Understanding and Modeling Agent Autonomy in Dynamic, Multi-Agent, Multi-Task Environments

Predrag Tosic[*], **Gul Agha**
Open Systems Laboratory, Department of Computer Science
University of Illinois at Urbana-Champaign, USA
{p-tosic, agha}@cs.uiuc.edu

Abstract

Applications where collection of agents reside in dynamically changing environments and are required to accomplish various tasks in real time pose a number of modeling, design and analysis challenges. These challenges include finding appropriate models for coordination and cooperation among the agents, and for individual agent's adaptability and autonomy. This paper chiefly focuses on understanding and modeling individual agent's autonomy. First we try to identify those characteristics that distinguish **autonomous agents** *from other types of agents. We consider autonomy to be a capability of goal-directed individual decision making in presence of uncertainty, incomplete knowledge and/or noise. An autonomous agent is viewed as a pro-active, goal-driven and generally selfish entity, and its autonomous decision making, together with other capabilities such as adaptability and cooperation with other agents, enables the agent to meaningfully strive to maximize its appropriately defined individual expected payoff. We illustrate the general ideas about agent autonomy with an example of autonomous unmanned aerial vehicles (UAVs) on a multi-task mission viewed as a real-time multi-agent system.*

Keywords: *multi-agent systems, autonomous agents, weak and strong agent autonomy, intelligent agents, distributed online scheduling*

1 Introduction

Autonomous agents are a growing and increasingly exciting research area in many scientific disciplines, from economics to social sciences to "hardcore" computer science to artificial intelligence. Different disciplines have different needs, use different terminologies and may have different notions of what exactly they mean by an *agent*. However, agents in economics and those in (distributed) artificial intelligence, for example, nonetheless tend to have at least some shared properties that are essential for why are the underlying entities considered *agents* in the first place.

This preliminary report has two main goals. One is to identify some such general properties of *autonomous agents*, and to look at agent autonomy in the context of goal- or utility-driven agents embedded in complex, dynamic environments from a high level of abstraction, that is, independently of any particular application domain. The second goal is to define a meta-problem of agents in multi-task environments, where the relationship between agent's autonomy, pro-activeness, and goal-orientedness, and simple characterization of these basic agent properties, can readily be established. To illustrate the applicability of the notions and concepts defined at this meta-level, we then look at a concrete example of what can be viewed as a multi-agent system - viz., a system of autonomous, "intelligent" vehicles that are both cooperating and competing in order to serve a number of dynamically changing tasks.

1.1 What are autonomous agents ?

As an introduction, we first reflect some more on what is it, in most general terms, that distinguishes agents from other, non-agent entities. Indeed, there has been much of debate, what property or set of properties exactly qualifies an entity, such as a computer program, for an *autonomous* or an *intelligent* agent. The good survey articles, such as [1] for intelligent agents or [2] for autonomous agents[1], while trying to clarify and unify the terminology, as well as offer broad agent taxonomies, also illustrate the heterogeneity and lack of agreement on the definition and the required (as opposed to optional) properties even in case of certain perhaps somewhat restricted classes of agents, such as autonomous agents that are required to be computer programs (rather than, say, humans or social insects)[2]. Whether the unification of different terminologies as well as various agent semantics is to be seen any time soon, is anyone's guess.

Some of the most frequently encountered general properties of agents found in the literature include reactiveness, pro-activeness, ability to execute autonomously, goal-orientedness or goal-directedness, a capability of sensing the environment

*Contact author. Phone: 217 - 244 - 4418; fax: 217 - 333 - 9386 (for calls/faxes from within the U.S.). Mailing address: Department of Computer Science, 3270 Digital Computer Lab, 1304 W. Springfield Ave, Urbana, IL 61801 USA

[1]The notion of *intelligent agents* in [1], as well as most definitions of intelligent agents, do subsume an appropriate notion of *autonomy* as a necessary but not sufficient capability that an agent needs to have in order to be considered *intelligent*.

[2]Incidentally or otherwise, both [1] and [2] insist that an intelligent (respectively, autonomous) agent be a computer program.

and being affected by the environment, a capability of affecting the environment, sociability, ability to communicate, persistence, purposefulness, and ability to learn and/or reason about the world. Not all agents have to possess all of the above properties, of course. As we are primarily interested in *autonomous agents*, we would like to make an attempt to identify those properties that can be reasonably considered *necessary* for an entity, such as a computer program or an unmanned vehicle or a social insect, to "qualify" for an autonomous agent. In case of computer programs, being capable of autonomous execution seems to be the most natural requirement. However, a question then arises, is this enough? For instance, a *finite state machine (FSM)* executes autonomously and reactively, but we find it hard to consider *FSMs* an appropriate abstraction of autonomous agents. In case of reactive situated agents that are adequately representable by (coupled) finite state automata, all possible states of the environment are mapped to a *pre-determined and fixed* finite set of *equivalence classes* corresponding to the agent's *"internal states"*; and each of these equivalence classes is then mapped into an action from a *fixed* finite set of the possible agent actions.

We shall argue in §4 that *pro-activeness* is central in distinguishing an autonomous agent from, say, a merely reactive situated agent. We also argue that pro-activeness is typically a result of the agent being, in some form, *goal-driven* or *utility-driven.* We shall outline in the sequel a fairly generic (yet unavoidably not fully general) multi-agent, multi-task framework where the agent's autonomous decision making is a capability used for a well-defined purpose, viz., to maximize agent's utility defined in terms of agent's successfulness in completing those tasks. To illustrate the usefulness and broad applicability of the meta-framework, we then outline how to apply this generic framework to an application domain that has recently arisen in our work - the system of unmanned vehicles on a multi-task mission.

We would like to focus specifically on entities that are capable of some form of individual decision making in non-trivial (and possibly quite complex) environments, where those decisions may impact both the entities themselves and their environment in certain ways that are not necessarily easy (or even possible) to always reliably predict. That is, our subject are *autonomous agents* in as general a sense of this term as possible; in particular, rather than restricting these autonomous agents exclusively to computer programs *alone*, we are more inclusive insofar as what kinds of entities, under appropriate circumstances, could qualify for the label of *an autonomous agent.* In particular, we would like to allow the entities that, in addition to certain computing capabilities based on an appropriate computer program, may also possess various *sensors*, *communication links* and *effectors* for interaction and information exchange with their environments.

The rest of the paper is organized as follows. In the next section, we define a multi-agent, multi-task meta-framework within which we intend to define and study some important concepts related to agent autonomy. As a step towards designing parametric models of autonomous agents and multi-agent systems, we identify some of the critical design parameters. We then focus on some of the main ingredients in modeling agents' autonomous behavior strategies[3]. Subsequently, we outline an application of the meta-framework to a system of UAVs that are set to accomplish a number of independent tasks. Finally, we summarize and briefly indicate some directions for future work.

2 Meta-Level Problem Formulation

We now give a high-level framework for a kind of multi-agent, multi-task problems where agent's capability of autonomous decision-making can be expected to play a central role. This framework is not completely general, but we feel it is general enough to capture many situations where one has a collection of autonomous agents and a collection of independent tasks these agents need to complete.

Let us consider a collection of N *agents* that need to serve a collection of M *tasks.* Each task T_j has a dynamically changing value associated with it, that we denote $V_j(t)$. Each agent is assumed to be a pro-active, goal-driven entity. That is, an agent, A_i, is driven by the desire to increase its own (expected) utility, U_i, by consuming as much of *value* of different tasks as possible. In particular, an agent is not simply embedded into its environment, where it may undertake different actions merely as a *reaction* to the (observed) changes in the environment. Instead, it actively seeks to improve its well-being, that is, to increase its (expected) utility. In order to be able to meaningfully and effectively pursue the increase in (expected) utility, the agent must have some idea of what its goals or tasks are, and some estimated *utility function* associated with completing each of these tasks. This is not the same as saying, that the agent needs to know all of its tasks ahead of time, or all of those tasks' values. However, some *a priori*, i.e., *built-in*, basic knowledge about the agent's own self (i.e., its identity), the agent's goals, and awareness of the available capabilities and resources for accomplishing those goals, has to exist. This awareness can be expected, in general, to evolve with time, as the agent goes along in exploring its environment, learning more about its tasks, resources and other agents. In particular, we know of no natural or artificial agents that, however adaptable and/or autonomous and/or intelligent they may be, do not have some of this initial, basic, *a priori* awareness about self, the ability to distinguish between self and one's surrounding, and some idea of one's goals in the world, and how to go about pursuing those goals.[4]

Back to our meta-model, we feel that assigning a time-dependent value-function $V_j(t)$ to each task T_j is a simple and generic enough way of capturing any relatively simple heterogeneity among the subproblems or tasks that are *mutually independent* and that have to be individually solved or completed by the agents. When an agent discovers a particular task, it gets attracted by its value. If the agent happens to be simultaneously aware of two or more tasks at a given

[3]These considerations naturally open up the issue of what is the proper role of game theory and decision theory in studying *MAS* made of autonomous agents. Due to space limitations, however, we leave that discussion for another occasion.

[4]In case of natural agents, i.e., living organisms, one of the basic goals is certainly survival, and all species that can be reasonably considered autonomous agents have some, however rudimentary, sense or survival instinct, i.e., in our terminology, some *a priori* or built-in idea or intuition or instinct of how to pursue the basic goal of survival.

(discrete) time step[5], the agent needs to decide which of the tasks is currently *most attractive* to it. Based on that decision, the agent then chooses one of finitely many possible actions at its disposal; this action may include sending messages to other agents. Whatever action the agent chooses, the idea is that, the agent acting rationally, it is agent's desire and hope that this action will maximize agent's (expected) payoff. In our simple meta-problem, the choice of action basically boils down to which task the agent chooses to tackle next. In particular, we observe that this model implies the agent does no (long-term) *planning*. Indeed, the *classical planning* framework familiar from AI does not apply here, as the assumption is that the tasks and environment in general are only partially observable, dynamic and possibly nondeterministic (or at least appear so to the agent); see, e.g., §11 in [3]. More complex models of planning, and, in particular, those that include *decision making under uncertainty*, however, could well be applicable - but we would expect them to be prohibitively expensive for nontrivial problem sizes. We return to decision making under uncertainty in the context of models of agent autonomy in §§4.2.

The total amount of the available value in the system in our meta-problem is assumed to be bounded at all times. Consequently, the agents can be expected to end up competing for this limited resource. Agents are assumed to be *selfish* by default, that is, without an explicit or implicit *incentive*, they are not going to be inclined to cooperate. Each agent's *sole goal* is to maximize its own utility function. However, various forms of cooperation, coordination and collaboration may nonetheless arise. An explicit incentive is, for example, an agent's realization that, with its own resources alone, it cannot accomplish a particular badly desired task, and therefore, such agent may be able to use whatever communication channels it has on its disposal, to contact other agents and offer collaboration. This phenomenon is well-known in multiplayer game theory, where the incentive for forming coalitions is still purely egotistic - including the realization by the rational agent (player) that the spoils, if any, will have to be split with the participating collaborators, for some spoils are still better than none. A generic example of implicit incentive to cooperate with competitors is any kind of a conflict resolution scenario where the agents, unless they coordinate, can expect mutual destruction and, therefore, no hope of fulfilling their goals or maximizing utility. Concrete such examples are readily available, for instance, in the context of intelligent transportation systems and autonomous vehicles.

3 Some Critical Design Parameters in Multi-Agent Systems

We now briefly discuss some of the most critical design parameters featuring in our generic model of a multi-agent, multitask system. One ultimate goal of our current and future investigations is to cover as much of this multidimensional parameter space as possible, analytically whenever feasible and via computer simulation and experiments when analytical approach is not promising, and subsequently to compare and contrast the individual agent behaviors as well as the system performance in various scenarios, thereby increasing our understanding of autonomous agents and *MAS* made of such agents, and, in particular, our prediction ability of the *MAS* behavior and performance under various conditions.

In the meta-model we outlined above, while we do consider our formulation fairly general and high-level, nonetheless a number of simplifying assumptions has been tacitly made. For instance, in much of the work on *MAS* (including our own [4], [5]), the existence of global time and the global clock is assumed. The agents are assumed to communicate with one another (or, when applicable, with any kind of *central control*) exclusively via *message passing*. We also assume that all communication is perfectly *synchronous*. It is well-known that the more realistic assumption of asynchronous communication renders many important distributed coordination and agreement problems formally undecidable [6].

In our current framework ([4], [5]), the agents are assumed all to be equally capable - in particular, no agent specialization has been considered, where only certain agents can perform some particular task or set of tasks. The tasks appear identical to all agents - except possibly for their (true or estimated) values. We further assume that the tasks are *mutually independent* of one another; in particular, in case that an agent has several available choices, which tasks it is going to pick to service, and in what order, is entirely driven by the agent's appropriate estimate of those tasks' values and their expected utilities or payoffs for that agent.

With these simple yet necessary stipulations in place, we can proceed to identifying the main *MAS* parameters. The parameters of interest include the ratio of the number of agents to the number of tasks, agents' sensor and communication ranges, the model of agent-to-agent communication, the model of the agent's knowledge of the tasks (i.e., whether any uncertainty and/or noise are allowed), a choice of a coordination model and strategy (if any), a model of the agent individual behavior strategy (or, in heterogeneous scenarios where agents are distinguishable, *strategies*), a model of individual agent's *adaptability*, etc. For instance, by *adaptability* we mean the ability to change the individual strategies and coordination models based on observed changes in the environment, including but not limited to any form of a *feedback*, such as an appropriate reinforcement via varying, action-dependent *payoffs*, received from the environment. Similarly, coordination may refer to any form of communication or other way of information exchange among different agents, with agents purposefully mutually adjusting their behaviors in order to individually benefit from these adjustments. In other words, in our model of *selfish, individual utility-driven autonomous agents,* each agent may choose to coordinate only because it feels it would be beneficial for its own expected utility to coordinate. Hence, coordination may be simply about conflict avoidance and resolution, or about coalition forming among different groups of agents, or a form of self-organization among the agents, who may adaptively *learn* that, under certain circumstances, they are each individually better off coordinating and cooperating than acting strictly on their own. These general observations are usually more appreciated if some concrete examples from relevant application areas are offered. For that purpose, we refer the

[5] Agent's awareness of tasks' existence, and some, not necessarily accurate, idea of those tasks' values are results of either the agent having sensed those tasks, or because it got the (not necessarily reliable) information about the tasks from other agents.

reader to **§5**, and a brief discussion on agent autonomy, and its relationship to coordination among *UAVs* as an illustrative example. More details on the *UAV* coordination *per se* can be found in [4] and [5].

4 Models of Agent Autonomous Behavior Strategies

We now try to justify the class of generic models of agent autonomy in the context of the meta-problem of an autonomous agent acting in a multi-agent, multi-task environment as outlined in earlier sections. The meta-model of agent autonomy will be an appropriate (class of) mathematical function(s), that, depending on the situation, can be deterministic, nondeterministic or probabilistic[6].

Before we begin a discussion of what properties this class of functions (that we'd like to argue fairly generally captures the notion of agent autonomy) should have, however, we first reflect on what types of agents lack a kind of autonomy we are after.

During the late 1980s and much of 1990s, perhaps the most dominant thread of agent-related research in AI was the study of *situated agents.* These are agents that are embedded in dynamic, possibly quite complex environments, a kind of agents that are capable of sensing the environment, reacting to the observed environmental changes, and also affecting their environment via their own actions. In [7], such persistent situated agents are defined as *"Intelligent agents [that] are systems that have a complex, ongoing interaction with an environment that is dynamic and imperfectly predictable".* While most ingredients of what we consider autonomous agents are indeed present in such situated agents - reactiveness, situatedness, ability to change the environment, persistence - clearly something is missing. It is not surprising that [7] concludes that *"it is more appropriate to think of an agent embedded in an environment as performing a* **transduction"**, and to proceed with a formal, transduction model of such situated "intelligent" agents based on *coupled finite state automata.* What appears confusing, however, is what exactly makes such a transducer *autonomous*, let alone *intelligent*? It is our view that the coupled automata model in [7] is an appropriate abstraction for reactive, persistent agents acting in dynamic, complex and partially or not fully accurately observable environments, but this abstraction, in our view, fails to fully capture the notion of agent *autonomy.*

Let us now consider the notion of *intelligent agents* as defined in [1]. The four necessary properties that a computer program or system has to possess in order to be considered an intelligent agent (in the *weak sense* of [1]) is that it must be *autonomous, responsive, pro-active and sociable.* The authors find responsiveness, pro-activeness and sociability to be characteristics of agent's *flexibility* as a necessary prerequisite for agent intelligence, but, interestingly enough, they do not require *adaptability.* While we do agree with [1] that *intelligent agents* are a (proper) subset of *autonomous agents* - in order for any reasonable, however weak, notion of *intelligence* to hold for an agent, that agent certainly has to have some degree of *autonomy* in its behavior - we also feel that, together with autonomy, responsiveness and pro-activeness, an *intelligent agent* has to be capable of *adjusting* to the changes in its environment and/or in its goals, that is, such agent has to be *adaptable.* This adaptability, however, may but need not always mean a capability of a highly complex behavior such as, e.g., some form of *symbolic learning.* Indeed, in our view an agent may be called *adaptable* even if its capability of autonomous, dynamic behavior adjustment is only rather rudimentary. For example, the ability of an appropriate parameter adjustment based on the feedback from the environment or the estimated distance from fulfilling one's goal would qualify an autonomous agent to be considered *adaptable*, and hence one step closer to *intelligent.* Due to space constraints, we leave modeling agent adaptability for another occasion.

What are, then, the critical ingredients needed to capture agent autonomy? In our view, *reactivity* (or what in [1] is referred to as *"responsiveness"*) is clearly necessary, as the agent has to be able to (i) notice changes in the environment, and (ii) appropriately respond to those changes - but *reactivity alone* does not appear to be sufficient. An existing, *weak notion* of agent autonomy can be approximated as

$$\textit{(weak) autonomy} \approx \textit{reactivity} + \textit{persistence} \tag{1}$$

For example, a "proper" software agent can be expected to be at least somewhat persistent, unlike say a *subroutine* of a computer program whose "turning on and off" is controlled from outside of that subroutine. This view is strongly related to the usual requirement of *(weak) autonomy*, namely, the requirement that an entity (agent) have the control of its internal state.

We identify below two other key ingredients that we find crucial for a notion of autonomy that is stronger than mere "$reactivity + persistence$". It is the presence of these two new ingredients, in some form, that we consider necessary requirements, in addition to reactivity and persistence, for a natural or artificial agent to be autonomous in a stronger, or *"more AIish"*, sense.

One feature that we see in virtually all known genuinely autonomous agents, biological and computational alike, is some form of *goal-orientedness* or *goal-drivenness.* In case of living organisms, that is, "biological (autonomous) agents", the highest level driving mechanisms are the instincts or desires of *survival* and *reproduction*[7]. At lower levels, the driving mechanisms - finding food or a sexual partner - are those that are expected to provide, promote or enhance the two highest-level goals, survival and reproduction. In case of computational agents such as a web crawler or a robot or an autonomous unmanned vehicle, these agents are designed and programmed with a particular *task* (or *a set of tasks*) in mind. For instance, while there certainly are simple robots (e.g., those working in a car factory assembly line) that are merely reactive, such robots lack autonomy, and vice versa: those robots that can be reasonably viewed as autonomous have nontrivial tasks to accomplish, a kind of tasks where mere pre-determined and fixed reactions to changes in the input stream from the outside world is not enough. That is, autonomous robots (and other computational agents that are

[6]That is, a *probability distribution* over a well-defined, *finite* set of possible actions, plans or strategies.

[7]The latter instinct being related to the survival of the species or perhaps, in the "Dawkinsian" terms, of the particular genes and gene patterns [8], as opposed to the survival proper, that is, the survival of individuals.

truly autonomous) are always driven by some built-in notion of a goal or a set of goals that need to be accomplished.[8]

Finally, in addition to responsiveness, persistence and goal-drivenness or goal-orientedness, one more characteristic found in virtually all autonomous agents, not altogether unrelated to goal-orientedness, is that of *pro-activeness.* Biological examples abound: a very hungry lion, even if say physically rather tired, will pursue finding some food, or at least get younger lions to get some food for him[9]. Such a lion, if it were merely reactive, would just lie down and wait for an antelope to come near enough so that he can grab her - something antelopes usually don't do. Now lions sometimes are lazy - but usually not to the point of starving to death. This is so because the lions are *goal-driven* (survival, reproduction) and, in order to be able to fulfill their goals, being just passively reactive is quite often simply not good enough; so, they have to be pro-active in pursuing their goals, i.e., doing what it takes in order to be able to survive and reproduce. In case of computational agents, on the other hand, identifying where pro-activeness comes from may be more complicated but, as a general rule of thumb, it is always closely related to the agent having some *a priori* notion of its goals, and being driven by the desire to achieve those goals. In other words, we argue that *pro-activeness* is related to, and (at least in part) stems from, some form of *goal-orientedness.*

To summarize, we find that it is precisely the properties of (i) reactiveness or responsiveness, (ii) persistence, (iii) pro-activeness, and (iv) goal-directedness or goal-orientedness that, together, make an agent *truly autonomous*:

$$\begin{aligned}\textit{(strong) autonomy} \;\approx\; & \textit{reactivity} + \textit{persistence} + \\ & \textit{goal-orientedness} + \textit{pro-activeness} \qquad (2)\end{aligned}$$

Granted, much of the agent literature has identified properties (i) - (iv) as common to autonomous and/or intelligent agents (see [1], [2] and references therein). We claim, however, that these four properties are *the necessary* properties that are all found in nearly every reasonable model of autonomy, whereas other characteristics, including sociability, adaptability, mobility, "mental states" (beyond whatever is necessary to have some notion of one's goals), beliefs-desires-intentions, etc., are non-essential, and are found in (or can reasonably be attributed to) only *some*, but by no means *all* or *nearly all*, of the known autonomous agents, whether biological or computational.

After this brief excursion into the ontology of agent autonomy, we return to more practical issues: what kind of mathematical and computational tools and paradigms are needed to adequately capture agent autonomy at an abstract, and therefore hopefully broadly applicable level? As we are arguing that the agent autonomy is, ultimately, a capability of (pro-active, goal-driven) decision making, based on one's internal state and the input stream from the outside world, of how to (rationally and effectively) pursue one's agenda or goals, the critical question arises, how to model this decision-making process, and what are the critical parameters that it depends on? We turn to these questions next.

4.1 Agent Autonomy: Noiseless Case

In the quest for generic models of autonomous agent decision making, we first consider the simplest, noiseless case. The agents are assumed to have *perfectly reliable* sensors and communication links. An agent's knowledge of the environment, and of the tasks in particular, while assumed (locally) accurate, is still not necessarily complete, however. Due to the ontological assumptions of (i) no central control, and (ii) bounded sensor and communication ranges, each agent necessarily has only a *local* picture of the tasks, as well as the other agents and their whereabouts. If the agents work in unison, i.e., if they have the common goal that they are striving to achieve (that is, a single *joint utility function* that they all as a system are trying to maximize), then the problem of how to split up the tasks among the agents, and in what order, approximately reduces to a well-known problem of *(distributed) online task allocation and scheduling* (see, e.g., [9]). One substantial difference between classical scheduling problems and this multi-agent division-of-labor problem is that, in case of the latter, instead of some concrete metrics such as the task priority or the expected duration of task completion, a more general notion of the *task value* is the quantitative metric the agents are trying to optimize. In [9] one can find a useful taxonomy of various flavors of (real-time) scheduling problems, some well-known solutions, and a rich bibliography on the subject. Our problem of the choice of an agent's optimal action or strategy in the multi-agent, multi-task meta-framework defined herein would fit into the category of distributed dynamic pre-emptive scheduling problems.

However, the problem we desire to model is conceptually and analytically considerably more complex than mere dynamic scheduling. The main reasons behind the complexity of the problem at hand are the inherently distributed nature of individual agent information, knowledge, and interests - in particular, the assumption of agents striving to maximize any sort of *joint utility* is usually unrealistic, and one ought to assume agents' quests for maximizing (expected) *individual utilities* instead. That is, even under the usually unrealistic assumptions of perfectly reliable communication and sensing, we identify the following generic sources of considerable additional complexity:

(i) each agent only has a local knowledge of the tasks, and of other agents;

(ii) typically, each agent is trying to maximize its own individual payoff, and there is no guarantee, that individual selfishness would necessarily lead to a satisfactory efficiency of the system as a whole;

(iii) even if communication links are perfectly reliable, the information that an agent receives from other agents need not be reliable, as the agents, in general, can be expected to compete for tasks and therefore the *veracity* assumption, in any such competitive scenario, need not hold.

Observation (i), and models and analysis of this paradigm in various application domains, is a subject of *distributed artificial intelligence.* Similarly, (ii) and its generalization - how to reconcile the quest for maximizing individual vs. joint utility functions - is a subject matter of *incentive engineer-*

[8] We remark that payoff- or utility-driven agents, as we see it, are nothing but a little more complex version of goal- or task-driven agents, a kind of agents that, in addition to some notion of their goal(s), also have some "performance metric", i.e., quantitative measure of those goals' values, importance, resource requirements, etc.; see also §2 of [3].

[9] How exactly this kind of coordination takes place is another interesting subject, but beyond the scope of this work.

ing [10]. Finally, (iii) takes us into the realm of *multiple-player game theory*. Therefore, we share the view that the ideas, paradigms and tools of *distributed AI*, *incentive engineering* and *N-person game theory* are all highly applicable and, in case of multi-agent systems of considerable complexity along all the indicated dimensions, even *necessary* for accurate and systematic model development and analysis of such systems. Moreover, this list will be further expanded in the next subsection, when we introduce the additional complexity stemming from faulty communication links and less-than-perfectly-accurate sensors.

4.2 Agent Autonomy: Noisy Case

In the previous subsection, we have assumed that the agents' knowledge of their environment and their tasks, while local, is perfectly accurate. Most of the time, a realistic agent model has to drop the assumption of perfect accuracy in favor of an imperfect, noisy model of agent knowledge of the world. In particular, an agent's sensors, in general, cannot be assumed to be perfectly accurate and reliable; whatever properties of the environment they measure, these measurements likely bring about some noise in the agent's local picture of the world. Likewise, an agent's communication links, in practice, hardly ever can be safely considered perfectly reliable: they may be faulty and they may experience delays. Hence, even if the agent veracity assumption did hold, the delays in communication can lead to an agent basing its decisions on outdated information received from other agents, which, in complex, dynamic environments, often can be worse than not receiving any information from other agents at all - however well-intentional these other agents may be.

Allowing an agent's sensors and communication links not to be perfectly reliable brings in additional complexity to studying autonomous agents and multi-agent systems. In particular, in the single-agent framework, once agent's (local) knowledge of his surrounding becomes noisy, one moves into the realm of *decision making under uncertainty* [11] - a subject of *Decision Theory*. It is worthwhile observing that the environment in which the agent is embedded and acting may be perfectly deterministic and predictable; however, as long as it does appear nondeterministic to the agent who is probing that environment with his limited and imperfect resources, from the agent's perspective, this is all that matters. In case of multiple agents, the stated assumptions cast the problem into the category of multi-stage (or *repeated*) *multi-player games of imperfect information*, but the new element, "borrowed" from distributed computing, is that of possible *asynchrony* in the communication between the agents due to network delays.[10]

It is, therefore, an unavoidable conclusion that multi-agent systems where self-interested, goal-driven or utility-driven autonomous agents act and communicate with one another in dynamic, unpredictable, only partially and imperfectly accurately observable environments, agents whose sensors and communication links need not be reliable, and who in general cannot trust each other's veracity, represent an incredibly complex class of dynamical and computational systems to model and analyze, and pose a number of considerable mathematical and computational challenges.

5 An Example: UAVs on a Multi-Task Mission

A collection of *Unmanned Aerial Vehicles (UAVs)* on a mission provides an ideal framework for identifying, modeling and analyzing many interesting paradigms, design parameters and solution strategies applicable not only to autonomous unmanned vehicles, but to *Multi-Agent Systems (MAS)* in general. UAVs are finding their use in a variety of military and law-enforcement operations, e.g., in various surveillance, reconnaissance, and search-and-rescue tasks. These UAVs carry sophisticated payloads in order to be able to fulfill their increasingly complex missions. In particular, a typical UAV is equipped with certain *sensors* such as, e.g., radars. With these sensors, a UAV probes its environment and forms a (local) "picture of the world" on which its future actions may need to be based. A UAV is also equipped with some *communication capabilities*, that enable it to communicate with other UAVs and/or the ground or satellite control. This communication enables a UAV to have an access to the information that is not local to it - that is, the information not directly accessible to the UAV's sensors.

While trying to accomplish their mission, these UAVs need to respect a heterogeneous set of constraints on their physical and communication resources. The UAVs also need to be able to communicate and cooperate with each other. Their cooperation can range from merely assuring that they stay out of each other's way (collision avoidance) to enabling themselves to adaptively and dynamically divide-and-conquer their tasks.[11]

Not all kinds of UAVs can be reasonably considered genuine agents; e.g., those that are remotely controlled throughout their mission would not qualify for *autonomous agents* in the usual sense. However, for the reasons of system scalability, dependability and robustness, increasingly complex and autonomous unmanned vehicles are being studied, designed, manufactured and employed. We are interested in UAVs that are not remotely controlled and that have the ability to make their own decisions in real time. We are also assuming either no central control, or only a limited central control. In particular, the knowledge of the world that each UAV possesses is, in general, assumed to be *local*, possibly *noisy*, to vary with time, and to be augmentable, at a certain cost, via communication with other UAVs.

Some of the problems that have been extensively studied in the context of UAVs include motion planning and conflict detection and resolution (see, e.g., [12-14]). What has drawn considerably less attention, to the best of our knowledge, is modeling and analysis of the task-driven behavior of the UAVs that can be reasonably viewed as autonomous agents.

We now turn to the *MAS* formulation of the system of autonomous UAVs problem and how it fits into the frame-

[10] In this context, knowing whether the messages that an agent receives are bugged due to other agents trying to mislead this agent, or due to noise in the communication links, may or may not matter to the agent, but dwelling any further into this is beyond the scope of this work. The new element not present in classical game theory, however, is that of communication delays and the asynchrony resulting from such delays.

[11] In [4], this latter, higher form of cooperation (coordination) we also call *goal-driven* cooperation (respectively, coordination).

work outlined in the previous sections. A collection of N UAVs needs to accomplish a certain complex, multi-task mission. We model this mission with a collection of M *interest points (IPs).* An interest point is a semantic extension of the more common notion of a *target* - in addition to targets proper, an IP may also refer to, e.g., a small local region of interest, that may or may not include "real targets", but is nonetheless worth while exploring. Each interest point j has a dynamically changing value associated with it, $\Pi_j(t)$. An IP may be static or mobile. A mobile IP j, at any time step t, is completely and uniquely specified by its position and velocity vectors, $\psi_j(t)$ and $\xi_j(t)$, respectively, and its value $\Pi_j(t)$. Each UAV i is driven by the desire to increase its own utility, U_i, by consuming as much of value of various IPs as possible. The total amount of value is assumed to be bounded at all times. Consequently, the UAVs may end up competing for this limited resource.

From an individual UAV's perspective, the goal is to maximize its own utility, by visiting as many interest points and consuming as much of their value as possible. This is accomplished by following a certain either fixed or dynamically changing (adaptable) *individual behavior strategy.* This strategy can be specified by an appropriate *individual behavior function*, Θ_i, that UAV i follows as long as there is no outside signal telling the UAV it should start doing something else. An example of such outside signal is a request to a given UAV to join a newly formed group; if such request comes from a leader whose supremacy in authority is recognized, the follower UAV will have to abandon its current behavior and comply with the leader's desires, thereby, in a sense, giving up its individual autonomy. Thus, we can observe in this case an instance of a fundamental *tradeoff* between individual autonomy (viewed as a *capability* that enables an agent to strive for maximizing its individual utility), and group coordination which may require (partial or complete, temporary or permanent) sacrifice of the agent's autonomy.

5.1 Simple Models of UAV Autonomy

In order for any type of *"intelligent" vehicles* to be considered *autonomous agents*, they have to be capable of *autonomous decision making* without direct assistance of a human or other outside operator. We outline a simple model of autonomy applicable to UAVs that would render UAVs proper autonomous agents, albeit of a perhaps fairly restricted kind. That is, UAVs are modeled herein as *utility-driven entities.* They fulfill their goals and thus increase their utilities by servicing tasks. In our modeling framework, tasks are given as *interest points* and the UAVs, loosely speaking, strive to consume as much of interest points' *value* as fast as possible. As we assume that a single UAV can consume value from at most one IP at any single time step, the question arises: among several candidate IPs, how should a UAV choose in what order it is to visit these IPs? Therefore, it can be argued that each UAV faces an *online scheduling problem.* For simplicity, we consider a simplified version of dynamic online scheduling, and only ask, given a set of interest points whose current positions and (estimated) values are known to a particular UAV, which IP among those points should the UAV select to visit first?

We model the individual UAV's autonomous decision-making with UAV's *individual behavior functions*, Θ_i. Given a set of IPs with their current positions and values[12], a UAV i evaluates its behavior function Θ_i that returns the index $j^\star$ of the IP such that, if the UAV selects that IP as its next task to service, this choice, from that UAV's perspective, is expected to maximize the *estimated increase* in the UAV's *utility.* A generalization of this short-term, "single-shot" action selection mechanism given by Θ_i to the individual behavior functions that, instead of a single index $j^\star$, would return partial or complete *schedules* $(j_1^\star, ..., j_k^\star)$, is immediate.

In particular, each UAV is assumed to behave *greedily* - unless and until ordered differently from the outside. However, a great variety of greedy strategies can be specified via different choices of the functions Θ_i.

Some variables that individual behavior functions can be expected to depend on are the UAV's distance from the given IP, the IP's current value (or its estimate), and the estimated competition for that IP and its value - viz., the number of other UAVs in the IP's vicinity. Let $x_i(t)$ and $v_i(t)$ be the i-th UAV's position and velocity vectors at time t, respectively, let $\psi_j(t)$ be the position of IP j at time t, $\xi_j(t)$ its velocity, and let $n_{j,r}$ be the total number of UAVs within the distance r from IP j. Then one class of models of the i-th UAV's target selection strategy can be specified by

$$\Theta_i(t) = arg\{max_j\ G(\Pi_j, ||x_i - \psi_j||, ||v_i(t) - \xi_j(t)||, n_{j,r}, t)\}$$

where G is a function that is increasing in Π_j and non-increasing in the distance of the UAV from the IP j, $||x_i(t) - \psi_j(t)||$, and the (estimated) relative velocity between the two, $||v_i(t) - \xi_j(t)||$. This function specifies what IP $j^\star$ should UAV i pick as the estimated short-term (expected) optimal choice.

One example of a simple greedy individual behavior that fits the given general framework is given by

$$\Theta_i(t) = arg\{max_{1 \leq j \leq M}\ [\tfrac{\Pi_j(t)}{n_{j,r}(t) \cdot ||x_i(t) - \psi_j(t)||}]\ \} \qquad (3)$$

where it is assumed that the minimal allowable distance of any UAV from any IP is strictly positive, and that distances are appropriately normalized so that one may pretend they are "dimensionless".

Another example of a greedy individual behavior that either assumes the IPs are stationary or else that their velocities can be neglected is given by

$$\Theta_i(t) = arg\{max_{1 \leq j \leq M}\ [\tfrac{\Pi_j(t) - n_{j,r}(t) \cdot d}{||x_i(t) - \psi_j(t)||}]\ \} \qquad (4)$$

where d stands for *the (constant) consumption rate* of an IP's value, and similar assumptions hold as in (1).

Yet another, slightly more elaborate UAV choice-of-action function that we have considered is given by

$$\Theta_i(t) = arg\{max_{1 \leq j \leq M}\ [\tfrac{\Pi_j(t) - t^j_{est} \cdot n_{j,r} \cdot d}{n_{j,r} + 1}]\ \} \qquad (5)$$

where t^j_{est} is the *estimated time* that UAV i will need to reach close enough to IP j in order to start consuming its value. The heuristic assumptions made - for example, that the number of other UAVs in vicinity of j is going to remain fixed in the meantime, or that all those UAVs are precisely after the IP j (and not some other IP that is in its vicinity) - are a kind of oversimplifications that may not hold in

[12] In case of *mobile targets*, current velocities would be also needed.

a given situation even as crude approximations. Yet, some such simplifying assumptions are likely necessary, in order for the agent's choice-of-action function Θ_i to be readily and quickly computable "on the fly", at each discrete time step t.

Everything said thus far about UAVs' individual behavior strategies rests on the assumption that each UAV acts strictly selfishly, and largely independently (except for the dependence of Θ_i on $n_{j,r}$) from what other UAVs do. Once UAV-to-UAV communication and coordination are taken into account, modeling a UAV's autonomous behavior becomes more complex. In particular, in addition to the already mentioned parameters, each $\Theta_i(t)$ would be expected to also depend on *the set of messages that the* i*-th UAV has received from other UAVs* at time steps $t' \leq t$. These issues, crucial for the design and analysis of UAV communication and coordination models, however, are beyond the scope of this work.

An important point about the IP value function $\Pi_j(t)$ is that this function, for each IP j, represents the j-th task's true (or objective) value, irrespective of which UAV may have observed IP j, and from how far away. In other words, the tacit assumption is that, once an IP is discovered by the UAVs, all UAVs who are aware of this IP's existence immediately also know its exact (current) value. While this *perfect knowledge* assumption is a good starting point for modeling the UAV task-driven behavior, the more realistic and more general assumption is that each UAV has its own, *local* and, in general, imperfect (i.e., *noisy*) and probabilistic knowledge of each of IPs' values. The generalization of this observation to other types of agents and their tasks, in lieu with the discussion in §§4.1. and §§4.2., is immediate.

6 Concluding Remarks

The subject of this paper are *autonomous agents*. First, we offer some thoughts on agents in general, and autonomous agents in particular. We try to identify the critical and necessary characteristics that one should expect to find in any autonomous agent. We also briefly discuss the relationship and importance of these properties, and we try to make a clear distinction between them on one, and various "optional" properties that one may or may not find in different kinds of the actual autonomous agents, on the other hand.

In order to be able to model agent autonomy as an appropriate mathematical function (deterministic or otherwise), we define a meta-problem for autonomous agents' decision-making in multi-task environments. This meta-problem, while very high level, is unavoidably not fully general; in particular, in terms of the common AI classifications, our meta-problem is "scheduling-flavored" rather than "planning-flavored". Nonetheless, we find this meta-problem to provide a sufficient abstract framework for a variety of autonomous agent models where planning can be considered impractical, impossible or prohibitively computationally expensive, and where agents have to make their decisions *online* and/or in real time, and possibly in presence of various types of uncertainties and/or noise. To illustrate the practical usefulness of our meta-problem framework, we dedicate a section to an application of this meta-problem to a concrete example of a *MAS*, and to modeling and analysis problems of appropriately identifying parameters and functionally representing agent autonomy in the context of that application.

While the ideas presented herein clearly represent work in progress, we do hope that this report will shed some light and provide yet another perspective on agent autonomy, and a modest contribution to understanding autonomous agents. We also hope that these ideas, once further developed, will provide some useful insights into how to proceed in pursuing effective ways of mathematically and computationally modeling agent autonomy at a high level, that would make these models general enough to be applicable and useful in a great variety of practical multi-agent system situations.

Acknowledgment: Many thanks to all members of the Open Systems Laboratory's TASK team, for many useful discussions. This work was supported by the *DARPA IPTO TASK Program* under the contract *F30602-00-2-0586*.

Bibliography

[1] M. Wooldridge, N. Jennings, "Intelligent Agents: Theory and Practice", Knowledge Engin. Rev., 1995

[2] S. Franklin, A. Graesser, "Is it an Agent, or just a Program? A Taxonomy for Autonomous Agents", Proc. 3rd Int'l Workshop on Agent Theories, Architectures & Languages, Springer-Verlag, 1996

[3] S. Russell, P. Norvig, "Aritifical Intelligence: A Modern Approach", 2nd ed., Prentice Hall Series in AI, 2003

[4] P. Tosic, M. Jang, S. Reddy, J. Chia, L. Chen, G. Agha, "Modeling a System of UAVs on a Mission", Proc. SCI 2003 (invited session), Orlando, Florida, July 2003

[5] M. Jang, S. Reddy, P. Tosic, L. Chen, G. Agha, "An Actor-based Simulation for Studying UAV Coordination", Proc. 15th Euro. Symp. Simul. (ESS 2003), Delft, The Netherlands, October 2003

[6] G. Tel, "Introduction To Distributed Algorithms", 2nd ed., Cambridge Univ. Press, 2000

[7] S. J. Rosenschein, L. P. Kaelbling, "A Situated View of Presentation and Control", Artificial Intelligence vol. 73, 1995

[8] R. Dawkins, "The Selfish Gene", Oxford Press, 1990

[9] H. Kopetz, "Scheduling", §18 in "Distributed Systems" (ed. S. Mullender), ACM Press & Addison-Wesley, 1993

[10] D. H. Cansever, "Incentive Control Strategies For Decision Problems With Parametric Uncertainties", Ph.D. thesis, Univ. of Illinois Urbana-Champaign, 1985

[11] S. Parsons, M. Wooldridge, "Game Theory and Decision Theory in Multi-Agent Systems", AAMAS, vol. 5, 2002

[12] J. K. Kuchar, L. C. Yang, "A review of conflict detection and resolution modeling methods", IEEE Trans. Intelligent Transport. Systems, vol. 1, Dec. 2000

[13] L. Pallottino, E. M. Feron, A. Bicchi, "Conflict Resolution Problems for Air Traffic Management Systems Solved With Mixed Integer Programming", IEEE Trans. Intelligent Transport. Systems, Vol. 3, No. 1, March 2002

[14] A. Bicchi, L. Pallottino, "On optimal cooperative conflict resolution for air traffic management systems", IEEE Trans. Intelligent Transport. Systems, vol. 1, Dec. 2000