

Modeling A System Of UAVs On A Mission

Predrag Tomic*, Myeong-wuk Jang, Smitha Reddy, Joshua Chia,
Liping Chen, Gul Agha

Open Systems Laboratory, Department of Computer Science
University of Illinois at Urbana-Champaign

{p-tomic, mjang, sreddy1, jimchia, lchen2, agha}@cs.uiuc.edu

Abstract

We outline a parametric model of a system of **unmanned aerial vehicles (UAVs)** on a mission. The UAVs have to accomplish their mission composed of several tasks as efficiently as possible, while satisfying a heterogeneous set of physical and communication constraints. UAVs can be viewed as an example of a highly dynamic **multi-agent system (MAS)**. These UAVs may be required to autonomously make decisions, communicate, coordinate, adapt to rapidly changing environments and efficiently perform their tasks in real-time and under the limitations of local, incomplete and/or noisy knowledge of their surroundings. In particular, an individual UAV in our work can be viewed as an **agent**: it is autonomous, goal-driven, can affect and be affected by its environment, has its own behavior strategy, can communicate with its peers, and may find it beneficial to cooperate and coordinate not only to avoid collisions, but also in order to accomplish its set of tasks more effectively. We focus herein on two aspects of agent-based modeling of UAVs: modeling autonomous decision-making of the individual vehicles viewed as autonomous agents, and different models of UAV coordination.

Keywords: *unmanned aerial vehicles, agent-based modeling, agent coordination, agent autonomy*

1 Introduction & Motivation

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 application in a variety of contexts, e.g., they are being increasingly used for various surveillance, reconnaissance, and target-and-rescue missions. These UAVs carry sophisticated payloads, as they are designed to accomplish increasingly complex, multi-task missions. In particular, a typical UAV is equipped with certain *sensors* such as, e.g., radars. With these sensors, each 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 UAV’s sensors.

While trying to accomplish their mission (typically, a set of pre-defined and/or dynamically arising tasks as in the examples above), 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. This latter, higher form of cooperation (coordination) we also call *goal-driven* cooperation (respectively, coordination).

Not all kinds of UAVs can be reasonably considered *agents*; e.g., those that are remotely controlled

*Contact author

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 and manufactured. We are interested in UAVs that are not remotely controlled and that have the ability to make their own decisions in real time. This ability of autonomous decision making would “qualify” UAVs to be considered *autonomous agents* in the usual, computer science sense of the word.¹ 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., [BIC], [KUC], or [PAL]). What has drawn considerably less attention is modeling and analysis of the task-driven, goal-oriented behaviors of UAVs viewed as agents. Herein, we focus on some critical agent-based modeling paradigms applied to UAVs, namely, models of agent-to-agent coordination and the individual agent autonomy.

The rest of the paper is organized as follows. In §2, we give a high-level problem formulation, introduce some terminology, and reflect on some of the main assumptions of our modeling framework. The central part of the paper, section §3, is dedicated to identifying and discussing some of the most crucial design parameters of the model - those pertaining to modeling tasks, and UAV coordination and autonomous decision-making. We outline some possible extensions of the modeling framework in §4, and briefly summarize in §5.

2 Problem Formulation

In this section, we first briefly discuss the model at a high level: what are the UAVs trying to accomplish, both individually and as a single multi-agent system with a common goal, how we model these goals, and how we model UAVs strategies and mechanisms for accomplishing their goals. We also introduce the necessary terminology along the way. We conclude the section with a brief discussion about the main assump-

tions made in our model, and some of their implications.

A collection of N UAVs needs to accomplish a certain complex mission - such as any combination of surveillance, reconnaissance, target detection and/or target identification. 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 $\Pi_j(t)$. A UAV V_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. The *coordination model* describes how much UAVs cooperate and divide-and-conquer the IPs in order to function efficiently as a system.

At one extreme, if there is no coordination, each UAV acts entirely autonomously and, assuming no central control or other outside mechanisms, *selfishly*. At another extreme, in the leader-based coordination models, the UAVs that are not leaders obey their respective leaders, thereby sacrificing (temporarily or permanently) their autonomy as agents. Assuming unbounded communication radius for the UAV-to-UAV communication, a single-leader model becomes equivalent to a centralized model, with a possibly incomplete and/or noisy knowledge of the “world”. If, however, there are several leaders and different groups of “followers” associated with each leader, and if the radius of communication is non-trivially finite, then one may expect to encounter many of the fundamental challenges in distributed computation and communication, such as dynamic leader election and group formation problems, distributed consensus reaching, limits to local individual or group knowledge and their implications, and other issues (see, e.g., [TEL]).

Thus, 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

¹We leave aside the fact that there is no general agreement within the agent research community on *what exactly qualifies an entity* (such as a computer program, or a UAV with its sensors, effectors and software) *to be considered an agent* [FRA].

either fixed or dynamically changing (adaptable) *individual behavior strategy*. This strategy can be specified by an appropriate *individual behavior function*, Θ_i , that UAV V_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.

From a system's perspective, on the other hand, it is the successfulness of the entire mission that matters - not the gratification of individual agents. How to translate the individual utilities into the global utility maximization in the framework where, in general, both cooperation and competition are to be expected, is a challenging *incentive engineering* problem [CAN]. We address these issues elsewhere, where our constraint optimization based framework for modeling UAV missions is the central theme [TOS].

We conclude this section with stating some basic assumptions made both in the agent-based UAV model outlined herein, and in our simulation platform based on this model². One important assumption pertains to the nature of *time*. First, the time steps are discrete. Second, we assume the existence of the *global clock* and, therefore, *global time*. The *global clock* could be provided, say, by a central satellite-based control, and we also assume that all UAVs at any time have an instantaneous access to this global clock. The assumption about global time is (tacitly) made in most of the work on *MAS*, where the existence of a global clock is taken for granted. Without it, modeling and analysis of UAV-like distributed systems becomes considerably more difficult. The UAVs are assumed to communicate with one another (or, when applicable, with the 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 [TEL].

²For more on our UAV simulation and some experimental results, see [JAN].

³By *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 *payoff*, received from the environment.

⁴*DCS* by itself is, in general, computationally intractable, as even its centralized version, being a nontrivial generalization of the well-known *Satisfiability problem*, is **NP-hard**.

3 Modeling UAVs' Tasks, Coordination and Autonomy

We now discuss in some detail what we consider to be the most critical design parameters in our agent-based model of UAVs on a multi-task mission. The parameters of interest include ratio of the number of UAVs to the number of IPs, sensor ranges, communication ranges, a choice of a coordination model and strategy, a model of UAV's individual behavior strategy (or, in heterogeneous scenarios where UAVs are distinguishable, *strategies*), models of both individual and system *adaptability*³, and a choice of the model of UAVs' knowledge of their environment - such as whether this knowledge is local or global, perfect or noisy, etc.

Due to space limitations, we focus herein on three issues: quantitative models of interest points, models of individual UAV's autonomy, and UAV coordination models.

3.1 A Simple Model of UAV Tasks

Our model of UAVs on a mission emphasizes the goal-orientedness of UAVs as agents. UAVs are not merely flying around and trying to avoid colliding with one another or with other obstacles, but are actually trying to accomplish some set of tasks. Thus our model is trying to capture paradigms beyond the usual motion planning and obstacle avoidance problems. Mathematically, rather than having merely to solve an instance of *Distributed Constraint Satisfaction (DCS)* problem⁴ [YOK], our UAVs have, in addition to obeying a number of physical and communication constraints, also an *objective* or *utility function* that they strive to maximize. We discuss some possible ways of translating this desire to maximize utility into UAV's individual decision-making strategies in the next section. First, however, we outline our quantitative model of UAVs' tasks.

As our goal is to capture UAVs acting in possibly heterogeneous and highly dynamic environments, where not all tasks need be (i) identical, or (ii) known ahead of time, and where, in general, no central control is available to provide each UAV with the information about each of the tasks, a natural starting point in our quest to quantify the mission successfulness is to specify a simple *quantitative* notion of a *task*. We also need

a simple model for possible heterogeneity of different tasks. The basic task to be serviced in our model is an *interest point*.

Since not all tasks, or interest points, are necessarily the same, one simple way to capture this heterogeneity is to assign a time-dependent value function, $\Pi_j(t)$, to each interest point j , $1 \leq j \leq M$. When a UAV discovers an IP, it gets attracted by its value. Assuming a UAV is aware of two or more IPs at a given time step (by either having sensed those IPs, or because some other UAV has broadcast the IPs' locations and (estimated) values), the UAV needs to decide which IP is currently *most attractive* to it. The exception to this general rule is when, in task-driven coordination models, the UAV gets instructions from another UAV or the ground or satellite control what it is supposed to do next.

When a UAV arrives to an IP's location (or within a specified small distance from it), it starts consuming its value, thereby increasing its own utility or payoff. This value is consumed at some rate, d . In our simulations (see [JAN]), d was assumed to be a constant; in general, various models where d may depend on time, UAV's index i , and/or IP's index j , are worth considering.

To illustrate the general usefulness of the concept of IPs and their *value functions*, we sketch two special cases as examples.

First, let's assume that UAVs are on a surveillance mission. That is, each IP (or a set of IPs) needs to be revisited repeatedly. Some regions (represented by IPs) may be so important that they require presence (i.e., one or more UAVs essentially hovering or circling in their vicinity) at all times. The simple way to represent that in our model is to make the function $\Pi_j(t)$ independent of time: even though some UAVs are "consuming" the value of such an IP, while those UAVs' utility is increasing, the IP's value actually stay the same, thereby assuring that those IPs keep their attractiveness. If some IPs need to be revisited periodically but do not require ceaseless surveillance, then the function $\Pi_j(t)$ of such an IP j can be made periodic: once a UAV arrives to j , the value starts going down until UAV leaves; after some number of time steps, the value jumps back up again, thereby making j (more) attractive again, and thus prompting the UAVs to come back to this IP.

The second example is to consider an IP that is an actual target. Once its location has been discovered, one or more UAVs approach this target. Once a particular UAV, V_i , gets within some pre-specified

distance from the target, with probability p , the UAV consumes the target's entire value at once, with probability p_i . That is, Π_j of this IP goes to zero in one time step, and the UAV's payoff increases accordingly with probability p_i , whereas, with probability $1 - p_i$, V_i "misses" the target, so that the value Π_j remains unaltered. Whether the UAV V_i stays "at the target" j for one time step irrespective of the outcome, or for as many time steps as is needed for a success to occur, are different possibilities that can be modeled with different choices of the UAV's individual behavior functions.

What are, then, the general properties of function Π , and what parameters does it depend upon? At time step $t + 1$, the value $\Pi_j(t + 1)$ of IP j can be reasonably expected to depend on the value at the previous time step, $\Pi_j(t)$, the number of UAVs servicing this IP at time t , that we denote by $n_j(t)$, and the value consumption rate, d . In addition, the value function may explicitly depend on time. For instance, in the surveillance example above, the parameters $\Pi_j(t)$, d and $n_j(t)$ alone cannot capture a jump in $\Pi_j(t + 1)$ value due to some form of *aging*. Thus, the general form of the IP value functions we are interested in can be written as

$$\Pi_j(t + 1) = F(\Pi_j(t), n_j(t), d_{i,j}(t), t)$$

for some integer-valued or real-valued function F . A particular choice of Π that we have extensively experimented with [JAN] is

$$\Pi_j^*(t + 1) = \max\{\Pi_j^*(t) - d \cdot n_j(t), 0\},$$

where d is an integer constant. While this particular IP value function is always nonnegative, we also consider $\Pi(t)$ that can be negative; such value functions are useful whenever one wishes to model certain regions that UAVs should strive to avoid - such as, e.g., dangerous regions in the mission area of little actual value. Similarly, not all IP value functions need be nonincreasing in time like Π_j^* is; depending on what kind of IPs are modeled, value functions may be chosen to be, e.g., periodic or nondecreasing in time, and the like.

3.2 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. UAVs are

goal-driven entities. They fulfill their goals by servicing tasks. In our modeling framework, tasks are given as interest points and 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 values⁵, a UAV V_i evaluates its behavior function Θ_i that returns the index j^* of the IP that, if the UAV selects that IP as its next task to service, this choice is expected or to maximize the *estimated increase* in UAV's *utility*. Therefore, each UAV is assumed to behave *greedily*. 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 $\psi_j(t)$ be the position of IP j at time t , 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 is specified by

$$\Theta_i(t) = \arg\{ \max_{1 \leq j \leq M} G(\Pi_j, \|x_i - \psi_j\|, n_{j,r}, t) \}$$

where G is an integer-valued function that is increasing in Π_j and nonincreasing in distance of the UAV from the IP j given by $\|x_i(t) - \psi_j(t)\|$. This function specifies what IP should UAV V_i pick as the estimated short-term optimal choice. Notice that, for simplicity, *relative velocity* of a UAV with respect to the interest point is not taken into account. One example of a simple greedy individual behavior that fits the given description is

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

where it is assumed that the minimal distance of any UAV from any IP is strictly positive.

⁵In case of *mobile targets*, current velocities would be also needed.

⁶The infinite communication range assumption applies whenever this range is unbounded for all practical purposes - which is the case if, for example, the diameter of the entire mission area is less than the diameter of the UAV-to-UAV communication range.

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 UAV's autonomous agent 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 i -th UAV has received from other UAVs* at time steps $t' \leq t$. We discuss some models of UAV coordination next.

3.3 Models of UAV Coordination

We now discuss some possible design choices for modeling UAV-to-UAV coordination.

At one extreme, a single UAV becomes the "group leader", and this leader then broadcasts to other UAVs what it wants them to do. Typically, the leader is the *first* UAV that detects one or more IPs in a particular region of the *mission area*. In case of a tie (where two or more UAVs announce their claim to leadership simultaneously), the tie is broken according to some pre-specified rule (e.g., the UAV with the lowest index wins). Assuming the UAV-to-UAV communication radius is infinite⁶, and if the bandwidth availability is not an issue, this, single-leader scenario is very similar to a centralized control model. The one difference is that the leader's knowledge about the environment, in general, can be expected to be incomplete and/or imperfect (i.e., noisy), and that this knowledge is likely to dynamically change in often unpredictable ways.

While the single-leader model is perhaps the simplest to analyze and relatively easy to simulate, it also suffers from a number of shortcomings. These shortcomings can be divided into two general categories. One category are the usual problems with centralized or quasi-centralized control models, such as "ungraceful" degradation (due to a single point of failure), and the possible communication bottleneck at the leader node. The second category of potential troubles is peculiar to any situation where a single leader is itself "just another agent", whose sensors or communication links could be unreliable, whose local and possibly noisy "picture of the world" is imposed onto everyone else even though perhaps other agents have more accurate knowledge or more reliable links, and the like. Any

satisfactory solution, therefore, has to provide mechanisms for UAVs not only to “talk back” to their leaders, but also for the *ad hoc network* of UAVs to be able to detect possible troubles with the leader, and, if need be, dynamically reconfigure itself and elect a new leader.

At the other extreme, we consider models where there is no explicit task-driven coordination. In these models, not only is the whole group of UAVs autonomous, but also each individual UAV within this system acts as an entirely autonomous agent. In other words, each UAV simply follows its own strategy by re-computing its *individual behavior function*, irrespective of what others do. This individual behavior function is the agent’s strategy for maximizing its own individual payoff. The UAVs may still wish to communicate with one another, but there is no explicit coordination as to how to accomplish the mission more efficiently, how to divide-and-conquer tasks, etc. While, in the context of individual utility maximization this situation may be considered a default scenario, we view it as an extreme in the more appropriate, *joint utility* framework, where all UAVs have a single mission to accomplish, and where the successfulness of the entire system in accomplishing that mission - rather than that of the individual vehicles - is what matters [TOS]. Thus the “no goal-driven cooperation” scenario can serve as a yard stick with respect to which the effectiveness of various coordination strategies can be measured.

In between the two tentative extremes - the leader-based coordination on one, and the no explicit coordination model on the other hand - are many intermediate cases, and many possible coordination strategies. These intermediate coordination models are, typically, more flexible but also more complex than the leader-based approaches. We refer to this broad class of *intermediate* coordination strategies as *leaderless coordination models*.⁷

Let us summarize regarding the possible choices of a coordination model, and the tradeoffs these choices entail. Without any explicit goal-driven coordination, each UAV follows its own, pre-specified (but possibly adaptable) individual behavior strategy. The parameters to this strategy, encapsulated in each UAV’s *individual behavior function*, are provided by the environment, i.e. by the (possibly noisy) data about the environment gathered by the vehicle’s sensors, by the UAV’s knowledge of its “internal state”, and by the communication with other UAVs and/or the ground control that helps the UAV navigate, detect and avoid

possible collisions, and the like. Goal-driven coordination, then, entails different UAVs occasionally having to, temporarily or permanently, *abandon their individual behavior strategies*, and to begin doing what they have reached an agreement with other UAVs they ought to be doing. In general, coordination should aid the system to do better, and therefore any reasonable coordination strategy should be expected, given the same mission and the same set of tasks and other environment parameters, to perform at least as well as the corresponding coordination-free, purely autonomous individual behavior strategy. This intuition has been confirmed in some restricted scenarios that we have simulated [JAN]. We briefly argue in favor of a broader validity of our claim regarding the expected benefits of *the goal-driven coordination*. Assuming perfectly reliable communication and agent’s perfect knowledge of the environment - two typically unrealistic assumptions in practice - any model of goal-oriented coordination implies, at the very least, some *knowledge sharing*, i.e., more information available to each UAV than what is provided by the UAV’s sensors and communication about possible conflicts *alone*. More information, on the other hand (and leaving aside for the moment the issues of communication, storage and processing overheads), should not make the system or its components do any worse than without that additional information. Once the assumptions of perfect knowledge about the IPs and perfect communication links are dropped, however, the problem of choosing the right coordination strategy becomes both intuitively and analytically overwhelming, and therefore our view is that, in that case, there is no substitute for computer simulation and intense experimentation in various scenarios.

We also point out that, irrespective of the assumptions about reliability of communication links and UAVs’ sensors, any coordination *necessarily* requires more communication - and, in the real world, communication does not come for free. It also means that the agents have to execute an appropriate coordination mechanism, which may mean a considerable computational overhead - possibly prohibitively costly in a real-time application. Thus, in general, one can expect a tradeoff between the amount and nature of coordination, and the extra cost of this coordination. Hence, for example, in those scenarios where goal-driven coordination is expected or experimentally shown to be of a little benefit, by the Occam’s Razor principle, the simpler and less costly strategies with no explicit coordination would likely be preferable.

⁷Our UAV simulator has an implementation of a variant of each of these three categories of coordination models [JAN].

4 Some Future Plans

We now outline some possible extensions of the agent-based modeling framework presented in the previous sections. Regarding the nature of tasks, in addition to the interest points' values, we consider introducing different IP *types*. This naturally extends the model so that it can capture more heterogeneous scenarios, where both UAVs and IPs are *distinguishable* and, in particular, the UAVs become specialized: a UAV can handle only those IPs that are of a compatible type.

An important remark, leading to a considerable model extension, is that currently, for each IP j , its Π_j is assumed to represent the j 's true (i.e., objective) value, irrespective of which UAV may see IP j , and from how far away. That is, the tacit assumption is that, once an IP is discovered by the UAVs, all the UAVs who are aware of this IP's existence immediately also know its *exact* current value. This *perfect knowledge* assumption can be appropriately relaxed for the sake of the model's realism. The more reasonable assumption is that each UAV has its own, *local* and *imperfect* estimate of Π_j of those IPs j that the UAV knows about. In case of imperfect and/or incomplete local knowledge of the tasks, each UAV's individual behavior function, Θ_i , naturally becomes probabilistic, i.e., UAVs are required to make decisions in the presence of uncertainty. In particular, Θ_i is now a function of an *estimated attractiveness* of each IP j that UAV V_i is aware of, rather than of the exact actual value Π_j as this actual value, in general, need not be known to V_i . Ability of decision-making under uncertainty is often considered one of the hallmarks of most autonomous agents one finds in the literature [PAR].

5 Summary

We have presented an agent-based approach to modeling UAVs on a mission made of multiple tasks. Assuming no remote control, the UAVs share many properties characteristic of autonomous agents, such as goal-orientedness, pro-activeness, the ability to affect and be affected by the environment, autonomous decision-making under uncertainty, and peer-to-peer communication, coordination and cooperation. The current emphasis of our modeling is on some simple yet interesting classes of autonomous agent behavior strategies and of goal-driven agent coordination. It is our hope that

future extensions and improvements of the modeling framework presented herewith, and simulations guided by similar models, will increase the understanding and enhance future design of both autonomous UAVs and other intelligent vehicles, and multi-agent systems in general.

Acknowledgment: Work presented herewith was supported by the **DARPA IPTO TASK Program**, contract number **F30602-00-2-0586**.

Bibliography

- [BIC] A. Bicchi, L. Pallottino, "On optimal cooperative conflict resolution for air traffic management systems", IEEE Trans. Intelligent Transportation Systems, vol. 1, Dec. 2000
- [CAN] D. H. Cansever, "Incentive Control Strategies For Decision Problems With Parametric Uncertainties", Ph.D. thesis, Univ. of Illinois Urbana-Champaign, 1985
- [FRA] 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
- [JAN] M. Jang, S. Reddy, P. Tosic, L. Chen, G. Agha, "An Actor-based Simulation for Modeling Coordination among UAVs", submitted to 15th Euro. Simulation Symposium & Exhibition, Delft, The Netherlands, 2003
- [KUC] J. K. Kuchar, L. C. Yang, "A review of conflict detection and resolution modeling methods", IEEE Trans. Intelligent Transport. Systems, vol. 1, Dec. 2000
- [PAL] 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
- [PAR] S. Parsons, M. Wooldridge, "Game Theory and Decision Theory in Multi-Agent Systems", AA-MAS, No. 5, 2002, Kluwer Academic Publishers
- [TEL] Gerald Tel, "Introduction To Distributed Algorithms", 2nd ed., Cambridge Univ. Press, 2000
- [TOS] P. Tosic, M. Jang, G. Agha, "Constraint Optimization Framework for UAV Surveillance Mission Problem", in preparation
- [YOK] M. Yokoo, K. Hirayama, "Algorithms for Distributed Constraint Satisfaction: A review", AA-MAS, Vol. 3, No. 2, 2000