

Relating Quantified Motivations for Organizationally Situated Agents^{*}

Thomas Wagner and Victor Lesser

Department of Computer Science
University of Massachusetts, Amherst, MA 01003
{wagner, lesser}@cs.umass.edu

Abstract. To scale agent technologies for widespread use in open systems, agents must have an understanding of the organizational context in which they operate. In this paper we focus on the issue of task valuation and action selection in socially situated or organized agents – specifically on the issue of quantifying agent relationships and relating work motivated by different sources.

1 Introduction

We believe that in order to scale-up agent technology [25] for use in open application domains, e.g., electronic commerce on the web, agents must model their organizational relationships with other agents and reason about the value or utility of interacting and coordinating with particular agents over particular actions. For example, a database management agent owned and operated by IBM¹ might have an extremely cooperative relationship with an information gathering agent owned by Lotus (Lotus is a subsidiary of IBM), but an entirely different type of relationship with a Microsoft information gathering agent – the IBM agent might prefer to service requests for the Lotus agent over the Microsoft agent or it might be willing to cooperate with the Microsoft agent if a higher fee is paid for its services. The agents might even coordinate via different protocols; the IBM agent might haggle with the Microsoft agent over delivery time and price whereas it might simply satisfy the Lotus request in short order and with a nominal or zero profit margin. Representing situations such as these is one aspect of our current research agenda. The overall objective is to expand the contextual information used by agents to make control decisions. Space limitations preclude a full description

* Effort sponsored by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory Air Force Materiel Command, USAF, under agreement number F30602-97-1-0249 and by the National Science Foundation under Grant number IIS-9812755 and number IRI-9523419. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA), Air Force Research Laboratory, National Science Foundation, or the U.S. Government.

¹ Disclaimer: The use of actual corporate entities in this example is for illustration purposes only.

of the modeling or knowledge structures under consideration, however, the structures specify, or partially specify factors such as: the (multiple) organizations to which an agent belongs, the different organizational roles an agent is likely to perform for the organization (in a task-centered sense, akin to [19]), the relationships between agents within the organization and without, the importance of a given role to an organization, the importance of a role to the agent, the coordination protocols to use in different circumstances, etc.

Broadening the scope of an agent's understanding of the organizational context in which it operates affects the agent control equation in two primary ways. *Structural* information affects the *scope* of the agent control process. For example, information that specifies with which agents a given agent is likely to interact, with respect to a particular goal, affects the scope of the agent's coordination dialogue. Structural information is particularly important in large multi-agent systems (MAS) because it helps control the combinatorics – it may constrain the distributed search space for any coordination episode. In contrast, *value* information pertains mainly to representing, and reasoning about, complex agent relationships. Value information affects the way in which a given agent evaluates its candidate tasks and actions; information that describes the objective function [37] of an organization, and thus the relative importance of tasks performed for the organization, falls into this category. In this paper, we focus on the *value* side of the problem, i.e., on the agent's *in context* task valuation and selection process.

To ground the discussion, consider a simple example. Figure 1 shows an organized network of financial information agents in the WARREN [15] style. The network is a subset of a larger organization of agents that is populated by three types of agents. *Database Manager (DBM)* agents are experts in data maintenance and organization. These agents maintain repositories of information, e.g., D&B reports, Value Line reports, financial news, etc., and act as the interface between a repository or digital library and the rest of the network. The repositories may be simple databases, collections of databases, or even entail lower-level database management agents with which the primary database manager interacts. Thus the manager's functions are not simply to query a single existing database, instead they conform to the properties of agency, having multiple goals, multiple ways to achieve the goals, and so forth. *Information Gathering (IG)* agents are experts in particular domains. They know about databases (and database managers) pertaining to their area of expertise, or know how to locate such databases. Their task is to plan, gather information, assimilate it, and produce a report, possibly accompanied by a recommendation to the client about a particular action to take based on the gathered information. The *IG* agents pictured in the figure are both experts at collecting and assimilating financial news to build investment profiles of different companies. *Personal Agents (PA)* interface directly with the human client, perhaps modeling the client's needs. These agents also decide with which information specialists to interact to solve a client's information need. *PAs* for a given company may interact with specialists outside of the company, however, interaction styles may differ, i.e., different protocols may be used, different fee structures may apply, etc. The edges in the figure denote interactions between agents. We will focus on the interactions and relationships between the *IG* expert for Merrill Lynch, denoted IG_{ML} , and the *IG* expert for Schwab, and the multiple pictured *PAs*.

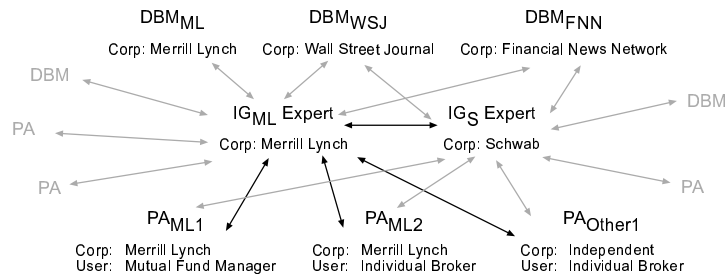



Fig. 1. A Network of Organized Information Agents

Agent IG_{ML} is organizationally situated. The agent belongs to multiple different organizations and it has different relationships with the other agents, stemming from the different organizations, different organizational objectives within and without the organizations [5], and from different relationships within the organizations. Figure 2(a) shows IG_{ML} 's organizational relationships. It is part of the Merrill Lynch corporate structure and thus shares this organization with PA_{ML1} and PA_{ML2} . It is also part of the set of IG agents that specialize in financial information gathering and shares this in common with IG_S . IG_{ML} also belongs to the organization of financial information agents and shares this in common with PA_{Other} . Note, we view organizations as hierarchical structures that can be specialized (i.e., subclassed). In this figure, the organization shared by IG_{ML} and PA_{Other} may have the same root as the organization shared by IG_{ML} and IG_S , however, the specializations differ (in fact, all the agents are members of a root organization pertaining to financial information agents). In addition to its organizational positioning, IG_{ML} also has different relationships within these organizations. Figure 2(b) shows the agent's different relationships. This figure differs from Figure 2(a) in that IG_{ML} has a different relationship with PA_{ML1} and PA_{ML2} . While PA_{ML1} and PA_{ML2} are both members of the Merrill Lynch organization, PA_{ML1} represents a mutual fund manager from the funds division and PA_{ML2} represents an individual broker associated with the retail division.

One of the issues that arises when examining a scenario like this is the need to relate the different motivational factors that influence agent decision making. For example, IG_{ML} interacts with PA_{ML1} and PA_{ML2} for cooperative reasons. In contrast, IG_{ML} interacts with PA_{Other} for self-interested reasons, namely profit for itself, its division,



(a) IG_{ML} 's Organizational Memberships

(b) IG_{ML} 's Inter-agent Relationships

Fig. 2. Different Relationships Complicate Action Choice

or Merrill Lynch. Agents situated in open, social, environments interact with different agents, and different organizations of agents, for different reasons. The ability to relate these different motivations is a requisite for the agents to act rationally, or approximately so, given their social context. Without this ability, how can IG_{ML} determine which requests to service, and in what order? Assuming a model in which agents are rationally bounded, tasks/requests arrive dynamically, and deadlines or service times on requests are present, the agent cannot simply perform all the tasks, but must instead select a subset of the tasks to perform and then determine an appropriate sequence in which to perform them. It is important to note that the agent decision process is contextual. Since the environment is dynamic, and the state of problem solving changes over time, given a set of tasks from which to choose, the choice of which tasks are appropriate is dependent on the current situation. For instance, if IG_{ML} has spent the last n units of time problem solving for PA_{ML1} , and new requests from PA_{ML1} and PA_{ML2} arrive, even if PA_{ML1} requests generally take precedence over PA_{ML2} requests (as specified by the organizational structure), it may be appropriate for IG_{ML} to service the PA_{ML2} request before servicing the PA_{ML1} request.

Figure 3 shows IG_{ML} 's candidate actions at some point time, t . The tasks are structured in a TÆMS [16] network, though the $sum()$ function simply specifies that any number of the tasks may be performed in any order. IG_{ML} 's candidate tasks include servicing requests from PA_{ML1} , PA_{ML2} , and PA_{Other} , as well as doing a local-only task (updating its source models). It also has the option of contracting out its *update sources* task to IG_S . In order to compare these actions the agent requires a framework that quantifies and relates the different motivational reasons for performing particular tasks, as well as relating the costs/benefits of doing tasks for others, and doing local work, to the costs/benefits associated with contracting out the local update task. The complexity of the relationships that the agent has with other agents requires this complex approach to evaluation. The rationale for keeping the different motivational concerns separate is that they represent quantities that are not interchangeable, e.g., progress toward different problem solving objectives, akin to [33]. They are not reducible at all agents to some uniform currency and not all quantities have value to all other agents. For example, doing a favor for someone cannot in turn be used to purchase something at the local store. Another intuitive example: work done on one's yard has no intrinsic value to a professional peer, unless said peer is your neighbor. With respect to computational agents, partitioning of concerns like these maps to the balancing of local work with non-local work, but also to the balancing of work done to satisfy some




Fig. 3. IG_{ML} 's Abstracted Task Structure

joint goal [26, 21, 13, 38] JG_α in contrast to work done to satisfy joint goal JG_β . The idea of this research is not wholly to partition different activities, and the evaluation of their *worth* to the agent, but rather to support ranges of representations, e.g., tasks and actions that have both self-interested and cooperative motivations, or work relating to multiple different joint goals held by multiple agents related, at least partially, through different organizations.

In the sections that follow we present a model for relating different motivational factors, and different measures of progress, that enables agents to compare different types of actions, and the costs and benefits of particular courses of action. We then discuss the issue of interfacing this model with our existing agent control technologies and present ideas about how agents will make decisions based on this model.

2 Quantifying and Comparing Motivations

There are three different classes of tasks that a socially situated agent, such as IG_{ML} , must reason about: 1) tasks that are of local concern only and do not have direct value or repercussions in any non-local context; 2) tasks that other agents wish the local agent to perform; and 3) tasks that other agents may perform for the local agent. Obviously, there are graduations or tasks that pertain to more than one of these classes. For example, a task may produce a result that is valuable locally as well as having value to another agent. Additionally, each task may be performed for cooperative reasons, self-interested reasons, or ranges of these. For example, performing a task for an associate for a nominal fee may pertain to both cooperative concerns and self-interested motivations. It is important to note that even actions performed for cooperative reasons actually have different motivations. For example, doing a favor for one's superior at work is evaluated differently than doing a favor for a peer, which is treated differently than doing a favor for persons unknown, and so forth. In order to address these concerns, we have developed a model for agent activities that quantifies these different motivational factors and enables the local agent to compare the factors via a multi-attributed utility function. Definitions:

Agents are autonomous, heterogenous, persistent, computing entities that have the ability to choose which tasks on which to operate, and when to perform them.² Agents are also rationally bounded, resource bounded, and have limited knowledge of other agents. Agents:

- Can perform tasks locally if they have sufficient resources.
- Interact through communication with other agents to perform tasks. For presentation clarity, we will cast discussion in terms of two basic interaction models: the local agent asking other agents to perform tasks, or the local agent performing tasks for other agents.
- Agents interact via multiple different mediums of exchange known as *motivational quantities* (MQs) that are produced by performing tasks, i.e., a given agent has

² This is by no means the only definition of agency [25, 28, 12, 38, 31, 23, 18, 44].




Fig. 4. Role of MQ in Agent Control

a set of MQs that it accumulates and exchanges with other agents, as shown in Figure 4(a).³

- Not all agents have the same MQ set. However, for two agents to interact, they must have at least one MQ in common (or have the means for forming an MQ dynamically).
- For each MQ_i belonging to an agent, it has a preference function or utility curve⁴, U_{f_i} , that describes its preference for a particular quantity of the MQ , i.e., $\forall MQ_i, \exists U_{f_i}()$ such that $U_{f_i}(MQ_i) \mapsto U_i$ where U_i is the utility associated with MQ_i and is not directly interchangeable with U_j unless $i = j$. Different agents may have different preferences for the same MQ_i .
- An agent’s overall utility at any given moment in time is a function of its different utilities: $U_{agent} = \gamma(U_i, U_j, U_k, ..)$, as shown in Figure 4(b). We make no assumptions about the properties of $\gamma()$, only that it enables agents to determine preference or dominance between two different agent states with respect to MQs .
- For simplicity of presentation, let us assume that $\gamma()$ is not a multi-variate utility function and instead that for each U_i there is an associated function $\omega_i()$ ⁵ that translates MQ specific utility into the agent’s general utility type, i.e., $\forall U_i, \exists \omega_i()$ such that $\omega_i(U_i) \mapsto U_{agent}$. Thus U_{agent} may take the form of Equation 1.⁶

$$U_{agent} = \sum_{i=0}^n \omega_i(U_i) \quad (1)$$

³ If agents are allowed to contract with other agents via a proxy agent, and the proxy agent translates MQs of one type to another, it is possible for the agents to be viewed as sharing a common MQ . However, this is limited by the availability of MQs of the proper type. If we ignore the issue of MQ quantity, the general issue of reducibility of MQs via proxy can be viewed as a graph connectivity problem.

⁴ We currently view these as continuous functions but are exploring the possible need for step-wise utility functions that describe “saving-up” for a potential future event.

⁵ Astute readers will note that $\omega_i()$ could be combined with $U_{f_i}()$. We partition these concerns to provide separate places for mapping different organizational and relationship-centered influences.

⁶ This simple model assumes that all utilities associated with different motivational quantities can be mapped to a common denominator at the agent. This does not mean that the same mapping is possible at all agents, nor do we feel this property is necessary for the model. It is intended to simplify presentation and model manipulation at this time.

- Change in agent utility, denoted ΔU_{agent} , is computed through changes to the individual utilities, U_i, U_j , etc. Let U_i denote the utility associated with MQ_i before the quantity of the MQ changes (e.g., as the result of task performance). Let U_i' denote the utility associated with the changed MQ quantity. The change in overall utility to the agent, in this simplified model, is expressed in Equation 2.

$$\Delta U_{agent} = \left| \sum_{i=0}^n \omega_i(U_i') - \omega_i(U_i) \right| \quad (2)$$

Tasks are abstractions of the primitive actions that the agent may carry out. We return to the issue of abstraction in Section 4. Tasks:

- Require some time or duration to execute, denoted d_i .
- May have deadlines, $deadline_i$, for task performance beyond which performance of said task yields no useful results. (This could also be defined via a function that describes a gradual decrease in utility as $deadline_i$ passes.)
- May have start times, $start_i$, for task performance before which performance of said task yields no useful results. (This could also be defined via a function that describes a gradual increase in utility as $start_i$ approaches.)
- Produce some quantity of one or more MQ s, called an *MQ production set* ($MQPS$), and is denoted by: $MQPS_{i,j,k} = \{q_i, q_j, q_k, \dots\}$, where $\forall i, q_i \geq 0$. These quantities are *positive* and reflect the benefit derived from performing the task. They may be the direct outcome of performing the task, i.e., some result produced by doing the actual work, or they may be quantities that another agent is paying for the work to be performed. In this model, the two are equivalent.
- Tasks may have multiple MQ production sets; that is a given task may produce different groups of MQ s. This models the idea that agents may interact with multiple different mediums-of-exchange. For instance, agent IG_{ML} may service a request for agent IG_S in return for some financial compensation, or by IG_S “calling-in” a favor, or for some combination of these. The multiple MQ production sets are represented: $\{MQPS_{i,j}, MQPS_{l,m}, \dots\} = \{\{q_i, q_j\}, \{q_l, q_m\}, \dots\}$. Note that $MQ_x \cap MQ_y$ may $\neq \phi$ as different $MQPS$ sets may have common members. To simplify presentation, we concentrate on tasks that have a single $MQPS$, though we return to the issue of different $MQPS$ in Section 3.⁷
- Akin to the $MQPS$, tasks may also consume quantities of MQ s. The specification of the MQ s consumed by a task is called an *MQ consumption set* and denoted $MQCS_{i,j,k} = \{q_i, q_j, q_k, \dots\}$, where $\forall i, q_i \geq 0$. As with $MQPS$ s, a task may have multiple $MQCS$ sets. Consumption sets model the idea of tasks consuming resources, tasks hindering progress toward some objective, and agents contracting work out to other agents, e.g., paying another agent to produce some desired result or another agent accumulating favors or good will as the result of task performance. In contrast to production sets, consumption sets are the *negative* side of performing a particular task.

⁷ The issue of which $MQPS$ from the candidate sets will pertain to a given transaction can be viewed as an issue for explicit negotiation between agents [20, 7], or as a dynamic agent choice problem [40].




Fig. 5. Motivational Quantities and Utility

- All quantities, e.g., d_i , $MQPS$, $MQCS$, are viewed from an expected value standpoint. We return to the issue of uncertainty in Section 5.

To illustrate, Figure 5 shows a single utility curve for a single MQ . Assume some task, T , produces X amount of MQ_i . The agent reasons about task performance, and the utility thereof, by comparing the change in U_i associated with the change in MQ_i that performing T will produce. If this is the only task being considered, $\Delta U_{agent} = \Delta U_i$.

Figure 6 illustrates the model's application to the task structure of IG_{ML} pictured in Figure 3. The different problem solving options available to IG_{ML} are: 1) performing task $T_{P_{AML2}}$ for P_{AML2} ; 2) performing task $T_{P_{AOther}}$ for P_{AOther} ; 3) performing its local update task, T_{Local} ; 4) contracting its local update task out to IG_S , represented as T_{IG_S} . Recall that IG_{ML} has different relationships with P_{AML2} , P_{AOther} , and IG_S . As shown in Figure 4(a), the agents' different relationships translate into different MQs with which they interact. IG_{ML} services requests from P_{AML2} for cooperative reasons – it is part of IG_{ML} 's job description and it is recorded as an inter-company transaction for reporting purposes. This motivation is expressed as MQ_{ML2} in IG_{ML} 's $MQPS$. In contrast, IG_{ML} has a very different relationship with IG_S – per the two agents' $MQPS$ sets, they may interact via currency ($MQ_{\$}$) or via an MQ based on professional favors, classified as MQ_S . IG_{ML} has still another relationship with P_{AOther} and they interact via currency only. To compare the different candidate tasks, IG_{ML} reasons about the positive/negative changes in utility that result from carrying out the tasks. For example, to compare T_{Local} , $T_{P_{AML2}}$, and $T_{P_{AOther}}$ (assuming the single valued utility mapping shown in Equation 1):

1. For $T_{P_{AML2}}$: 1) The task consumes a local resource U_{R1} , e.g., monthly allotment of ppp connection time. Therefore, compute the negative change in U_{R1} that will result from the performance of $T_{P_{AML2}}$; 2) compute the positive change in U_{ML2} that is produced by performing the task for P_{AML2} (i.e., the increase in MQ_{ML2}); 3) $U'_{agent\ scenario:ML2} = \omega(U'_{R1}) + \omega(U'_{ML2})$.
2. For $T_{P_{AOther}}$: 1) compute the negative change in U_{R2} , another (different) local resource that is consumed by $T_{P_{AOther}}$; 2) compute the positive change in $U_{\$}$ that is produced by performing the task for P_{AOther} (i.e., the monetary payment from P_{AOther} to IG_{ML}); 3) $U'_{agent\ scenario:PAOther} = \omega(U'_{R2}) + \omega(U'_{\$})$.
3. For T_{Local} : 1) compute the positive change in U_L produced by the performance of task T_{Local} ; 2) $U'_{agent\ scenario:Local} = \omega(U'_L)$.
4. To select between the three, simply choose whichever has the highest gain in utility for the agent. For example, if $U'_{agent\ scenario:Local} \geq U'_{agent\ scenario:PAOther}$ and

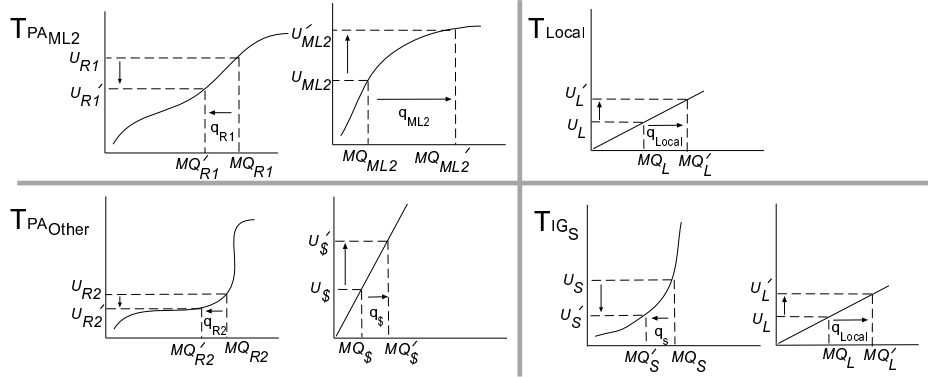


Fig. 6. Comparing Different Candidate Tasks

$U_{agent_{scenario:Local}} \geq U_{agent_{scenario:ML2}}$ then perform the local action. In other words, if the gain in utility achieved by performing T_{Local} exceeds the utility produced by performing $T_{PA_{ML2}}$, even when considering the resource cost of $T_{PA_{ML2}}$ (note that U'_{R1} is less than U_{R1} in Figure 6), then it is preferable to perform T_{Local} . Likewise with $T_{PA_{Other}}$.

If the agent's objective is to simply select which task to perform next, and tasks do not have associated deadlines, and the present and future value of MQ s are equivalent, then it can reason using the maximum expected utility principle and select the task at each point that maximizes immediate utility. However, this simple choose-between-available-tasks model does not map well to situations in which tasks have deadlines, or even situations with a temporal component. For example, consider choosing between T_{Local} and T_{IG_S} : if $\omega(U'_L) \geq \omega(U'_S) + \omega(U'_L)$ then perform the task locally, otherwise, contract it out. In this case, $\omega(U'_S)$, which is the cost of having IG_S perform the task for IG_{ML} , must be zero in order for IG_{ML} to consider allocating the task to IG_S . In order to properly assess the value of such an arrangement, the agents need to use the model presented in this section for comparisons, but, to add components such as *opportunity cost* or *future value* to the selection / decision process. We return to this issue in Section 5.

In this section we have presented a model for comparing tasks that are motivated by different factors. The model can support comparison between tasks that are performed for other agents in return for financial gain to tasks that are performed for other agents for cooperative reasons. Via the different preferences for the different quantities, agent control can be modulated and agents can reason about mixtures of different task types and different motivations. For example, a socially situated agent can reason about doing work in exchange for money as well as progress toward organizational objectives or the accumulation of goodwill, favors, and other non-currency exchanges. The use of state in the model also facilitates contextually dependent behaviors or adjustments to behaviors over time. Agent α performing cooperative work with a closely allied agent, β , for instance, may need to balance this work with cooperative work with others over time. As α accumulates goodwill (represented as one MQ) with β , its preference may shift to the accumulation of other MQ s. The use of utility for this application is flexible and very general, though to effectively use the model we must address how to meaningfully plan

and reason with the model and how to integrate it into existing agent control technology. We return to these issues in Sections 4 and 5.

3 Incorporating Organizational Structure and Influence

The MQ model enables the direct comparison of work motivated by variety of different sources, and it supports ranges of these. The model also supports the integration of certain classes of organizationally derived influence and structure. For instance, organizational relationships can be associated with particular MQ s, i.e., agents belonging to a particular organization and interacting for a particular organizational goal can track their contributions and joint progress (either by communication or by observation) toward the goal using an MQ explicitly for that purpose. Using the same means, agents can reason about their progress toward multiple different goals held by different organizations.

The selection of different $MQPS$ and $MQCS$ is another place where organizational structure integrates with the model. Organizations may have relationships with each other and this can be mapped into the selection of MQ s in particular $MQPS / MQCS$ sets. For instance, if organization α related to organization β in such a way that members of α are willing to coordinate in a cooperative fashion, though to a limited extent, with members of β , agents belonging to α can exchange MQ_α as well as $MQ_{\alpha\beta}$. The notion of “limited extent” in the previous sentence points to another place where organizational structure maps into the MQ -centric model; the preference functions or utility curves of the agent reflect the relative importance of particular types of problem solving activities to the agent. For example, a type of problem solving that is very important to the agent will have a steep utility curve relative to its other concerns; this approach also pertains to power relationships between agents. Organizational influences and relationships can also be mapped to γ , or to the ω functions used in the utility mapping of Equation 1.

Organizational structure imposed on the computation also comes into play in the initial assignment of q_i 's (quantities of MQ s) to agents. Note that since work is produced over time, the system is not a *zero sum game*, but instead is a growing economy. The initial allocation of MQ s to agents predisposes the system to initialize in a particular way and biases the flow of the distributed computation, as in [33]. Agent communication also has roles in this model. Negotiation [20, 30, 7] between agents can be used to select which MQ s, from a set of candidate $MQPS / MQCS$, will be used for a given exchange or produced by a given task execution. Negotiation can also be used to determine the “price” (in MQ s) or quantity that a particular transaction will produce. Auctions or other market mechanisms [43, 9, 6] can be integrated with the model through this avenue.

4 Integration with Detailed Agent Control

The MQ model is deliberately abstract to simplify control reasoning at the *meso*-level of agent control [34]⁸, i.e., the computational organizational level rather than the micro-level. While it could be used directly at the micro-level of agent control, the agent would

⁸ We are currently exploring the relationship between the meso-level and the *social-level* [24].

be unable to reason about a wide class of issues that are important for socially situated, resource bounded, agents. The model lacks features such as explicit representation and quantification of interactions⁹ between tasks and a detailed view of the actions that may be used to carry out the tasks. We generally subscribe to a model where agents have alternative ways to perform tasks (or achieve goals), and that part of the agent control problem is to evaluate the different possible ways to perform a task, taking into consideration the different trade-offs or performance characteristics, and to select one or more from the set of alternatives. Additionally, detailed and complex interactions between agent activities, such as chains of interactions, motivate detailed coordination between agents. This detailed, quantitative, temporal, constraint and interaction based view of the world is embodied by research in TÆMS [17], Design-to-Criteria (DTC) [41, 39] agent scheduling, and GPGP [16] agent coordination.

The existence of such sophisticated, quantitative, machinery for agent control begs the question of why the MQ -centered model is necessary. The detailed technologies are well suited to representation and control at a particular level of detail (micro-level). However, TÆMS is designed to represent a quantified view of the problem solving process of an agent – it does not lend itself to organizational level issues in its current form. Enhancing TÆMS for organizational level application may be possible, though because the class of issues is inherently different at the organizational level, we believe a new structure coupled with a new reasoning process is appropriate. The integration of the organizationally centered MQ framework with the detailed tools is akin to other recent work in integrating process-program controllers [27] and opening the detailed tools for use with BDI problem solvers [8, 35] and others [42, 29]. The general view is that the higher-level components are responsible for influencing the selection of candidate tasks for the agent, while the detailed tools (GPGP/DTC) reason about satisficing, real-time, detailed, temporal control or *implementation* of the selected tasks and goals. Space precludes a detailed discussion, more information is available in [39].

5 Conclusion and Future Directions

The model presented here is currently under development and integration. Recent extensions [39] include the addition of multiple *alternative* performance profiles for MQ tasks and support for an approximate MQ scheduling process. The scheduling process [39] includes facets that factor-in the future value of MQs , temporal issues, and opportunity costs. The potential importance of future value is illustrated in tit-for-tat agent coordination [36] and other cooperative games [32]. Opportunity cost plays a role in task selection, as well as the evaluation of long-term contracts [30] and negotiation [20, 7] over the terms (time and MQs) of said contracts. Reasoning about decommitment penalties or costs [1] also factors into the model at this level.

Regardless of the underlying scheduling technology, the model stands on its own merits as a way to quantify and relate hereto unrelated concerns like cooperative and self-interested motivational factors. Using the model, agents can reason about different

⁹ However, we are considering certain classes of interaction modeling at this level; the issue is expressiveness versus tractability.

concerns like self-interest, favors, altruism and social welfare [14]¹⁰. The model also frames the problem of balancing these different motivations, as well as balancing work between different organizational entities, the individual and the community [24], and balancing different agent relationships. It is important to note, however, that the model requires detailed information about tasks, organizational goals, *MQs*, and the utility functions of each individual agent. Certain classes of this information could be learned though in the general case this falls on the designer(s) of the multi-agent system. Obviously, design principles that guide such a process are desirable.

While the *MQ* model is related to research in social welfare, utility, and choice [22, 14, 4, 7], the model differs in its use of a local, state-based (contextual), view of the larger organizational issues. In the *MQ* framework, agents reason about the utility of particular actions based on their local view of organizational objectives expressed via *MQs* and utility functions. Inherent in the framework are the assumptions that: 1) agents have imperfect knowledge of the problem solving taking place at other agents; 2) the utility function of a given agent cannot generally be shared and computed by other agents because it is dependent on the agent's problem solving state;¹¹ 3) globally appropriate behavior can be approximated through local reasoning in the spirit of [10]. In this latter case, the precision of the approximation is dependent on the degree to which agents can *communicate* or *observe* problem solving toward organizational objectives. Distinctions made, there is a relationship between the model and research in social welfare, utility, and choice. In a sense, *MQs* might be used to approximate and implement social utility functions in multi-agent systems populated by complex problem solvers. It might also be reasonable to combine the technologies online, where formal views of social utility are used to determine *MQ* allocations and *MQ* utility functions, or where social utility is used in the organizational design phase to weight organizational objectives for the *MQ* level. There are also important empirical lessons that can be learned from the large body of research in social utility and social welfare.

The model also relates to other recent work in the multi-agent community, such as agents interacting via obligations [2], or notions of social commitment [11], but it differs in its quantification of different concerns and its dynamic, contextual, relative, evaluation of these. The model resembles MarCon [33] as the different degrees-of-satisfaction afforded by the *MQ* model is related to MarCon's constraint optimization approach, and MarCon too deals with utilities/motivations that cannot always be commingled. MarCon, however, views constraints as agents, assigning particular roles to particular agents, and the issue of which tasks to perform do not enter into the problem space.

¹⁰ All mapped to different *MQs* or groups of *MQs*. However, the issue of how to specify system-wide goal criteria, or organizational-level goals, that characterize acceptable ranges of these must also be addressed to employ *MQs* to concepts like social welfare in a meaningful fashion.

¹¹ To share such a function requires full exchange of the agent's knowledge structures and its objectives and that the receiving agent engage in the same (generally) exponential planning/scheduling computation that the sending agent uses to decide on its course of action (and that the receiver thus does this for every agent with which it interacts). In other words, we take the view that the computation of the utility that a different agent associates with a particular task is not generally feasible in complex real-time resource-bounded problem solving agents (there are also obvious issues of privacy and heterogeneity).

Evaluation of the MQ framework has two facets: modeling and scheduling. Evaluating the scheduling of MQ s is straightforward. Though not generally tractable, the space of possible MQ schedules can be produced exhaustively and the output of any approximate MQ scheduling process can be directly compared to the provably optimal solution. Evaluating the modeling aspects of the MQ framework is more subjective. The questions that must be answered are 1) does the model express desired situations, 2) does reasoning with the model enable the agent to act appropriately given the situation described in the model. If we assume no calculation errors in computing utilities for MQ tasks, case two reduces to case one. The real question is whether or not the model maps well to the situations for which it was designed. Because of the model's somewhat unique integration of local control combined with temporal constraints and utility, it is difficult to compare it directly to other work in social choice. We are currently experimenting with the representational strength of the model. Preliminary results can be found in [39].

Many other research questions remain. Aside from the obvious (and deliberate) lack of prescriptive semantics for the model, one of the outstanding issues is how to best leverage the model from a decision making standpoint, i.e., how to incorporate the model into a high-level decision process that can then be integrated with the rest of our agent control technology as discussed in Section 4. Another obvious question is how to translate organizational goals and objectives into MQ allocations, assignments of $MQCS/MQPS$, and local agent utility curves. Currently, this process is being explored by hand, though an automated organizational design process [3] is a future possibility once the issues are better understood. In terms of limitations, the primary issue is the relative "youth" of the MQ framework. While the local, state-based view appears appropriate for certain classes of agent control, it has yet to be employed in a wide range of projects and situations.

6 Acknowledgments

We would like to thank Professor Lawrence Zacharias of the Isenberg School of Management at the University of Massachusetts for his expertise in social economics and other related bodies of work. We would also like to acknowledge Professor Lee Osterweil for his thoughts on representational validation and Professor Shlomo Zilberstein for his expertise in flexible computation. Their input, combined with the views of the authors, shaped the formation of the MQ model.

References

1. Martin Andersson and Tumas Sandholm. Leveled commitment contracts with myopic and strategic agents. In *Proc. of the 15th National Conf. on AI*, pages 38–44, 1998.
2. Mihai Barbuceanu. Agents that work in harmony by knowing and fulfilling their obligations. In *Proc. of the 15th National Conf. on AI*, pages 89–96, 1998.
3. Brett Benyo and Victor R. Lesser. Evolving Organizational Designs for Multi-Agent Systems. MS thesis, University of Massachusetts, 1999.

4. Guido Boella, Rossana Damiano, and Leonardo Lesmo. Cooperating to the group's utility. In N.R. Jennings and Y. Lespérance, editors, *Intelligent Agents VI — Proc. of the Sixth Intl. Workshop on Agent Theories, Architectures, and Languages (ATAL-99)*, Lecture Notes in AI. Springer-Verlag, Berlin, 2000. In this volume.
5. Jonathan Boswell. *Social and Business Enterprises*. George Allen & Unwin Ltd., London, England, 1976.
6. Craig Boutlier, Yoav Shoham, and Michael P. Wellman. Economic Principles of Multi-Agent Systems. *Artificial Intelligence*, 1-2(1-6), 1997.
7. Sviatoslav Brainov. The role and the impact of preferences on multiagent interaction. In N.R. Jennings and Y. Lespérance, editors, *Intelligent Agents VI — Proc. of the Sixth Intl. Workshop on Agent Theories, Architectures, and Languages (ATAL-99)*, Lecture Notes in AI. Springer-Verlag, Berlin, 2000. In this volume.
8. M.E. Bratman. *Intention Plans and Practical Reason*. Harvard University Press, 1987.
9. Donald E. Campbell. *Resource Allocation Mechanisms*. Cambridge University Press, 1987.
10. Norman Carver, Zarko Cvetanovic, and Victor Lesser. Sophisticated cooperation in FA/C distributed problem solving systems. In *Proc. of the Ninth National Conf. on Artificial Intelligence*, pages 191–198, Anaheim, July 1991.
11. Cristiano Castelfranchi. Commitments: From individual intentions to groups and organizations. In *Proc. of the 1st Intl. Conf. on Multi-Agent Systems (ICMAS95)*, pages 41–48, 1995.
12. Phillip R. Cohen, Adam Cheyer, Michelle Wang, and Soon Cheol Baeg. An open agent architecture. In Michael N. Huhns and Munindar P. Singh, editors, *Readings in Agents*, pages 197–204. Morgan Kaufmann, 1998.
13. P.R. Cohen and H.J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42(3):213–261, 1990.
14. James Coleman. *Individual Interests and Collective Action*. Cambridge University Press, Cambridge, England, 1986.
15. K. Decker, A. Pannu, K. Sycara, and M. Williamson. Designing behaviors for information agents. In *Proc. of the 1st Intl. Conf. on Autonomous Agents*, pages 404–413, 1997.
16. Keith S. Decker. *Environment Centered Analysis and Design of Coordination Mechanisms*. PhD thesis, University of Massachusetts, 1995.
17. Keith S. Decker. Task environment centered simulation. In M. Prietula, K. Carley, and L. Gasser, editors, *Simulating Organizations: Computational Models of Institutions and Groups*. AAAI Press/MIT Press, 1996.
18. Robert Doorenbos, Oren Etzioni, and Daniel Weld. A scalable comparison-shopping agent for the world-wide-web. In *Proc. of the 1st Intl. Conf. on Autonomous Agents*, pages 39–48, 1997.
19. Edmund H. Durfee and Thomas A. Montgomery. Coordination as distributed search in a hierarchical behavior space. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1363–1378, 1991.
20. P. Faratin, C. Sierra, and N. Jennings. Negotiation Decision Functions for Autonomous Agents. *Intl. Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1997.
21. Barbara J. Grosz and Sarit Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86:269–357, 1996.
22. Lisa Hogg and Nick Jennings. Variable sociability in agent-based decision making. In N.R. Jennings and Y. Lespérance, editors, *Intelligent Agents VI — Proc. of the Sixth Intl. Workshop on Agent Theories, Architectures, and Languages (ATAL-99)*, Lecture Notes in AI. Springer-Verlag, Berlin, 2000. In this volume.
23. Michael N. Huhns and Munindar P. Singh. Agents and multiagent systems: Themes, approaches, and challenges. In Michael N. Huhns and Munindar P. Singh, editors, *Readings in Agents*, pages 1–23. Morgan Kaufmann, 1998.

24. N. R. Jennings and J.R. Campos. Towards a social level characterisation of socially responsible agents. *IEE Proc. on Software Engineering*, 144(1):11–25, 1997.
25. Nicholas R. Jennings, Katia Sycara, and Michael Wooldridge. A roadmap of agent research and development. *Autonomous Agents and Multi-Agent Systems*, 1(1):8–38, 1998.
26. N.R. Jennings and E.H. Mamdani. Using joint responsibility to coordinate collaborative problem solving in dynamic environments. In *Proc. of the 10th National Conf. on Artificial Intelligence*, pages 269–275, 1992.
27. David Jensen, Yulin Dong, Barbara Staudt Lerner, Eric K. McCall, Leon J. Osterweil, Stanley M. Sutton Jr., and Alexander Wise. Coordinating agent activities in knowledge discovery processes. In *Proc. of Work Activities Coordination and Collaboration Conf. (WACC) 1999*, 1999. Also as UMASS Tech Report UM-CS-1998-033.
28. Henry Kautz, Bart Selman, Michael Coeh, Steven Ketchpel, and Chris Ramming. An experiment in the design of software agents. In Michael N. Huhns and Munindar P. Singh, editors, *Readings in Agents*, pages 125–130. Morgan Kaufmann, 1998.
29. Victor Lesser, Bryan Horling, Frank Klassner, Anita Raja, Thomas Wagner, and Shelley XQ. Zhang. BIG: A resource-bounded information gathering agent. In *Proc. of the 15th National Conf. on AI (AAAI-98)*, July 1998. See also UMass CS Tech. Reports 99-13, 98-03, 97-34.
30. Victor Lesser, Thomas Wagner, and Abhijit Deshmukh. Negotiating via MQs. Ongoing Research, 1999.
31. Victor R. Lesser. Reflections on the nature of multi-agent coordination and its implications for an agent architecture. *Autonomous Agents and Multi-Agent Systems*, 1(1):89–111, 1998.
32. Martin A. Nowak, Robert M. May, and Karl Sigmund. The Arithmetics of Mutual Help. *Scientific American*, pages 76–81, June 1995.
33. H. Van Dyke Parunak, Allen Ward, and John Sauter. A Systematic Market Approach to Distributed Constraint Problems. In *Proc. of the Third Intl. Conf. on Multi-Agent Systems (ICMAS98)*, 1998.
34. Michael J. Prietula, Kathleen M. Carley, and Les Gasser. A Computational Approach to Organizations and Organizing. In Michael J. Prietula, Kathleen M. Carley, and Les Gasser, editors, *Simulating Organizations: Computational Models of Institutions and Groups*, pages xiv–xix. AAAI Press / MIT Press, 1998.
35. A.S. Rao and M.P. Georgeff. Modelling rational agents within a BDI-architecture. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proc. of the 3rd Intl. Conf. on Principles of Knowledge Representation and Reasoning*, pages 473–484. Morgan Kaufmann, 1991.
36. Sandip Sen and Anish Biswas. Effects of misconception on reciprocative agents. In *Proc. of the 2nd Intl. Conf. on Autonomous Agents (Agents98)*, pages 430–435, 1998.
37. Herbert Simon. Organizational objective functions. Personal Conversation, 1997.
38. Milind Tambe. Agent Architectures for Flexible, Practical Teamwork. In *Proc. of the 14th National Conf. on AI*, pages 22–28, July 1997.
39. Thomas A. Wagner. *Toward Quantified Control for Organizationally Situated Agents*. PhD thesis, University of Massachusetts at Amherst, Amherst, Massachusetts, February 2000.
40. Ludwig von Auer. *Dynamic Preferences, Choice Mechanisms, and Welfare*. Springer-Verlag, Heidelberg, Germany, 1998. Chapter 8.
41. Thomas Wagner, Alan Garvey, and Victor Lesser. Criteria-Directed Heuristic Task Scheduling. *Intl. Journal of Approximate Reasoning, Special Issue on Scheduling*, 19(1-2):91–118, 1998. A version also available as UMASS CS TR-97-59.
42. Thomas Wagner and Victor Lesser. Toward Generalized Organizationally Contexted Agent Control. In *AAAI Workshop on Reasoning in Context*, 1999. Also as UMASS CS TR 99-18.
43. Michael Wellman and Peter Wurman. Market-Aware Agents for a Multiagent World. In *Proc. of the 2nd Intl. Conf. on Autonomous Agents (Agents98)*, pages 115–125, 1998.
44. M.P. Wellman, E.H. Durfee, and W.P. Birmingham. The digital library as community of information agents. *IEEE Expert*, June 1996.