co-allocation, (iv) discussing the experimental designs for evaluating the proposed mechanism, and (v) discussing why a relaxed-criteria G-negotiation mechanism may enhance the success rates of agents in Grid resource co-allocation (based on preliminary results in [7–8]).

ACKNOWLEDGEMENTS

K. M. Sim gratefully acknowledges financial support for this work from the Hong Kong Research Grant Council, under project code: RGC/HKBU210906. The author would like to thank the referees for their comments and suggestions.

REFERENCES

[1] A. Ali et al. A Taxonomy and Survey of Grid Resource Planning and Reservation Systems for Grid Enabled Analysis Environment. Proc. of the 2004 Int. Sym. on Distributed Comp. and Appl. to Business Eng. and Science, Wuhan, China, pp. 1-8.

[2] R. Buyya, D. Abramson, and S. Venugopal, The Grid Economy. Proceedings of the IEEE, Volume 93, Issue 3, pp 698-714, IEEE Press, New York, USA, March 2005.

[3] K. Czajkowski, I. Foster, et al. *SNAP*: A Protocol for Negotiation of Service Level Agreements and Coordinated Resource Management in Distributed Systems. Job Scheduling Strategies for Parallel Processing: 8th Int. Workshop Edinburgh, 2002.

[4] K. M. Sim. From Market-driven Agents to Market-Oriented Grids. ACM SIGecom: E-commerce Exchanges, Vol. 5, No. 2, November 2004, Pages 45–53.

[5] K. M. Sim. G-Commerce, Market-driven G-Negotiation Agents and Grid Resource Management. IEEE Transactions on Systems, Man and Cybernetics, Part B, 2006, Vol. 36, No. 6, Dec. 2006, pp 1381-1394..

[6] K. M. Sim. From Market-driven e-Negotiation Agents to Market-driven G-Negotiation Agents. Proc. of the IEEE Int. Conf. on e-Technology, e-Commerce and e-Services, pp. 408-413, Hong Kong, 2005.

[7] K. M. Sim and K. F. Ng. A Relaxed-Criteria Bargaining Protocol for Grid Resource Management. In Work on Agent-based Grid Computing, held in conjunction with the IEEE International Symposium on Cluster Computing and the Grid 2006, Singapore.

[8] K. M. Sim and K. F. Ng. Relaxed-criteria negotiation for G-commerce. In Proc. of the Business Agents and the Semantic Web Workshop, 2006 held in conjunction with the Fifth Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems, Hokadate, Japan.

[9] H. Gimpel et al. PANDA: Specifying Policies for Automated Negotiations of Service Contracts In ICSOC 2003, LNCS 2910, pp. 287–302, 2003, Springer-Verlag.

[10] F. Lang. Developing Dynamic Strategies for Multi-Issue Automated Contracting in the Agent Based Commercial Grid. In Workshop on Agent-based Grid Economics, held in conjunction with the IEEE Int. Sym. on Cluster Computing and the Grid, 2005, UK.

[11] R. Lawley et al. Automated Negotiation between publishers and consumers of grid notifications. Parallel Processing Letters, 13(4):pp. 537-548, 2003.

[12] P. Ghosh et al. "A Game Theory Based Pricing Strategy for Job Allocation in Mobile Grids," Proc. Int. Parallel and Distributed Processing Sym., 2004.

[13] K. M. Sim. A Survey of Bargaining Models for Grid Resource Allocation. ACM SIGecom: E-commerce Exchanges, Vol. 5, No. 5, January, 2006, pages 22-32.

[14] Mandal et al. Scheduling Strategies for Mapping Application Workflows onto the Grid. HPDC-14. Proceedings of 14th IEEE International Symposium on High Performance Distributed Computing, 2005, July, 2005, pages 125- 134.

[15] A. Rubinstein, "Perfect equilibrium in a bargaining model," Econometrica, vol. 50, no. 1, pp. 97–109.

[16] P. Faratin et al. Negotiation Decision Functions for Autonomous Agents. Int. J. Robotics and Autonomous System. Vol.24, No.3: 159-182, 1998.

[17] K. M. Sim. Equilibria, Prudent Compromises, and the "Waiting" Game. IEEE Transactions on Systems, Man and Cybernetics, Part B: Cybernetics, Vol. 35, No. 4, Aug. 2005, pp. 712-724.

[18] K. M. Sim and S.Y. Wang. Flexible Negotiation Agent with Relaxed Decision Rules. IEEE Transactions on Systems, Man and Cybernetics, Part B, Vol. 34, No. 3., pp. 1602-1608, Jun., 2004.

[19] K. M. Sim and C.Y. Choi. Agents that React to Changing Market Situations. IEEE Transactions on Systems, Man and Cybernetics, Part B: Cybernetics, Vol. 33, No. 2, pp 188-201, April 2003.

[20] K. M. Sim and E. Wong. Towards Market-driven Agents for Electronic Auction. IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans, Vol. 31, No.6, pp 474-484, Nov. 2001.

[21] K.M. Sim. Negotiation Agents that make prudent compromises and are slightly flexible in reaching consensus. In Computational Intelligence, Special issue on Business Agents and the Semantic Web, Nov. 2004, pp. 643-662.

[22] K. M. Sim. A Market-driven Model for Designing Negotiation Agents. In Computational Intelligence, Special issue in Agent Technology for E-commerce, vol. 18, no. 4, 2002.

[23] K. M. Sim. Equilibrium Analysis of Market-driven Agents. ACM SIGecom: E-commerce Exchanges, vol. 4.2, Summer, 2003, pp. 32-40.

[24] J. C. Harsanyi. Bargaining, In The New Palgrave: Game Theory, edited by John Eatwell, Murray Milgate, and Peter Newman, 1st edition, The Macmillan Press Limited, 1989.

[25] K. Czajkowski, I. Foster, and C. Kesselman. Agreement-based Resource Management. Proceedings of the IEEE, Vol. 93, No. 3, pp 631-643, IEEE Press, USA, March 2005.

[26] M. Wooldridge. An Introduction to Multiagent Systems. 2002, John Wiley & Sons.

APPENDIX

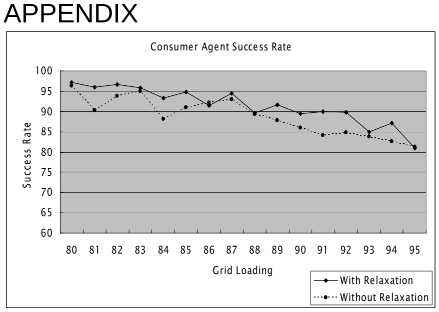

Fig. 2a. Success rates

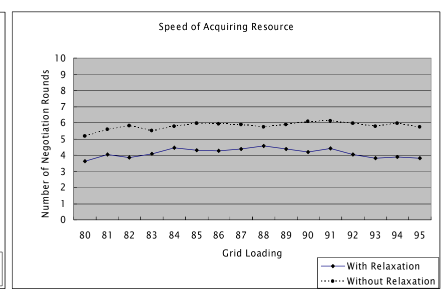

Fig. 2b. Number of negotiation rounds

Fig. 2. AC-Protocol vs. RC-Protocol [7-8]