

Verification of cooperative multi-agent negotiation with the Alloy Analyzer

Rodion Podorozhny¹, Sarfraz Khurshid², Dewayne Perry², and Xiaoqin Zhang³

¹ Texas State University, San Marcos, TX 78666,
rp31@txstate.edu

² The University of Texas, Austin, TX 78712,
{khurshid, perry}@ece.utexas.edu

³ The University of Massachusetts, North Dartmouth, MA 02747,
x2zhang@umassd.edu

Abstract. Multi-agent systems provide an increasingly popular solution in problem domains that require management of uncertainty and high degree of adaptability. Robustness is a key design criteria in building multi-agent systems. We present a novel approach for the design of robust multi-agent systems. Our approach constructs a model of the design of a multi-agent system in Alloy, a declarative language based on relations, and checks the properties of the model using the Alloy Analyzer, a fully automatic analysis tool for Alloy models. While several prior techniques exist for checking properties of multi-agent systems, the novelty of our work is that we can check properties of coordination and interaction, as well as properties of complex data structures that the agents may internally be manipulating or even sharing. The suggested work is the first application of Alloy to checking properties of multi-agent systems. Such unified analysis has not been possible before.

1 Introduction

As software systems grow in complexity, and reliability becomes a more vital concern, the need for robust designs of systems grows more critical. The Alloy tool-set provides a software design framework that enables the modeling of crucial design properties as well as checking them. Alloy [7] is a first-order, declarative language based on relations. The Alloy Analyzer [9] provides a fully automatic analysis for checking properties of Alloy models.

The Alloy language provides a convenient notation based on path expressions and quantifiers, which allow a succinct and intuitive formulation of a range of useful properties, including rich structural properties of software. The Alloy Analyzer performs a bounded exhaustive analysis using propositional satisfiability (SAT) solvers. Given an Alloy formula and a *scope*, i.e., a bound on the universe of discourse, the analyzer translates the Alloy formula into a boolean formula in conjunctive normal form (CNF), and solves it using an off-the-shelf SAT solver.

The Alloy tool-set has been used successfully to check designs of various applications, such as Microsoft's Common Object Modeling interface for interprocess communication [3], the Intentional Naming System for resource discovery in mobile networks [1], and avionics systems [5], as well as designs of cancer therapy machines [8].

Most of the prior applications of Alloy have abstracted away from properties of multi-threaded systems. We explore the use of Alloy in designing a rich class of distributed systems, known as the *multi-agent systems* (MAS). A common definition of a multi-agent system [15] stipulates that an agent is an autonomous, interacting, intelligent (i.e. optimizing its actions) entity. Any MAS is a distributed system but not every distributed system can be categorized as an MAS by the above mentioned definition.

Management of uncertainty via adaptability and ability to provide a satisficing solution to otherwise intractable problems are distinguishing features of multi-agent systems compared to centralized or other distributed systems. An agent knows of a great variety of methods to solve their local tasks, thus an agent can tailor a method of achieving a goal according to resource availability for data processing, information exchange and sources of information. Agents, due to their interaction, are capable of influencing the choices of methods both by themselves and by other agents due to recognition of various kinds of relationships between their subtasks that can be generalized as redundancy, facilitation and enabling [12]. Agents can decide the degree to which an environment state, their own state, and their partial knowledge about states of other agents influence the amount of their contribution to the solution of a task imposed on the whole MAS. Unlike components of other distributed systems, an agent can refuse a request or can choose not to answer. At the same time, other agents are prepared to deal with a possibility that their requests will be refused or not answered. This freedom of choice, in a way, defines an agent's autonomy and distinguishes it from a component of a conventional distributed system. Thus, due to the above mentioned capabilities, agents are able to adapt their solution methods to the dynamics of the environment [11].

Some MAS have explicit specifications of interaction protocols between the agents. There has been a plethora of work on verification of MAS systems. Such approaches as model-checking ([16], [13], [10], [2]), Petri-nets and situation-calculus [6] have been applied to MAS verification. The vast majority of recent work on MAS verification are various applications of model checking that take into account peculiarities of properties that are desired to be verified in MAS. The peculiarities of such properties usually are a consequence of bounded rationality in agents. Thus the set of operators (modalities) for property specifications is often extended to include operators as *agent beliefs*, *desires*, *intentions*. Once such additional operators are introduced, usually a method is suggested to map a property specification that uses these MAS-specific operators into a formalism understood by off-the-shelf model-checkers, e.g. into the propositional LTL.

Examples of properties might be: "every request for a quote is answered within 4 time steps" [2], "for all paths in each state if agent Train1 is in the tunnel then agent Train1 knows that agent Train2 is not in the tunnel" [10], "when sender is about to send an acknowledgment then it knows that the receiver knows the value of the bit that was most recently sent" [13] and "some agent i eventually comes to believe that agent i intends that i believes variable a has the value 10" [16].

As we can see from these examples most properties are some sort of reachability properties on a state transition model of a MAS. It is understandable as model checking is essentially an efficient brute-force global state transition graph reachability analysis. ConGolog uses situation calculus which is also most suited for the specification and analysis of event sequences, not data structures.

We expect that the Alloy analyzer is better suited for checking data rich properties than model checking approaches. In case of a model checking approach one needs to create a number of particular instances of data structures. For some kinds of data structures the size of such an enumeration can be prohibitively large, not to mention the fact that a generator of instances of that data structure has to be created [Vaziri thesis ref]. The Alloy approach allows verification of data rich properties via capturing them in a simple first-order logic formula. An example of such a property might be acyclicity in trees. Even though the size of a formula needed by the Alloy is somewhat larger than that needed by a model checker, a chance that in a non-worst case we can find a counter-example earlier than before having checked all the instances of data structures makes the Alloy approach a better choice (provided the likelihood of the worst case is low).

We explore an application of Alloy with its relational logic specification language to multi-agent systems specifically focusing on properties of data structures in addition to event sequences. We expect to be able to check properties of the following format: "if agent a receives a data structure that satisfies property p then eventually agent a will enter state sa if it believes that agent b is in state sb", "if agent a is in state sa and its task structure t1 satisfies property p1 then on reception of data structure msg1 (from agent b) agent a will modify t1 with some part of msg1 such that t1 will preserve property p1" and so on. In addition to being able to specify such properties, we can check the adequacy of a testcase set by determining if a property was evaluated.

2 Brief overview of Alloy

As software systems steadily grow in complexity and size, designing such systems manually becomes more and more error-prone. The last few years have seen a new generation of design tools that allow formulating designs formally, as well as checking their correctness to detect crucial flaws that, if not corrected, could lead to massive failures. The Alloy tool-set is one such design tool that is rapidly gaining prominence [7, 9]. The user formulates their design in the Alloy language, which is a first-order logic (with transitive closure) based on relations, and checks the correctness properties using the Alloy Analyzer.

The Alloy tool-set has been used successfully to check designs of various applications, such as Microsoft's Common Object Modeling interface for interprocess communication [3], the Intentional Naming System for resource discovery in mobile networks [1], and avionics systems [5], as well as designs of cancer therapy machines [8].

The Alloy language provides a convenient notation based on path expressions and quantifiers, which allow a succinct and intuitive formulation of a range of useful properties, including rich structural properties of software. Much of Alloy's utility, however, comes from its fully automatic analyzer, which performs a bounded exhaustive analysis using propositional satisfiability (SAT) solvers. Given an Alloy formula and a *scope*, i.e., a bound on the universe of discourse, the analyzer translates the Alloy formula into a boolean formula in conjunctive normal form (CNF), and solves it using an off-the-shelf SAT solver.

We present an example to introduce the basics of Alloy.
Let us review the following Alloy code for a DAG definition:

```

module models/examples/tutorial/dagDefSmall

sig DAG {
  root: Node,
  nodes: set Node,
  edges: Node -> Node
}
sig Node {}

```

The keyword `module` names a model. A `sig` declaration introduces a set of (indivisible) atoms; the signatures `DAG` and `Node` respectively declare a set of DAG atoms and a set of node atoms. The *fields* of a signature declare relations. The field `root` defines a relationship of type `DAG x Node` indicating that only one node can correspond to a DAG by this relationship. The absence of any keyword makes `size` a total function: each list must have a size. The field `nodes` has the same type as `nodes` but maps a DAG onto a set of nodes defining a partial function. Alloy provides the keyword `set` to declare an arbitrary relation. The field `edges` maps a DAG onto a relationship, i.e. on a set of tuples `Node x Node`, thus defining edges.

The following *fact* constrains a graph to be a DAG:

```

fact DAGDef {
  nodes = root.*edges
  all m: Node | m !in m.^edges
}

```

The operator `*` denotes reflexive transitive closure. The expression `root.*edges` represents the set of all nodes reachable from the root following zero or more traversals along the `edge` field. A universally quantified (`all`) formula stipulates that no atom `m` of signature `Node` can appear in traversals originating for that atom `m`. The operator `^` denotes transitive closure.

Here are some other common operators not illustrated by this example. Logical implication is denoted by `=>`; `<=>` represents bi-implication. The operator `-` denotes set difference, while `#` denotes set cardinality and `+` - set union.

To instruct the analyzer to generate a DAG with 6 nodes, we formulate an empty predicate and write a `run` command:

```

pred generate() {}

run generate for 6 but 1 DAG

```

The scope of 6 forces an upper bound of 6 nodes. The `but` keyword specifies a separate bound for a signature whose name follows the keyword. Thus we restrict a generated example to 1 DAG.

3 Subject system details

As the subject of our analysis we have chosen a cooperative multi-agent system with explicit communication and with a utility-based proactive planning/scheduling.

A multi-agent system is cooperative if it can be assumed that agents strive to collectively contribute to reaching some common goal. In such a cooperative MAS, agents are willing to sacrifice their local optimality of actions if they are convinced that such a sacrifice will help increase the global optimality of the combined actions in the whole MAS. For simplicity we also assume there are no malicious agents in the chosen MAS.

The explicit communication implies that agents interact with one another by sending explicit messages instead of passive observation (implicit communication). The interaction involves exchange of partial information about agents' states. The MAS chosen for analysis also provides for proactive planning (choice of a set of actions to reach a goal) and scheduling (creating a partial order of the chosen actions, tying them to a timeline, assigning resources). A certain utility function is used to measure goodness of produced schedules. By exchanging information about the utilities of local schedules the agents can reason about the sets of local schedules that can increase the combined utility of agents' actions. The negotiation protocols in MAS with explicit communication embody the decision-making of a trade-off between the amount of exchanged partial state information and the degree to which a combined utility can be increased thanks to the communication. Thus negotiation protocols regulate the efficiency of MAS operation somewhere between the two ends of a spectrum. On the one end - a completely centralized system that knows complete information about the states of all components where it is computationally too costly (e.g. due to the computational complexity of scheduling algorithms). On the other end of the spectrum are the distributed systems without any such (explicit or implicit) communication between agents that impacts the set and sequence of their actions. Thus components of such a system are acting only according to their local optimality criteria. They do not act synergistically, they do not recognize possible ways in which they can facilitate their combined actions or eliminate redundancy.

3.1 Task allocation problem in cooperative negotiation

One of the topics of negotiation in a cooperative MAS can be task allocation. In this problem, the agents are working to find how and when to perform a task so as to maximize the combined utility. Suppose that one agent (A) knows that accomplishing a task of certain M4 type is beneficial for the utility of its local schedule. Yet, it cannot accomplish that task, e.g. due to the limitations of its functionality or resources. Thus such an agent would like some other agent to accomplish that M4 type task by a certain deadline. In this example we omit the reasoning behind determination of which agents are capable of accomplishing the M4. Let us assume that a certain agent B has the potential capability to accomplish a task of M4 type. Agent B has its own schedule that did not initially imply that M4 had to be performed. Once agent A notifies agent B of its request to accomplish an instance of M4 by a certain deadline these agents can engage in a negotiation of whether agent B is willing to perform M4 and by what time. Thus these two agents have to resolve a task allocation problem. It is a negotiation protocol that governs decision making behind a resolution of this task.

3.2 Property examples derived from requirements

We can describe several properties informally at this stage, before we fix the assumptions of the MAS design further.

Some of the informal properties that are likely to be useful for such a negotiation:

1. negotiation must terminate;

2. the utility of the agreed upon combination of schedules must eventually increase throughout the course of negotiation even though occasional decreases are allowed; i.e. the negotiation must eventually converge on some choice of schedules that provides a local optimum of the combined utility (here local is used in the sense of restrictions on action set and time deadline, not in the sense of local to a single agent);
3. if agent B (the one who is requested to do an additional task) agrees to accomplish the task at a certain point in negotiation then it cannot renege on that agreement in the course of subsequent negotiation (somewhat related to the need to converge); and,
4. the beliefs of one agent about an abstraction of partial state of another agent obtained as a result of negotiation should not contradict the actual state of that other agent.

3.3 Choice of the analyzed system

Next we will provide greater detail about the design of the chosen MAS. This detail will let us illustrate the task allocation problem introduced generally above and to formalize a property. The chosen system has been developed in the MAS laboratory headed by Prof. Victor Lesser at the University of Massachusetts, Amherst. It has been used as a testbed for a great number of experiments and technology transfer demonstrations in the area of MAS (sensor network by Horling, Lesser's refs). An agent is combined of several components that include a problem solver, a negotiation component, among others. The problem solver provides a schedule based on a current set of task structures assigned for execution. The negotiation component drives the execution of negotiation protocols, it is aware of protocol specifications and keeps track of current states of negotiation instances undertaken by its agent. The task structures are specified in the TÆMS language [4]. The schedules are provided by the Design-To-Criteria (DTC) scheduler ([14]) developed by Dr. Tom Wagner which is invoked as part of the agent's problem solver component operation. The DTC takes as input a task structure in TÆMS and a utility function specification and provides as output a set of schedules ranked by their utilities.

In this system a simplified description of an agent's cycle is as follows:

1. *Local scheduling*: in response to an event requesting a certain task to be performed, obtain a number of high ranked schedules by utility;
2. *Negotiation*: conduct negotiation(s) within a predefined limit of time; and,
3. *Execution*: start execution of the schedule chosen as a result of negotiation(s).

The actual cycle of agent's operation is more complex as an agent can react to various kinds of events that it can receive at any of the mentioned cycle stages.

3.4 Relation between protocol FSMs, task structures, offers and visitations

Next we describe the task allocation problem in terms of this design. More details about the cooperative negotiation example can be found in [17]. The negotiation protocol

of an agent starting the negotiation (agent A), the contractor, is given in Fig. 3. The negotiation protocol of an agent responding to the request (agent B), the contractee, is given in Fig. 4.

Let us assume that agent A needs a certain non-local task (this means that an agent is not capable to do that task even though it appears in one of its task structures) to be performed by some other agent. The negotiation's goal is to increase the combined utility of actions of both agents by choosing a particular way to perform the non-local task at a particular time.

In the description that follows we mention notions of a protocol FSM, task structures, offers and execution paths encoded in visitations. These notions are related to one another in the following way. The design of the particular MAS we are analyzing contains a module called an agent [ref on JAF in SELMAS]. This module itself is an aggregate of several submodules. One of these submodules is "Negotiation" submodule that is responsible for encapsulating knowledge about various protocols known to an agent. These protocols are encoded as FSMs with states corresponding to abstractions of states of an agent in negotiation and transitions attributed with a trigger condition and action, as is quite often the case for FSM representations. A sequence of visitations corresponds to a path from a start node of such a protocol FSM to one of the final nodes. A task structure of an agent captures its knowledge about multiple ways in which a certain task can be accomplished. The root of a task structure corresponds to a task that an agent is capable of accomplishing. The leaves of a task structure correspond to atomic actions both the set of which and partial order can vary to reflect the way to accomplish an assigned task in a "utility-increasing" (but not guaranteed to be optimal) way. As an agent progresses through a protocol according to an FSM, the agent's task structure changes to reflect the agent's changing knowledge about other agent's state throughout a negotiation. Thus there are certain properties imposed on a task structure that must hold while an agent is in certain states of a protocol FSM. A collection of task structures, in a way, determines an agent's functionality analogously to a set of function signatures that would define an interface of a module. The roots of task structures serve similar purpose to function signatures at the agent level of abstraction of describing a software system. An outside event corresponding to a request to accomplish a certain task triggers an agent's reasoning about whether it can accomplish that task considering an agent's knowledge about the way to accomplish that task, that agent's state, the environment state and partial states of some other agents in the same MAS. The result of that reasoning is the current schedule that "interweaves" instances of atomic actions from various tasks currently assigned to that agent in a time-oriented partial order. That current schedule can be changed dynamically, as it is being executed, in response to agents' changing opinion about most reasonable schedule for a certain moment in time. We do not consider execution of schedules, we are only focusing on the negotiation phase in which schedules always cover future time intervals. An offer is a data structure generated by actions associated with FSM transitions. An offer encapsulates the parameters of a particular schedule formed on the basis of the agents' task structures, such as quality achieved, start time and finish time. The agents negotiate over these parameters.

Another submodule of an agent module is "Communication". The "Negotiation" submodule relies on "Communication" in a fashion similar to how a networking ap-

plication relies on TCP/IP protocols. The design intentionally separated the concern of ensuring reliable communication and naming mechanisms from the concern of ensuring that a certain “utility-increasing” protocol is followed during a negotiation between a pair of agents. Thus the issues of identifying agents to communicate with for a particular purpose were separated from the “negotiation” submodule by the authors of the MAS system we analyze. This was done to simplify their own analysis, to separate concerns. Our Alloy specification reflects that separation.

In a way, the task structure specifies all possible behaviors of an agent responsible for achieving the goal embodied by a task structure’s root. During the stages of *Local scheduling* and *Negotiation* the task structure can be modified, thus modifying specification of a set of behaviors of an agent during an *Execution* stage. The behavior of an agent during the stages of *Local scheduling* and *Negotiation* is static, i.e. it is not modified during run-time. A schedule agreed upon as a result of *Negotiation* is a selected behavior (execution path) from a set of behaviors that was modified at run-time (represented by a task structure; to be performed in the *Execution* stage). Thus a property we describe below checks certain well-formedness of a behavior specification modified at run-time and correctness of an implementation responsible for the modification.

3.5 Details of the task allocation problem in the chosen design

On Fig. 1 we see two task structures. One task structure, with the root TCR, was assigned to agent A, the other, TCE, was assigned to agent B. This assignment was due to requests sent from the environment (e.g. a human or other automated system). TCE and TCR turned out to be non-leaf nodes with elaborations. So agent A sent TCR structure to its local scheduler, agent B did the same for TCE.

Thus Agent A receives the following schedules from its problem solver component: M1, M2, M3, M4 - highest utility M1, M2, M3, M5 - lower utility, feasible Agent B receives the following schedule: B3, B4 - highest utility

Next, agent A identifies M4 in its best schedule as non-local. It sends a request to agent B to do it. The fact that agent A knows that B can do M4 is hardwired for the example. The request initiates an instance of negotiation. Agent A plays the role of a contractor, agent B - that of a contractee. Agent B must see whether it can do M4 by the deadline agent A needs it, while accomplishing its current task TCE within the constraints. This is done by modifying the “currently reasoned about” structure and submitting it to the scheduler that will report if such a schedule is possible and, if yes, then with what utility.

The TCE structure must be modified preserving its well-formedness constraints (e.g. functional decomposition remains a tree); and forcing an M4 into a schedule by choosing appropriate quality of M4 that reflects the combined utility of both schedules (chosen by A and by B). Fig. 2 shows agent B’s task structure updated with an M4. The quality attribute of M4 must be such that the problem solver of agent B must produce feasible (though not necessarily high ranking) schedules that contain M4 and still accomplish the original TCE task.

Even if the agent B’s local scheduler returns an acceptable schedule (has M4 in it and the original TCE is accomplished with the constraints on time and quality), agent A can request to make a tighter fit.

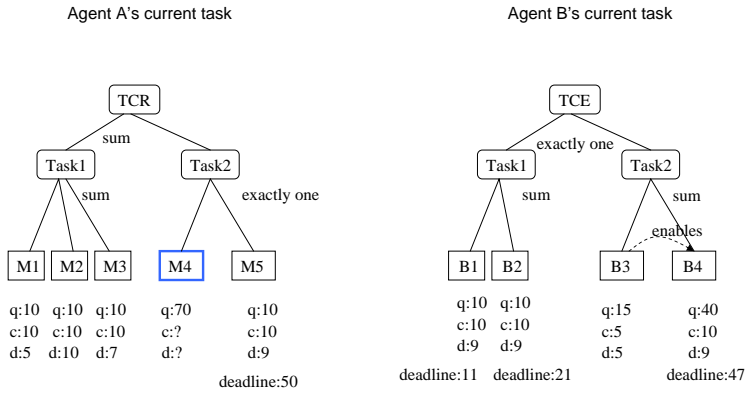


Fig. 1. Pre-negotiation task structures.

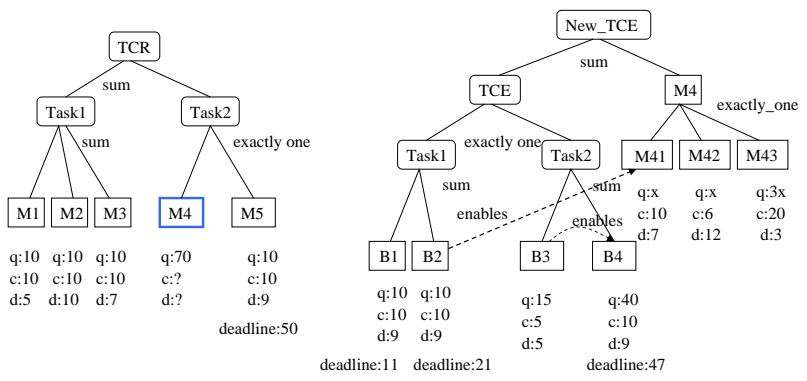


Fig. 2. Post-negotiation task structures.

With this description in mind we can rephrase property 3 in terms of the TÆMS structures and negotiation protocol specifications in Figures 3 and 4 as:

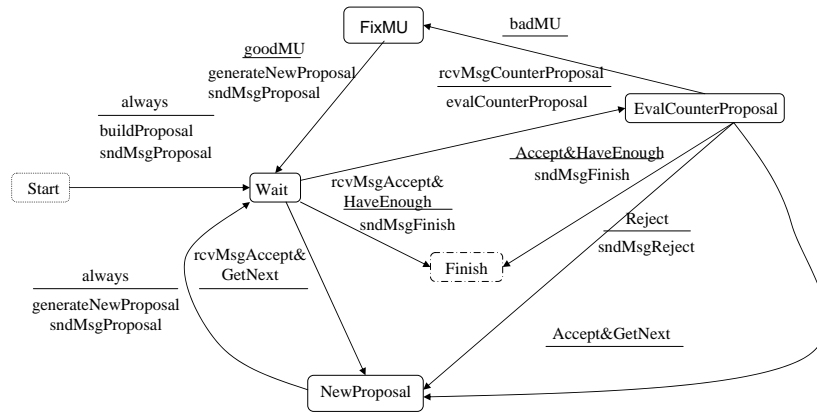


Fig. 3. Contractor's FSM.

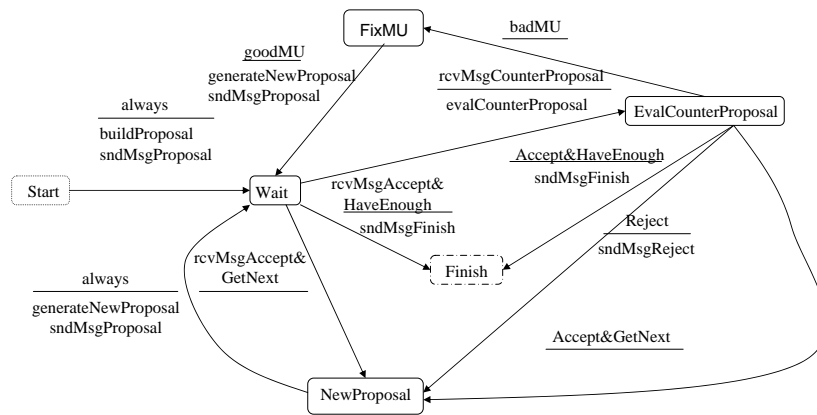


Fig. 4. Contractee's FSM.

Since agent B reaches state "Accept" its task structure must contain a subtree corresponding to task M4 and M4 must appear in a feasible schedule returned to agent B.

4 Alloy specification for the negotiation model

Our approach implies modeling particular paths traversed in the agents' negotiation finite state machines (FSMs) in response to certain testcases. Thus we check an abstraction of an execution path in a particular implementation. Both FSMs contain cycles. If a cycle diameter can be modeled with the scope that can be processed by the Alloy analyzer then we can iteratively check a certain property on an execution path that corresponds to multiple iterations of a cycle.

The negotiation protocols and goal trees described in section 3 had to be simplified to have a tractable scope for the Alloy analyzer. The simplifications include:

1. ignoring attributes of goal tree nodes (quality, duration, cost);
2. ignoring attributes of offers (mutual utility gain, cost, earliest start time);
3. ignoring attributes of schedules (start time and finish time of actions); and,
4. simplifying goal trees by removing intermediate nodes (e.g. no Task_1, Task_2) and reducing the number of leaf nodes (e.g. only B1 and B3 left in agent B's goal tree).

The actual models used for analysis also contain only those atoms that are necessary for verifying a property at hand. Thus transitions that were not traversed by a modeled execution path and associated states were removed.

This amount of simplification was necessary to make the analysis feasible. Earlier we constructed a more detailed Alloy specification of the analyzed system. The Alloy analyzer was not able to cope with such a specification. We had to reduce its size gradually while still keeping the analysis useful. We expect that the next generation of the Alloy analyzer, Kodkod, would be able to deal with a larger specification.

The resultant Alloy model of the MAS for the purpose of verifying our assertions consists of 3 modules. One module,

`negProtocol12_labridgeDataProp`, models the FSMs, Visitations of transitions through the FSMs (paths specified by transitions), and assertions. Two more modules model the data structures manipulated by the agents - their goal trees and schedules. Let us briefly go over the Alloy models in these modules.

The `negProtocol12_labridgeDataProp` defines signatures for `State`, `Transition`, `Visitation` and `Offer`. Thus an FSM is modeled by constraining atoms of `State` and `Transition` signatures via the "fact" construct. A `Transition` signature contains fields for source and destination states, a set of visitations of that transition by a path and a set of transitions outgoing from the destination state of the transition.

```
abstract sig State {}

abstract sig Transition {
  source, dest: State,
  visit: set Visitation,
  nextTrans: set Transition
}

fact Injection { all t, t': Transition | t.source =
  t'.source && t.dest =
  t'.dest => t = t' }

abstract sig Visitation {
  trans: lone Transition,
  nextVisit: lone Visitation,
```

```

    offer: lone Offer
  }
}
fact VisTransConsistent {
  all visitation: Visitation | visitation in
    visitation.trans.visit
}

```

The `treeDefSmall` module models a task structure (goal tree) of an agent.

```

module models/examples/tutorial/treeDefSmall

abstract sig Tree {
  root: Node,
  nodes: set Node,
  edges: Node -> Node
}

{
  nodes = root.*edges
  all m: Node | m !in m.^edges
}

abstract sig Node {}

one sig TCR, M3, M4, M5, TCE, B1, B3, New_TCE extends Node{}

one sig AgentB_preTaskStructTCE extends Tree {}
fact AgentB_preTaskStructTCEDef {
  AgentB_preTaskStructTCE.root = TCE
  AgentB_preTaskStructTCE.nodes = TCE + B3
  AgentB_preTaskStructTCE.edges = TCE->B3
}

one sig AgentB_postTaskStructTCE extends Tree {}
fact AgentB_postTaskStructTCEDef {
  AgentB_postTaskStructTCE.root = New_TCE
  AgentB_postTaskStructTCE.nodes = New_TCE + TCE + B1 + M4
  AgentB_postTaskStructTCE.edges = New_TCE->TCE +
    New_TCE->M4 + TCE->B1
}

```

The `schedDefSmall` module models a schedule data structure of an agent. It imports the `treeDefSmall` so that schedule items could point to the nodes of task structures.

```

module models/examples/tutorial/schedDefSmall
open models/examples/tutorial/treeDefSmall

abstract sig SchedItem {
  activity: Node
}

one sig SchedItemM3 extends SchedItem{}
fact SchedItemM3Def {
  SchedItemM3.activity = M3
}

one sig SchedItemM4 extends SchedItem{}
fact SchedItemM4Def {
  SchedItemM4.activity = M4
}

one sig SchedItemM5 extends SchedItem{}
fact SchedItemM5Def {
  SchedItemM5.activity = M5
}

```

```

one sig SchedItemB1 extends SchedItem{}
fact SchedItemB1Def {
  SchedItemB1.activity = B1
}

one sig SchedItemB3 extends SchedItem{}
fact SchedItemB3Def {
  SchedItemB3.activity = B3
}

abstract sig Sched {
  items: set SchedItem,
  precedenceRel: SchedItem -> SchedItem
}

one sig AgentAschedWithNL extends Sched {}
fact AgentAschedWithNLDef {
  AgentAschedWithNL.items = SchedItemM3 + SchedItemM4
  AgentAschedWithNL.precedenceRel =
    SchedItemM3->SchedItemM4
}

one sig AgentAschedWithOutNL extends Sched {}
fact AgentAschedWithOutNLDef {
  AgentAschedWithOutNL.items = SchedItemM3 + SchedItemM5
  AgentAschedWithOutNL.precedenceRel =
    SchedItemM3->SchedItemM5
}

one sig AgentBschedWithNL extends Sched {}
fact AgentBschedWithOutNLDef {
  AgentBschedWithOutNL.items = SchedItemB1 + SchedItemM4
  AgentBschedWithOutNL.precedenceRel =
    SchedItemB1->SchedItemM4
}

one sig AgentBschedWithOutNL extends Sched {}
fact AgentBschedWithNLDef {
  AgentBschedWithNL.items = SchedItemB3
}

```

The consistency of the model has been successfully checked with an empty stub predicate. The analyzer found a solution.

5 Alloy specification for the properties

The paths of execution of the two negotiation protocols are represented by atoms of the `visitation` signature. Thus it is via these atoms that we express a property that can be informally phrased as “If agent A is led to believe by a certain sequence of communications that agent B reaches a certain state then agent B should have indeed reached that state, having been subjected to the same changes of observed environment as agent A”. This informal statement pinpoints such feature of agents in a MAS as bounded rationality. The property checks for consistency between a certain abstraction of other agent’s state (agent B) that a certain agent (A) obtains via communication. In the case of the particular system we used the communication is explicit. By modeling the environment sensed by agents we could allow for checking such properties based on implicit communication.

More specifically, in view of the simplifications we made, a property of this kind can be informally restated as “if agent A reaches state `EvalCounterProposal` then agent

B should have reached state `wait2` and beginning since that state, agent B's current schedule data structure should have contained an instance of atomic action M4". Below we can see how this property is formally expressed in the Alloy's relational algebra.

```
assert AgentABeliefCompliesWithAgentBState {
  (some visitation: Visitation |
   visitation.trans.dest = EvalCounterProposal) =>
  (some visitation': Visitation |
   visitation'.trans.dest = Wait2 &&
   M4 in visitation'.offer.agentBTaskTree.nodes)
}
```

The assertion has been successfully checked. No counterexamples were found for the path containing visitations that corresponded to the expected states and data structure conditions. Conversely, once an inconsistency between agent A's belief and agent B's state and data structures has been introduced into visitations, the analyzer pinpointed a possible counterexample.

We have also translated an Alloy specification of this property into a dynamic assertion in Java using a systematic translation approach [undergrad's thesis ref]. Consequently we were able to test an actual implementation for partial correctness.

6 Specification difficulties

The main difficulty is keeping the Alloy model under a tractable scope while checking useful properties. In the case of the design of this particular MAS the protocols are specified via FSMs with loops. Thus we can check properties only within the scope of the FSM's diameter. Other difficulties are due to highly dynamic, hard to predict behavior of sensing agents. One has to classify the dynamics of the environment sensed by the agents and check the properties within each such situation. For instance, in the example used in this paper we can classify the situations based on combinations of "best" schedules of the 2 agents with regard to including the non local task (M4) into their schedules. Some of the possible combinations (for all cases agent A has M4 in its best schedule):

- agent B does not have M4 in its best schedule; the local utility of agent B's schedule outweighs the combined utility if agent B is forced to do M4;
- agent B does not have M4 in its best schedule; the local utility of agent B's schedule is below the combined utility if agent B is forced to do M4;
- agent B has M4 in its best schedule too, but not within the timeframe agent A needs M4 to be finished
- agent B has M4 in its best schedule too, it is within the timeframe agent A needs M4 to be finished

It should be possible to provide an Alloy model so that these combinations would not have to be specified explicitly. Instead, the Alloy analyzer itself would check over all the alternatives it sees in the model. A straightforward approach of modeling the attributes of the nodes in the agents' goal trees results in a too large scope for the Alloy to handle. Perhaps the attribute values should be abstracted as features of the structure of goal trees, not as numerical values.

7 Conclusions and Future Work

We have created and validated a model for verifying data structure rich properties of a cooperative multi-agent system using a manually created execution path. To our knowledge, our work is the first application of the Alloy analyzer for checking properties of a multi-agent system.

One of the next immediate steps is to check these properties on a simplified trace produced by a full scale MAS system. Another step is checking a property on all interior paths of a loop in an FSM. One more interesting property would involve checking if an elaboration of the non-local task is “interwoven” in one of the many alternative ways into the goal tree of agent B. After all our checking of mere presence of an atomic task M4 in a schedule is a simplification. We expect that checking such a more complicated and a more realistic case might highlight Alloy’s advantage due to the declarative nature of its relational algebra.

We also envision checking properties of a MAS in a model that recognizes the possibility of more than 2 agents interacting. In this case we will highlight the advantage of the Alloy over model checking approaches in that we will not have to fix the number of agents in a model, instead, for some properties that are invariants for MAS of various sizes in the number of agents, we can bound that number by a scope (such as in [Dewayne’s ref]) but enhanced with invariants on agents’ data structures). Such properties can check correctness of multiple simultaneous negotiations or self-organization mechanisms of multi-agent systems. It is not necessary to consider the problem of identifying agents that should communicate for the purpose of checking invariants of multiple simultaneous negotiations. That problem can be dealt with separately allowing for a smaller specification which is more likely to allow feasible analysis.

It would be interesting to see whether CSP-based models and tools (FDR) or B CSP models would be useful for checking properties of negotiation in MAS systems with explicit communication.

A larger future goal is to expand the Alloy analyzer’s language with operators suitable for checking properties specific to multi-agent systems such as belief, desire and intention. Thus we can generalize some of the Alloy model “design patterns” for specific needs of MAS and provide a set of abstractions more useful for checking MAS-specific properties. Having done that we can conduct comparative analysis of usefulness, applicability and current performance of the Alloy approach and model checking approaches to verifying MAS.

Acknowledgments

We would like to express our deep gratitude to Prof. Victor Lesser (UMass, Amherst) for his help with the negotiation protocol example implemented in the multi-agent system simulator and helpful comments.

References

1. W. Adjie-Winoto, E. Schwartz, H. Balakrishnan, and J. Lilley. The design and implementation of an intentional naming system. In *17th ACM Symposium on Operating Systems Principles (SOSP)*, Kiawah Island, Dec. 1999.
2. N. Alechina and B. Logan. Verifying bounds on deliberation time in multi-agent systems. In *EUMAS*, pages 25–34, 2005.
3. D. Box. *Essential COM*. Addison Wesley, 1998.
4. K. Decker. TAEMS: A Framework for Environment Centered Analysis & Design of Coordination Mechanisms. In *Foundations of Distributed Artificial Intelligence, Chapter 16*, pages 429–448. G. O’Hare and N. Jennings (eds.), Wiley Inter-Science, January 1996.
5. G. Dennis. TSAFE: Building a trusted computing base for air traffic control software. Master’s thesis, Massachusetts Institute of Technology, 2003.
6. G. Gans, M. Jarke, G. Lakemeyer, and T. Vits. SNet: A modeling and simulation environment for agent networks based on i* and ConGolog. In *CAiSE02*, pages 328–343, Canada, Toronto, May 2002. Springer.
7. D. Jackson. *Software Abstractions: Logic, Language and Analysis*. The MIT Press, Cambridge, MA, 2006.
8. D. Jackson and M. Jackson. *Separating Concerns in Requirements Analysis: An Example*, chapter Rigorous development of complex fault tolerant systems. Springer-Verlag. (To appear).
9. D. Jackson, I. Schechter, and I. Shlyakhter. ALCOA: The Alloy constraint analyzer. In *22nd International Conference on Software Engineering (ICSE)*, Limerick, Ireland, June 2000.
10. M. Kacprzak, A. Lomuscio, and W. Penczek. Verification of multiagent systems via unbounded model checking. In *AAMAS*, pages 638–645, 2004.
11. V. Lesser. Reflections on the Nature of Multi-Agent Coordination and Its Implications for an Agent Architecture. *Autonomous Agents and Multi-Agent Systems*, 1:89–111, January 1998.
12. V. Lesser, K. Decker, T. Wagner, N. Carver, A. Garvey, B. Horling, D. Neiman, R. Podorozhny, M. N. Prasad, A. Raja, R. Vincent, P. Xuan, and X. Zhang. Evolution of the GPGP/TAEMS Domain-Independent Coordination Framework. *Autonomous Agents and Multi-Agent Systems*, 9(1):87–143, July 2004.
13. W. van der Hoek and M. Wooldridge. Model checking knowledge and time. *Proc. of the 9th Int. SPIN Workshop*, 2318 of LNCS:95–111, 2004.
14. T. Wagner and V. Lesser. Design-to-Criteria Scheduling: Real-Time Agent Control. *Proceedings of AAAI 2000 Spring Symposium on Real-Time Autonomous Systems*, pages 89–96, March 2000.
15. G. Weiss, editor. *Multiagent systems : a modern approach to distributed artificial intelligence*. MIT Press, Cambridge, Mass., 1999.
16. M. Wooldridge, M. Fisher, M.-P. Huget, and S. Parsons. Model checking multi-agent systems with mable. In *AAMAS*, pages 952–959, 2002.
17. X. Zhang, R. M. Podorozhny, and V. Lesser. Cooperative, MultiStep Negotiation Over a Multi-Dimensional Utility Function. In *Proceedings of the IASTED International Conference on Artificial Intelligence and Soft Computing (ASC 2000)*, pages 136–142, 2000.