



A Real-Time Negotiation Model and A Multi-Agent Sensor Network Implementation

9 LEEN-KIAT SOH

lksoh@cse.unl.edu

5 *Department of Computer Science and Engineering, University of Nebraska, 256 Avery Hall, Lincoln, NE*
6 *68588-0115, USA*

7
8
12 COSTAS TSATSOULIS

Information and Telecommunication Technology Center (ITTC), Department of Electrical Engineering and
Computer Science, University of Kansas, Nichols Hall, 2335 Irving Hill Road, Lawrence, KS 66045-
7612, USA

16 **Abstract.** This paper describes a negotiation model that incorporates real-time issues for autonomous
17 agents. This model consists of two important ideas: a real-time logical negotiation protocol and a case-
18 based negotiation model. The protocol integrates a real-time Belief-Desire-Intention (BDI) model, a
19 temporal logic model, and communicative acts for negotiation. This protocol explicitly defines the logical
20 and temporal relationships of different knowledge states, facilitating real-time designs such as multi-
21 threaded processing, state profiling and updating, and a set of real-time enabling functional predicates in
22 our implementation. To further support the protocol, we use a case-based reasoning model for negotiation
23 strategy selection. An agent learns from its past experience by deriving a negotiation strategy from the
24 most similar and useful case to its current situation. Guided by the strategy, the agent negotiates with its
25 partners using an argumentation-based negotiation protocol. The model is time and situation aware such
26 that each agent changes its negotiation behavior according to the progress and status of the ongoing
27 negotiation and its current agent profile. We apply the negotiation model to a resource allocation problem
28 and obtain promising results.

29 **Keywords:** real-time negotiation, temporal logic, argumentative negotiation protocol, case-based rea-
30 soning, multi-agent system.

31 1. Introduction

33 While negotiation has been used in the past in problem solving in multi-agent sys-
34 tems, in our work we focus on negotiations and activities that must occur in real
35 time. The introduction of hard real-time constraints in negotiation and action exe-
36 cution complicates the problem greatly, and existing negotiation protocols cannot
37 provide an adequate solution. In this paper we describe a real-time negotiation
38 model that is used in resource allocation problems. As an example domain we use
39 multi-sensor target tracking, where each agent controls a sensor with a limited
40 sensing coverage area. As a target moves across space, agents have to cooperate to
41 track it. Each agent (together with the sensor it controls) consumes resources such as
42 time, battery power, bandwidth of the communication channel, and some percentage
43 of the CPU where the agent resides, and each agent strives to manage and utilize

	Journal : AGNT	Dispatch : 10-3-2005	Pages : 56
	PIPS No. : N000000539	<input type="checkbox"/> LE	<input type="checkbox"/> TYPESET
	MS Code : N000000539	<input checked="" type="checkbox"/> CP	<input checked="" type="checkbox"/> DISK

44 its resources efficiently and effectively. This motivates the agents to share their
45 knowledge about a problem based on their viewpoints in their effort of arriving at a
46 solution. The problem of global resource allocation becomes a problem of locally
47 negotiated compromises and local constraint satisfaction.

48 We propose a logical negotiation protocol that incorporates a real-time Belief-
49 Desire-Intention (BDI) model [43, 44] to dictate the rules of encounter among our
50 autonomous agents. A feature of our problem involves generating a “good-enough,
51 soon-enough” solution¹ to resource allocation. Since time is critical – for example, to
52 make a good triangulation for the location of a target, three different sensors have to
53 make a measurement within 2 seconds of each other – agents use time to guide their
54 negotiation behavior. We base our temporal model on [1] in which logical events or
55 propositions can be ordered consistently along a timeline and durations of events or
56 propositions holding true can be derived from their relationships with others. This
57 temporal logic allows us to define explicitly the transition of a BDI state to another,
58 including causality and co-existence. Equipped with the definition of time, we are
59 able to model our negotiation activity with more accuracy, spelling out how and
60 when a state changes and how and when it changes with other states, such as state s_1
61 triggers state s_2 , s_1 has to occur before s_2 , s_1 must hold true for some time during
62 which s_2 must hold true as well, and so on. Therefore, states may change their truth
63 values during a reasoning process as long as the states are needed to hold constant
64 during that time period do – releasing other states to be updated or changed by other
65 events or states. This is critical in our agent design as each agent is multi-threaded,
66 meaning that several threads may attempt to access and modify the same variable
67 (state) at the same time. To maintain data integrity, when a thread is accessing or
68 modifying a variable, other threads will be blocked, and this is the common
69 approach. However, software designers must find out how and for how long the
70 threads will be blocked; and this information is very important in a real-time system
71 like ours. With temporal logic, we know for how long threads will be blocked and
72 how these threads will be blocked awaiting which variables (states) to become
73 accessible. This allows us to fine-tune the system to increase the efficiency of the
74 negotiation process. Further, by incorporating temporal logic and BDI models into
75 our negotiation protocol, we can improve the temporal and logical structure that
76 facilitates the completion of a negotiation. We know what states are needed (and
77 when they are needed) for a negotiation to logically complete, and we also can model
78 the time distribution or usage needed for each step of the negotiation to complete
79 within certain time constraints.

80 Note that temporal components have been in place in the BDI model [9, 43] to
81 determine how the three modalities are related over time. For example, taking time
82 into consideration allows one to have persistent intentions, inevitable outcomes, and
83 so on. In our model, we use temporal logic to control the stability of a state, which in
84 turn facilitates our multi-threaded solution.

85 We further define two sets of communicative acts – one for handling incoming
86 messages and one for handling outgoing messages. A communicative act that han-
87 dles incoming messages is a function that turns an event (the arrival of a message) to
88 a set of BDI states. The function parses the incoming message and generates states
89 that are necessary for the agent reasoning during a negotiation process. Similarly, a

90 communicative act that handles outgoing messages is a function that composes a
91 message based on the agent's current BDI states and sends it out via a communi-
92 cation channel. We qualify these acts with temporal logic and incorporate them into
93 the negotiation protocol. In addition to the communicative acts, we utilize a suite of
94 real-time enabling functional predicates to assist agents in negotiations. These
95 predicates are events that take time to execute and they also generate or modify
96 states.

97 Since each agent is autonomous and reacts to its environment, each has its own
98 knowledge base and its own monitoring of the world events, including its sensor, its
99 neighbors, and the targets. To increase the fault tolerance of the multi-agent system,
100 each agent is responsible only for the resources it controls (in our example domain,
101 its sensor and associated components), and it controls the minimal set of resources it
102 requires to achieve its task. In our work agents have minimal knowledge and
103 information – they know how to perform their tasks, have a local, limited view of the
104 world provided to them by the equipment they control, and know of the existence of
105 other similar agents, but they do not have an explicit view of the information of the
106 other agents. There is some implicit knowledge, namely that the other agents control
107 a set of resources, that they are willing to cooperate, that they are capable of
108 negotiation for resource sharing, and that they are truthful. To establish a common
109 reasoning basis during a collaborative effort, an agent is required to communicate to
110 its potential partner why it needs to share the resources controlled by the partner.
111 This knowledge exchange can be done via different mechanisms such as a blackboard
112 where agents post information on a common site, or auctions where a contractor-
113 agent oversees the message passing among contractee-agents, or through agent-
114 based negotiations where agents exchange information directly. In our approach, we
115 use negotiations motivated by a global goal – to track as many targets as accurately
116 as possible – guided by a set of local optimization criteria that affect the strategies.
117 During a negotiation agents exchange information of their individual viewpoints of
118 the current (and relevant) world situation. In this manner, the agents are able to
119 argue and attempt to persuade each other explicitly, resulting in efficient knowledge
120 transfer. We thus do away with a centralized information facility that requires
121 constant updates and polling from agents, and, instead, knowledge is exchanged
122 when necessary resulting in less communication traffic. Knowledge inconsistencies
123 are resolved in a task-driven manner, making knowledge management easier.

124 One important part of the negotiation process is the determination of the nego-
125 tiation strategy based on the current task description. To do so, we use a model
126 derived from case-based reasoning (CBR) [23] that is time-constrained and that
127 retrieves the most similar cases, selects the best case based on utility theory, adapts
128 the case to the current situation, and then uses the case's negotiation strategy to
129 perform negotiations. The CBR approach limits the time needed to decide on a
130 negotiation strategy – selection (through retrieval) and generation (through adap-
131 tation) of a situation-appropriate strategy – and enables the agent to learn auton-
132 omously and adapt itself to different scenarios in the domain. We will only briefly
133 discuss our CBR approach in this paper.

134 In the following, we first describe briefly the characteristics of our agents and
135 multi-agent system and its real-time constraints. In Section 3, we describe in detail

136 the logical model of our negotiation protocol. Then we describe our implementation
137 of the negotiation protocol, including a suite of real-time enabling functional
138 predicates, in Section 4. Next, we discuss some results of our negotiation model
139 applied to the multi-sensor target-tracking problem. Then we study some related
140 work in agent-based negotiations. Finally, we conclude with some suggestions for
141 further work in the area.

2. Agent characteristics

143 Each agent has the following characteristics:

- 144 (1) Autonomous – Each agent runs without interaction with human users. It
145 maintains its own knowledge base, makes its own decisions, and interacts with its
146 sensor, neighbors and environment.
- 147 (2) Rational – Each agent is rational in that it knows what its goals are and can
148 reason and choose from a set of options and make an advantageous decision
149 to achieve its goal [57].
- 150 (3) Communicative – Each agent is able to communicate with others, by initiating
151 and responding to messages, and carrying out conversations.
- 152 (4) Reflective (or Aware) – According to [8], a reflective agent reasons based on its
153 own observations, its own information state and assumptions, its communication
154 with another agent and another agent’s reasoning, and its own control or reason-
155 ing and actions. By being reflective, each agent is time aware and situationally
156 aware. When an agent is time aware, it observes time in its decision making and
157 actions. Its reasoning takes time into account, and thus, the outcome of a reason-
158 ing process is partially dependent on time. When an agent is situationally
159 aware, it observes its current situation, the situation of its neighbors, and that of
160 the world and makes decisions based on these observations. In general, an agent
161 that is situationally aware observes the resources that it shares with other agents,
162 its current tasks, messages, profiles and actions of its neighbors, and the external
163 changes in the environment. In this paper, we require a stronger level of situa-
164 tional awareness. An agent also observes its own resources that *sustain the being*
165 of the agent. For a hardware agent, these resources may be the battery power, the
166 radio-frequency (RF) links, etc. For a software agent, these resources may be
167 CPU, RAM, disk space, communication channels, etc. Note that, for example, in
168 [48], a *bounded rationality* model is used where each agent has to pay for the
169 computational resources (CPU cycles) that it uses for deliberation, assuming that
170 the resources are available. In our model, however, we require an agent to be
171 aware of whether the resources are available before even starting a negotiation.
- 172 (5) Honest – Each agent does not knowingly lie or intentionally give false infor-
173 mation. This characteristic is also known as veracity [16].
- 174 (6) Adaptive – Each agent is able to adapt to changes in the environment and learns
175 to perform a task better, not only reactively but also from its past experience.
- 176 (7) Cooperative – Each agent is motivated to cooperate if possible with its neighbors
177 to achieve global goals while satisfying local constraints.

178 Generally, the agents in a multi-agent system may be controlling different
179 resources and use different reasoning and negotiation techniques. In our approach,
180 we require (1) that all agents be capable of negotiation in which they share a com-
181 mon vocabulary that enables message understanding, and (2) that each agent knows
182 what resources may be used or controlled by a non-empty subset of the other agents
183 in the environment so that it can determine whom to negotiate with. In our
184 particular domain of application, each agent controls the same resources, since each
185 one controls the same type of sensor. Also, each agent uses the same negotiation
186 methodology based on CBR, but the individual case bases differ.

187 Formally, our multi-agent system architecture is defined as follows. In our
188 study, a multi-agent system consists of a finite set of autonomous agents, denoted
189 by Ω . Suppose that we define a neighborhood of an agent α_i , Ψ_{α_i} such that,
190 $\Psi_{\alpha_i} \subseteq \Omega$, $\Psi_{\alpha_i} \neq \emptyset$ and that the agent α_i knows about all other agents in the
191 neighborhood.

192 For the given agent α_i and any agent from its neighborhood, Ψ_{α_i} , $\lambda(\alpha_i, \alpha_j)$, where
193 $\lambda(a,b)$ means agent a knows about the existence of agent b and can communicate with
194 agent b . A neighborhood is different from a team as defined in [55], which is task-driven
195 and formed among a set of agents to accomplish a task. A neighborhood is a subset of
196 agents of the multi-agent system that *could* form a team. In our particular domain of
197 application (multi-sensor target tracking), a neighborhood consists of a set of agents
198 that control sensors that are physically close and whose sensing beams overlap. So, in
199 our multi-agent system Ω , there is a set of neighborhoods, and each neighborhood can
200 form any number of teams. Neighborhoods do not necessarily have the same number of
201 members, and neighborhoods may share members.

202 When a target is sensed, an agent tracks the target, refers to its neighborhood
203 information, and dynamically forms a *negotiation coalition*, that is, a subgroup of its
204 neighborhood agents with which it *may* negotiate to request resources to assist it in
205 its task. For example, when an agent detects a target in its sensing area, the agent
206 immediately obtains an estimate on the position and velocity of the target. It then
207 projects the future positions of the target and identifies the neighbors whose sensors
208 are able to cover the target moving in the projected path. These are agents that
209 control resources (i.e. sensor beams) that it needs to track the target, and these are
210 the agents that will be part of the negotiation coalition.

211 Figure 1 shows the agent architecture. The agent consists of several threads:
212 communication, execution, core, and multiple negotiation threads, allowing parallel
213 processing of multiple tasks. The core thread maintains the lifecycle of the agent,
214 linking all managers together. The execution manager manages the execution thread,
215 scheduling radar-related tasks in a job queue and actuating them to affect the agent's
216 environment. The communication manager manages the communication thread,
217 sending and receiving messages through sockets, and storing messages in a mailbox
218 that is accessed by other threads. The negotiation manager manages the negotiation
219 threads, spawning, activating, and updating them. The core thread consists of several
220 managers. The profile manager maintains the status of the agent, radar, its neigh-
221 bors, and its environment. The CBR manager maintains the casebase and performs
222 CBR such as retrieval, adaptation, and learning. The Radar manager is a domain-
223 specific module that hosts a geometric model of the tracked target within the radar

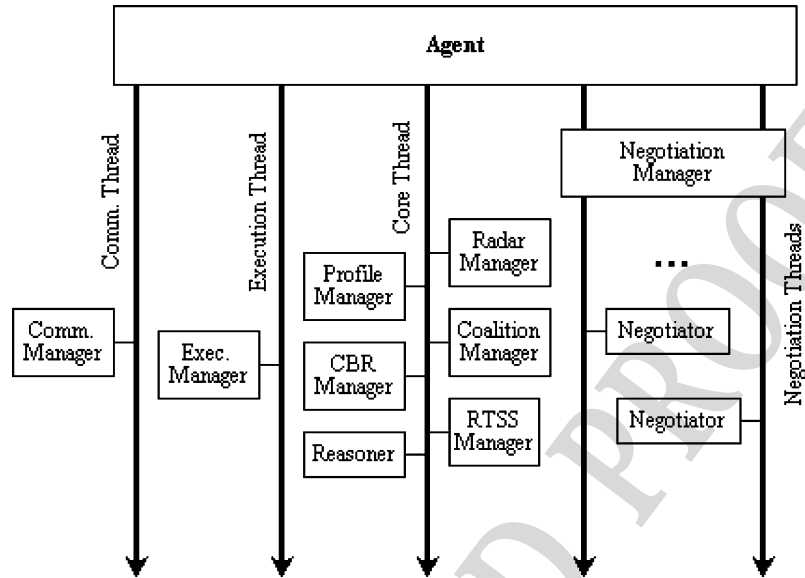


Figure 1. Our multi-threaded agent architecture.

224 coverage of an agent and its neighbors. When the agent encounters an event, the
 225 coalition manager forms a coalition, oversees the negotiations to recruit members,
 226 and confirms or discards a coalition depending on the status of the event and the
 227 outcome of the negotiations. The RTSS manager interacts with a real-time sched-
 228 uling service (RTSS) at the kernel level, allowing the agent to monitor and request
 229 for its CPU allocation and usage. The reasoner is responsible for the general thinking
 230 tasks of the agent. For each negotiation thread, a negotiator shoulders the respon-
 231 sibility of carrying out communicative acts, basing its negotiation strategy on the
 232 information provided by the profile and CBR managers, monitoring time using the
 233 RTSS manager, receiving coalition information from the coalition manager,
 234 retrieving negotiation messages by checking the mailbox maintained by the com-
 235 munication manager, and, most importantly, executing the real-time argumentative
 236 negotiation, the main focus of this paper. Readers are referred to [50] for a detailed
 237 description of the agent architecture.² Note that the multi-threaded design of an
 238 agent infrastructure has been implemented elsewhere in MAS research [41, 49, 54].

3. A logical protocol for real-time argumentative agent negotiations

240 In this section, we describe the logical protocol for our real-time argumentative agent
 241 negotiations. Argumentative negotiations differ from traditional negotiations
 242 because the agents conducting the former negotiate about *why* one of the agents
 243 needs to perform a certain task in addition to what the task is. Therefore, our work is
 244 similar to [40]. However, we assume that agents have the same inference rules.
 245 Parsons' work assumes that agents may have different ones, and thus his argu-
 246 mentation protocol requires agents to exchange inference rules as well. Moreover, we
 247 incorporate real-time issues into our design guidelines.

248 3.1. *Real-time negotiation design guidelines*

249 In this section, we outline several design objectives for our argumentative model.
250 Note that in our domain, each agent stores a local viewpoint of the world and
251 information is distributed among the agents. For an agent to convince another it
252 needs to exchange information with the partner during negotiations. In a real-time,
253 dynamic, and distributed environment, the agents do not have sufficient time to
254 perform detailed negotiations, leading to the following design objectives. We believe
255 that these guidelines are general and applicable to real-time argumentative inter-
256 agent negotiations, where time is critical.

257 We also offer a brief discussion of their operationalization in our system:

258 *Design objective 1* A negotiation should be bounded by time.

259 In a real-time environment a negotiation should complete (with success or failure)
260 within a predefined window of hard time, requiring time-bounded rationality from
261 the agents. In our model, before a negotiation begins, the agent determines the time
262 allotted for the negotiation based on the speed of the target and the coverage time
263 available. If a negotiation has run beyond the allotted time window, then the agent
264 terminates the negotiation since any deal reached after that and the subsequent task
265 performed will no longer be useful.³ Also, an agent should be able to react to the
266 pace of its negotiation. If the pace of the negotiation is slow, then, an agent counter-
267 offers to narrow the constraints further so that both agents can come to an agree-
268 ment or disagreement more quickly.

269 *Design objective 2* Each step of the negotiation should be fast.

270 A negotiation process consists of multiple actions (steps): parsing of and
271 responding to messages, evaluation of offers, submission of offers, etc. Each one of
272 these steps should be performed efficiently and quickly to achieve real-time behavior.

273 *Design objective 3* A negotiation should be kept short – the number of iterations
274 minimized.

275 If a negotiation takes more iterations, then it incurs more processing time, more
276 communication latency, more bandwidth usage, and a higher opportunity for noise
277 to distort the messages. Towards this end, we use ranked arguments such that the
278 most important argument is sent over to the responding agent first, hoping to
279 convince it sooner.

280 *Design objective 4* A negotiation-related message should be kept short.

281 This is meant to reduce loss and to improve communication speed.

282 *Design objective 5* An agent should be able to abort a negotiation unilaterally when
283 communication fails.

284 In a real-time environment, the communication channels may be jammed or
285 simply congested. An agent should be able to recognize when its partner's channel is
286 no longer receptive, terminate the negotiation immediately, and move on to the next
287 neighbor or task.

288 *Design objective 6* A negotiation strategy should be determined quickly and
 289 reflectively.

290 A negotiation strategy guides how an agent should conduct its negotiation. It
 291 consists of a set of parameters such as the number of negotiation iterations, time
 292 allotted, ranking of arguments, etc. These values should be derived quickly such that
 293 by the time the negotiation strategy is obtained, it still is reflective of the agent's
 294 status and the target. Slow determination of the negotiation strategy renders the
 295 negotiation behavior less situated to the agent's current status.

296 *Design objective 7* A negotiation should only be *efficiently reflective*.

297 During a negotiation, an agent's status changes constantly. If we check the agent's
 298 state at each time step⁴ and dynamically fine-tune the negotiation parameters, then
 299 the agent is fully reflective. However, both the monitoring of the agent attributes that
 300 might affect the negotiation, and the response to changing attributes, are time
 301 consuming and may slow down the negotiation. Hence, there must be a delicate
 302 trade-off and a negotiation should only be *efficiently reflective*, that is, consider only
 303 important changes to attributes and ignore others. In our model, an agent reflectively
 304 considers only time and the need for a negotiation.

305 3.2. Negotiation protocol

306 Figure 2 shows our negotiation protocol in a state diagram between two agents: *a*
 307 and *b*. State 0 is the initial state, the double-circle. State 1 is the first *handshake* state,
 308 indicating whether the initiated negotiation will be entertained. State 4 is the initi-
 309 ating state while state 5 is the responding state. The initiating state is where the
 310 initiating agent, *a*, returns to, basically the processing loop of the negotiator module.
 311 The responding state is where the responding agent, *b*, returns to, respectively. Agent
 312 *a* initiates a negotiation request to *b* by sending an INITIATE message (initiate
 313 (*a,b*)), the state transitions to 1. At this juncture, there are four possible scenarios.
 314 First, agent *b* may outright refuse to negotiate by sending a NO_GO message (no_go
 315 (*b,a*)). This results in a final state of failure (state 2, rejected). Second, agent *b* may
 316 outright agree to the requested task by sending an AGREE message (agree (*b,a*)).
 317 This results in a final state of success (state 3). Third, agent *b* may decide to entertain
 318 the negotiation request and thus sends back a RESPOND message (respond (*b,a*)).
 319 This transitions the state to 4. Fourth, there may be no response from agent *b*. Thus
 320 agent *a*, after waiting for some time, has no choice but to declare a no response
 321 (no_response (*a*)) and moves to a state of failure (state 8, channel_jammed).

322 When the agents move to state 4, the argumentative negotiation begins and iterates
 323 between states 4 and 5 until one side opts out or both sides opt out or both sides agree.
 324 During the negotiation, (1) agent *a* provides information or arguments to *b* by sending
 325 INFO messages (info(*a,b*)), (2) agent *b* demands information or arguments from *a* by
 326 sending MORE_INFO messages (more_info(*b,a*)), (3) if agent *a* runs out of argu-
 327 ments, it sends a INFO_NULL message to *b* (info_null(*a,b*)), (4) if agent *b* runs out of
 328 patience, it counter-proposes by sending a COUNTER message to *a* (counter(*b,a*)),
 329 and (5) agent *a* can agree to the counter offer (agree(*a,b*)) and move to the state of

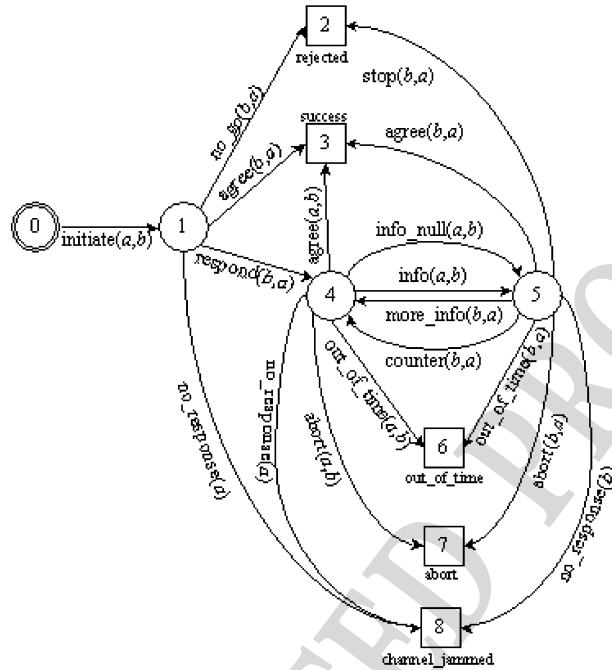


Figure 2. Our negotiation protocol. Squares are final states. The double-circle is the initial state. This state diagram shows how two negotiating agents interact based on responses to offers, and the real-time constraints of the environment.

330 success (state 3), or provide more information ($info(a,b)$) as requested, or provide no
 331 information ($info_null(a,b)$) if it has time to do so, hoping that agent b might come up
 332 with a better offer, or simply disagrees ($abort(a,b)$). Thus, an initiating agent will
 333 always negotiate until it has run out of time or when the responding agent opts out.
 334 However, an initiating agent may abort a negotiation, and this is where the conditions
 335 come into play. If the agent realizes that it has already obtained what it wants from
 336 other negotiations happening in parallel, then it aborts the current negotiation; or if
 337 the agent realizes that it no longer cares about the current negotiation, then it aborts.
 338 These conditions are based on desires and intentions, which in turns are based on
 339 beliefs of the agent. When an agent runs out of time, it issues an OUT_OF_TIME
 340 message to the other agent and quits the negotiation with a failure (state 6, out_
 341 of_time). When an agent aborts, it issues an ABORT message to the other agent and
 342 quits the negotiation with a failure (state 7, abort). Finally, whenever an agent does not
 343 hear from the other agent within an allocated time period, it assumes that the com-
 344 munication channel has been jammed or congested and quits with a failure (state 8,
 345 channel_jammed). Note that we distinguish NO_GO, STOP, OUT_OF_TIME, and
 346 ABORT in the above protocol. With the above different end states, agent a can
 347 determine whether the negotiation has failed because it has exhausted all its arguments
 348 (STOP) or otherwise and subsequently learn from the failure.

349 Note that our work does not attempt to address the issues with AI in Law,
 350 especially in the dialectics in legal arguments (e.g., [21, 42]). Our “persuasion” is a

351 negotiation protocol for information exchange between two agents to reach a deal in
 352 which the initiating agent provides arguments to support its request in order to
 353 convince the responding agent to agree to the request. Hence, we do not look into
 354 the maneuvers found in legal arguments such as defenses, proposals, positions, and
 355 attacks for our design.

356 3.3. Logical framework

357 We use an expanded multi-context BDI agents framework of [36, 40] to describe the
 358 logical framework of our negotiation protocol. As presented in [44], there are three
 359 modalities. First, beliefs (B) represent the states of the environment and the agent.
 360 There are also belief states that arise during negotiations, as agents learn each other's
 361 beliefs and intentions. Second, desires (D) represent the motivations of the agent.
 362 Third, intentions (I) represent the goals that the agent wants to achieve. We also
 363 assume the following axioms [40]:

- 364 (1) $B : B(p \rightarrow q) \rightarrow (B(p) \rightarrow B(q))$
- 365 (2) $B : B(p) \rightarrow \neg B(\neg p)$
- 366 (3) $B : B(p) \rightarrow B(B(p))$
- 367 (4) $B : \neg B(p) \rightarrow B(\neg B(p))$

368 Axiom 1 states that if one believes that p implies q , then if one believes p , then one
 369 believes q as a consequence. Axiom 2 states that if one believes p , then one does not
 370 believe the negation of p . Axiom 3 states that if one believes p , then one also believes
 371 that it believes p . This ensures that an agent *knows* and believes in what it believes.
 372 Similarly, Axiom 4 states that if one does not believe p , then one believes that it does
 373 not believe p . Similar axioms for desires and intentions are:

- 374 (5) $D : D(p \rightarrow q) \rightarrow (D(p) \rightarrow D(q))$
- 375 (6) $D : D(p) \rightarrow \neg D(\neg p)$
- 376 (7) $I : I(p \rightarrow q) \rightarrow (I(p) \rightarrow I(q))$
- 377 (8) $I : I(p) \rightarrow \neg I(\neg p)$

378 We adopt partially the strong realist BDI agent model of [40] such that (1) an
 379 agent only intends to do what it desires, and (2) an agent only desires what it
 380 believes. However, we do not adopt other rules presented in [40], regarding the
 381 communication unit, because in our case (1) an intention for performing a task does
 382 not necessarily imply a communication of the performance of the task – in our
 383 domain, a task may sometimes be performed by the agent itself, and (2) we do not
 384 require an agent to report a completion of a task to another agent – because in our
 385 current design, an agent assumes that if another agent agrees to perform a task, it
 386 knows that that agent will try its best to perform and complete the task, and that
 387 whether it believes the task has been performed is no longer important. In the future,
 388 however, we plan to include a monitoring mechanism in an agent to enrich our real-
 389 time modeling of events and agent behavior in making better decisions. At that time,
 390 an agent would have to care whether a task has been performed successfully to
 391 update its own belief states, but would still not depend on the communicated
 392 information for the update.

393 We incorporate temporal logic into our protocol to explicitly define the various
394 belief, desire, and intention states of an agent and when they are true. This is key to
395 the real-time implementation of the protocol. To satisfy real-time constraints, our
396 agent, as shown in Figure 1, consists of multiple concurrent processing threads. Each
397 thread carries out a set of tasks, and these tasks access and modify the same states at
398 different times. Some states must not be modified before certain actions have been
399 carried out; some states should not be accessed before certain results have been
400 obtained. To implement such a *state synchronization* across multiple concurrent
401 processes, we use temporal logic to define when a state must be true and the duration
402 for that state to stay true. Without the temporal logic component, it would have been
403 close to intractable to manage the inter-thread, real-time activities.

404 As previously mentioned, Cohen and Levesque [9, 10] and Rao and Georgeff [43,
405 45] have incorporated temporal components into the BDI model. In [9], intentions
406 are defined in terms of temporal sequences of an agent's beliefs and goals. Each
407 possible world extendable from a current state at a particular time point is a time line
408 representing a sequence of events. As such, the inter-modal relationships are stronger
409 than those in [43]. For example, an agent fanatically committed to its intentions will
410 maintain its goals until either they are believed to be achieved or believed to be
411 unachievable. Thus, intentions are seen as a *special* class of desires. Rao and
412 Georgeff [43], on the other hand, present an alternative possible-worlds formalism
413 for BDI-architectures. Instead of a time line, they choose to model the world using a
414 temporal structure with a branching time future and a single past, called a time tree,
415 where a particular time point in a particular world is called a situation. There are
416 three crucial elements to the formalism. First, intentions are on a par with beliefs and
417 goals. This allows them to define different strategies of commitment and to model a
418 wide variety of agents such as blinded, single-minded, and open-minded agents.
419 Second, they distinguish between the choice an agent has over the actions it can
420 perform and the possibilities of different outcomes of an action, factoring in the
421 uncertainty that the environment brings into the determination of the outcomes.
422 Third, they specify an interrelationship between beliefs, goals, and intentions that
423 allows them to avoid problems such as commitment to unwanted side effects. Fur-
424 ther, there has been work in for branching-time temporal logic such as the com-
425 putation tree logic (CTL) to carry out temporal resolution [7, 45].

426 In our model, the incorporation of the temporal elements is different. In [9], the
427 temporal component was used to define intentions through commitment and per-
428 sistence, derived from beliefs and goals. In [43], it was used to order possible worlds
429 from a situation in both the time and space dimensions. Each time tree denotes the
430 optional courses of events choosable by an agent in a particular world. For example,
431 an agent has a belief ϕ , denoted $B(\phi)$, at time point t if and only if ϕ is true in all the
432 belief-accessible worlds of the agent at time t . In our model, however, we use the
433 temporal component to define the temporal duration that a state needs to be stable,
434 or needs to occur in order for another state to take place. Thus, our motivation is to
435 help design and implement agents that are multi-threaded and multi-tasking. We do
436 not use temporal logic to define the BDI modalities.

437 In the following subsections, we formalize our theory of the actions depicted in the
438 state diagram of Figure 2.

439 **3.3.1. Temporal logic** To incorporate real-time concerns into our logical negoti-
 440 ation protocol, we use several interval relationships outlined in [1, 3]. Each interval t
 441 has a start time, t_s , and a finish time, t_f , and its duration is $t_f - t_s$. If $t_f - t_s$ equals the
 442 smallest amount within the resolution of the domain problem, then the interval
 443 becomes a *moment* or a *point*. There are seven basic relations between temporal
 444 intervals:

- 445 (1) *Before*(i, j) where interval i ends before interval j .
- 446 (2) *Meets*(i, j) where as soon as i finishes, interval j starts, i.e., the two intervals are
 447 consecutive.
- 448 (3) *Overlaps*(i, j) where a portion of interval i overlaps a portion of interval j in time
 449 and i starts before j and i ends before j .
- 450 (4) *Starts*(i, j) where interval i starts at the same time as interval j but interval i has a
 451 shorter duration.
- 452 (5) *Finishes*(i, j) where interval i finishes at the same time as interval j but interval i
 453 has a shorter duration.
- 454 (6) *During*(i, j) where interval i starts after interval j and interval i ends before
 455 interval j .
- 456 (7) *Equals*(i, j) where both intervals have the same durations and start and end at the
 457 same times. We also adopt the homogeneity axiom schema such that a propo-
 458 sition is homogeneous if and only if when it holds over an interval t , it also holds
 459 over any sub-interval within t . Within the framework of our negotiation proto-
 460 col, the BDI states are all homogeneous propositions, hence the use of strong
 461 negation \neg . The predicates such as the communicative acts are anti-homogeneous
 462 [2] since, for example, the action of composing and sending a message is a process
 463 that does not generally hold unless completed in the end.

464 We also introduce a notational convenience $[e]$ as the maximum interval of an
 465 event/action, or, in the case of homogeneous positions, as the interval of a propo-
 466 sition holding true.

467 **3.3.2 Communicative acts** In our protocol, we have two sets of corresponding
 468 communicative acts. One is for receiving and parsing an incoming message; the other
 469 for composing and sending an outgoing message. For example, a communicative act
 470 that composes a negotiation request and initiates contact with a potential negotia-
 471 tion partner is of the following form:

$$C_{out} : initiate(i, r, Do(r, \rho), t)$$

473 where the predicate *initiate* is the communicative act (composing and sending), i is
 474 the initiating agent, r is the responding agent, $Do(r, \rho)$ is the requested task, i.e., “ r do
 475 task, ρ ” and t is the time taken for the communicative act to start and finish.

476 A communicative act is an event that performs a set of tasks consecutively such that
 477 the sum of the durations of the tasks is the duration of the communicative act. A
 478 communicative act may generate new propositions, may cause a new state externally to
 479 another agent, and may be terminated by another proposition or event. For example, at
 480 the end of the interval t , if the communication is successful, then the responding agent
 481 will receive the request, $C_{in} : initiate(i, r, Do(r, \rho), k)$, where, in theory, *Meets*(t, k), and

482 in practice (due to communication latency), *Before* (t,k). On the other hand, if the
 483 communication is unsuccessful (e.g., due to communication channel being jammed),
 484 then the predicate *initiate* of the initiating agent will be terminated by a terminator [56],
 485 as will be discussed further in Section 3.3.3. To simplify our discussions, we use
 486 $C_{in} : f(\langle sender \rangle, \langle receiver \rangle, \langle request \rangle, t)$ and $C_{out} : f(\langle sender \rangle, \langle receiver \rangle, \langle request \rangle, t)$ to
 487 differentiate between incoming and outgoing message handling, where f is one of the
 488 acts defined in our protocol. Currently, we have the following communicative acts, as
 489 depicted in Figure 1: *initiate*, *respond*, *no_go*, *agree*, *abort*, *out_of_time*, *counter*, *mor-*
 490 *e_info*, *info*, *info_null*, and *stop*, for a total of 22. An example of a message is
 491 $C_{in} : agree(A1, A2, Do(A1, \rho), T_1)$, which means that $A2$ receives as an input message
 492 from $A1$ an agreement that states that $A1$ will perform the task ρ , and the receipt and
 493 processing of the message occurs at the time interval T_1 .

494 One of the objects generated by a C_{out} communicative act is the message. Our
 495 message syntax is $msg(\langle sender \rangle, \langle receiver \rangle, \langle type \rangle, \langle request \rangle, \langle contents \rangle)$ where the
 496 type of the message denotes one of the communicative acts and the contents consist
 497 of whatever pertinent to the request. In the following discussion, we often mention
 498 the messages in the same breath as the communicative acts and use them inter-
 499 changeably.

500 **3.3.3. Terminator** Vere [56] described a proposition as bounded by its holding
 501 true over a time interval, with a limited life span terminated by later contradictory
 502 assertions. Under Vere's definition, an assertion T is a terminator for an assertion A
 503 if and only if: (1) A and T are contradictory, (2) T follows A in time, and (3) no
 504 assertion T' exists satisfying the first two conditions such that T follows T' in time.
 505 Note that the third condition ensures that there are no conditions between A and T
 506 that might pre-empt T as a terminator. Basically, $Meets([A],[T])$.

507 Take our communicative acts for example. Suppose an initiating agent performs
 508 *initiate*. As we shall see later, the agent has a set of BDI states (including its belief that
 509 the communication channel is operating, $B : B_i(channel_good, t)$) that holds true such
 510 that the collective interval, Θ , of those states overlaps the communicative act's interval,
 511 $Overlaps([\Theta], [initiate])$. That is, the communicative act may continue after Θ no longer
 512 holds. Further, if the communication fails, which in this case means the message is not
 513 sent, the agent generates a belief $B : \neg B_i(channel_good, k)$. This becomes a terminator
 514 of *initiate* since one of the preconditions for the communicative act and the new belief
 515 state are contradictory and $Meets(t,k)$. As a result, the action is terminated.

516 Note that a proposition or an assertion in our logical negotiation protocol holds
 517 until terminated by a terminator, and so does a terminator. A termination may lead
 518 to the stoppage of a set of actions or tasks, which in turn may lead to the stoppage of
 519 a set of events. An action or task may, however, terminate by itself normally as it
 520 completes within a time interval.

521 **3.3.4. "CanDo," "Do," and "Doing," and time** Here we define the predicates
 522 *CanDo*, *Do*, and *Doing* and their association with time. Basically, if an agent believes
 523 that it can perform a requested task, then it is willing to be convinced to desire to
 524 perform that task. Once the agent has the desire, it intends to perform the task and
 525 believes that it will be performing the task soon. Once the agent is performing the

526 task, it continues to do so until the task is completed or until the agent no longer
 527 intends to perform the task or until it does not believe it can perform the task
 528 anymore. We present these three predicates since they are related to the underlying
 529 motivations of our negotiations.

530 *Predicate 1 CanDo*

531 When, at time t_1 , an α agent believes it can do a task ρ , and it holds this belief until
 532 time t_2 when a terminator contradicts the belief, we have $B_\alpha(\text{CanDo}(\alpha, \rho), t_B)$ where
 533 t_B is the interval with the start time t_1 and the end time t_2 .

534 *Predicate 2 Do*

535 When, an agent α desires to do a task ρ , then we must have $B_\alpha(\text{CanDo}(\alpha, \rho), t_B)$
 536 and $D_\alpha(\text{Do}(\alpha, \rho), t_D)$ such that $\text{During}(t_D, T_B)$. That is, the desire to do must be
 537 supported by a *CanDo* belief. Following a strong realist model, we also must have
 538 the following: $D_\alpha(\text{Do}(\alpha, \rho), t_D) \Rightarrow I_\alpha(\text{Do}(\alpha, \rho), t_I)$ such that $\text{During}(t_I, T_D)$. Note that
 539 we simplify this example by assuming that the only way to satisfy the desire to do a
 540 task is to perform the task. From this definition, an agent may no longer intend to
 541 perform a task even when it desires to. This is because our agents are susceptible to
 542 external states such as the communicative acts and the sensory information. In cases
 543 where the intended task is one of the many choices available to satisfy a desire, then
 544 agents are also susceptible to internal states. For example, suppose an agent desires
 545 to be not hungry, and it has two choices: *Cook* or *Go-To-Restaurant*. And, suppose
 546 initially, the agent chooses to cook at home: $D_\alpha(\neg\text{hungry}(\alpha), t_D) \wedge I_\alpha(\text{Cook}(\alpha), t_I)$. But
 547 after cooking for a time of t_c , such that $\text{During}(t_c, t_I)$, the agent realizes that the faucet
 548 stops working and there is no water in its kitchen, which constitutes a termination
 549 condition to *Cook*, thus $I_\alpha(\neg\text{Cook}(\alpha), t_I')$ but $D_\alpha(\neg\text{hungry}(\alpha), T_D)$. This allows the
 550 agent to replan and to $I_\alpha(\text{Go} - \text{To} - \text{Restaurant}(\alpha), t_I'')$.

551 *Predicate 3 Doing*

552 If an agent α intends to perform a task, then it believes that it is performing the
 553 task in the following temporal relationship: $I_\alpha(\text{Do}(\alpha, \rho), t_I) \Rightarrow B_\alpha(\text{Doing}(\alpha, \rho), t_B')$
 554 such that $\text{Overlaps}(t_I, t_B') \vee \text{During}(t_B', t_I)$. That means, an agent may still believe that
 555 it is performing a task that it no longer intends to perform. In Vere's [56] definition,
 556 $I_\alpha(\text{Do}(\alpha, \rho), t_I)$ is a trigger pre-condition. For example, an agent first intends to
 557 re-orient its sensors. During the re-orientation, it receives other information and
 558 decides to no longer intend to do so. However, once a re-orientation is launched,
 559 the agent has to wait for it to end. This type of task is considered to be *atomic* and
 560 *non-interruptible*, and this situation is observed in our domain of application.

561 **3.3.5. Coalition formation** When an agent, for example, senses a target, it needs to
 562 form a coalition from which to ask for help. To do so, it collects all its current BDI
 563 states and external information to obtain a list of potentially helpful neighbors. Since
 564 each task is new, that means in the beginning of the coalition formation, there exist
 565 no belief states such that $B : B_\alpha(\text{CanDo}(n, \rho), t_B)$ where n is a member of the set of all
 566 known neighbors, N_α , of the agent α and ρ is the new task. After a domain-specific
 567 search process,⁵ the agent obtains a list of potentially helpful neighbors, $N'_\alpha \subseteq N_\alpha$. At

568 this point, equipped with this list, the agent $\forall nB : B_x(CanDo(n, \rho), t_B)$ where $n \in N'_x$.
 569 Then the agent checks its needs and creates desires for only a certain number of these
 570 neighbors to perform the tasks (since enlisting everybody one knows to help out a
 571 task is counter-productive). In our actual agent design, these neighbors are ranked
 572 according to their utility values.⁶ The agent then checks the number of available
 573 negotiation threads that it may use to negotiate with these neighbors. That number
 574 trims N'_x to obtain N''_x . At this point, the agent $\forall nD : D_x(Do(n, \rho), t_D)$ where $n \in N''_x$.
 575 Subsequently, as the agent begins to send out requests, with each successful contact,
 576 the agent forms $I : I_x(Negotiate(n, Do(n, \rho)), t_I)$ with the neighbor n that has made
 577 contact. In the end, all neighbors contacted are in the coalition $C_x(\vec{\rho})$ such that
 578 $C_x(\vec{\rho}) \subseteq N''_x$ and where $\vec{\rho} = \{\rho_1, \rho_2, \dots, \rho_{|C_x(\vec{\rho})|}\}$ is a set of subtasks that contribute to
 579 the original task. In our model, we define subtasks of a task in terms of the resources
 580 required. For example, in our application to a distributed sensor network, the task of
 581 tracking a target requires three subtasks: each subtask requires an agent to turn on
 582 its sensor towards a particular direction at a given start time. In general, though not
 583 required in our model, subtasks could be seen as horizontal partitioning of the
 584 overall task. In this way, each neighbor is contacted to perform a subtask and a
 585 coalition formation drives an agent to negotiate with its neighbors.

586 In terms of the temporal interval relationships, since a negotiation process is
 587 stepwise and interruptible, if an agent does not have the desire for a neighbor to
 588 perform a task, then it does not intend to negotiate with that neighbor regarding that
 589 particular task. Thus, we have the following condition: $During(t_D, t_B) \wedge$
 590 $During(t_I, t_D)$.

591 Note that in our model, we have N 1-to-1 negotiations but do not conduct *direct* 1-to-
 592 N negotiations. That is, during a negotiation, a negotiation thread does not consult
 593 directly other negotiation threads of the same agent. However, the parent agent of the
 594 negotiation threads does examine the completion status of its negotiation threads and
 595 may change the beliefs, desires, and intentions of the negotiation threads through
 596 negotiation-related predicates (Section 3.4) due to the results of other negotiations.
 597 This design choice is motivated by real-time concerns. Instead of having the negotiator
 598 module of a negotiation thread monitoring the activities of other negotiation threads of
 599 the same agent, the core thread of the agent monitors the negotiation activities through
 600 the coalition manager. The coalition manager determines whether a coalition is still
 601 viable, whether a coalition has been achieved, and whether a coalition is to be aborted,
 602 and commands each individual negotiation thread accordingly. Thus, each negotiator
 603 can concentrate on their negotiation task at hand.

604 Here we briefly describe our coalition formation approach. Interested readers
 605 are referred to [50]. Our coalition formation consists of three stages: (1) coalition
 606 initialization where an agent obtains a ranked list of potentially helpful neighbors,
 607 $N'_x \subseteq N_x$ based on the current problem, (2) coalition finalization where the agent
 608 contacts the neighbors in $N'_x \subseteq N_x$ to negotiate, and (3) coalition acknowledgment
 609 where the agent concludes the success or failure of the coalition and inform
 610 neighbors who have agreed to help. In coalition initialization, a neighbor is
 611 ranked based on its potential utility in helping with the current problem. This
 612 potential utility is based on the past and current relationships between the agent
 613 and the neighbor, and the ability of the neighbor with the current problem. For

614 example, if the target is moving towards the sensor coverage of the neighbor and
 615 will be inside the coverage for a long time, then the neighbor has a high potential
 616 utility. During the coalition finalization step, the agent negotiates with the ranked
 617 neighbors concurrently. To negotiate, each agent uses CBR to derive an appropriate
 618 strategy to deal with each neighbor. As each negotiation thread reports its
 619 final status to the core thread (the parent agent), the parent agent decides whether
 620 to abort meaningless negotiations or to modify negotiation tactics. After the
 621 finalization step, the agent knows whether it has a coalition. If it does, it sends a
 622 confirmation message to all the neighbors who have agreed to help. If it does not,
 623 it sends a discard message to those who have agreed to help. This is the
 624 acknowledgment step.

625 **3.3.6. Implicit assumptions** As discussed in Section 2.1, our agents are cooperative
 626 and are also directed to satisfy global goals. Each is motivated to look for help from
 627 its neighbors and to entertain negotiation requests from its neighbors, and each
 628 genuinely wishes to have a successful negotiation. First, we have the implicit
 629 assumption $D : D_x(\text{Cooperate}, t_{\text{always}})$ where $t_{\text{always},s} = \text{time}0$ is the start time and
 630 $t_{\text{always},f} = \infty$ is the finish time, meaning that the desire is always true. Second, we
 631 have $D : D_x(\text{Satisfy} - \text{Global} - \text{Goal}, t_{\text{always}})$. These assumptions are the meta-level
 632 knowledge in an agent's context.

633 When an initiating agent, i , negotiates with a neighbor, r , the agent has the fol-
 634 lowing BDI states, as shown in the previous section:

$$D : D_i(\text{Cooperate}, t_{\text{always}}) \wedge B : B_i(\text{CanDo}(r, \rho), t_B) \wedge \\ D : D_i(\text{Do}(r, \rho), t_D) \wedge I : I_i(\text{Negotiate}(r, \text{Do}(r, \rho)), t_I).$$

636 Note that in [36, 40, 44], the mental states were *layered* or divided into different
 637 levels: beliefs, desires, and intentions, and bridge rules were used to transition across
 638 these levels. Here we simply outline a set of *propositions*, linking the multiple con-
 639 texts together in conjunctions. The key here is that each proposition is also quan-
 640 tified by a time interval. So, in a way, our model indicates that at certain time
 641 intervals, there are certain mental states (beliefs, desires, or intentions) that are true
 642 simultaneously (if their time intervals overlap). Thus, the above conjunction simply
 643 means that for some time interval there is a belief of CanDo, for some other interval,
 644 there is a desire of Do, and so on.

645 Combining the above states with $D : D_i(\text{Satisfy} - \text{Global} - \text{Goal}, t_{\text{always}})$, we as-
 646 sume that each agent:

$$I : I_i(\text{succeed}(\text{Negotiate}(r, \text{Do}(r, \rho))), t_I).$$

648 The above intention motivates an initiating agent to continue negotiating. In the
 649 later discussion, we use this intention to explicitly drive the negotiation axioms, while
 650 keeping other BDI states implicit.

651 When a responding agent, r , receives a request to negotiate from an initiating
 652 agent, i , the agent parses the message and examines its own current states. If it
 653 believes it can perform the requested task yet does not have the desire to do so, then
 654 because of the desire to cooperate, it has the following BDI states:

$$D : D_r(\text{Cooperate}, t_{\text{always}}) \wedge B : B_r(\text{CanDo}(r, \rho), t_B) \wedge \\ \neg \exists D : D_r(\text{Do}(r, \rho), t_D) \wedge I : I_r(\text{Negotiate}(i, \text{Do}(r, \rho)), t_I).$$

656 Similarly, combining the above states with $D : D_r(\text{Satisfy} - \text{Global} - \text{Goal}, t_{\text{always}})$,
657 we have the following that keeps the agent negotiating:

$$I : I_r(\text{succeed}(\text{Negotiate}(i, \text{Do}(r, \rho))), t_I).$$

659 A completely successful negotiation results in $D : D_r(\text{Do}(r, \rho), t_D)$. The negotiating
660 strategy of an initiating agent is to help the responding agent achieve
661 $D : D_r(\text{Do}(r, \rho), t_D)$ while that of the responding agent is to let itself be persuaded by
662 the initiating agent's arguments in order to achieve $D : D_r(\text{Do}(r, \rho), t_D)$. As we will
663 see in the next two sections, there are partially successful negotiations and different
664 types of failures.

665 Furthermore, we have to deal with the duration of a BDI state. A BDI state
666 holds true only sufficiently long: (1) no longer than the duration of the task to be
667 performed, or (2) until ended by a terminator. For the implicit assumptions, the
668 tasks *Cooperate* and *Satisfy-Global-Goal* have infinite duration. So, in
669 $D : D_r(\text{Do}(r, \rho), t_D)$, the agent r only desires to do the task for at most the
670 duration of the task ρ . Thus, we have $\text{During}(t_D, [\text{Do}(r, \rho)])$ where $[\text{Do}(r, \rho)]$ is the
671 time interval for r performing the task. This assumption applies to all BDI states
672 related to performing a task. On the other hand, when a negation of performing
673 a task is involved, such as $D : D_r(\neg \text{Do}(r, \rho), t_D)$, the agent cannot rely on
674 $[\text{Do}(r, \rho)]$ to quantify t_D . In our design, instead of making t_D a *moment* or a *point*,
675 we let it be until terminated by another assertion/proposition (e.g., generated by a
676 later coalition process). These assumptions allow an agent to negotiate regarding
677 the same task at two different times as long as the task negotiated first has
678 completed since the BDI states of performing a task are self-terminating and
679 those of *not* performing a task are terminated by assertions generated by other
680 agent activities.

681 Finally, we have to deal with arguments since our negotiation approach is argu-
682 mentation-based. Suppose the request is for the responding agent r to perform the task
683 $\rho : \text{Do}(r, \rho)$. The responding agent has a set of internal arguments, γ_r , for and against
684 performing the task. If the agent believes it can perform the task but
685 $\Gamma_r \not\models D_r(\text{Do}(r, \rho), t_D)$ (meaning the arguments do not support the desire for performing
686 the task), then it has to rely on the initiating agent for more arguments. The initiating
687 agent has its own set of arguments, γ_i , for the responding agent performing the task.
688 The underlying approach is to send over a subset Γ'_i of γ_i to the responding agent until

$$\Gamma_r \cup \Gamma'_i \models D_r(\text{Do}(r, \rho), t_D) \text{ (in which case the negotiation succeeds),}$$

690 or until

$$\Gamma_r \cup \Gamma_i \not\models D_r(\text{Do}(r, \rho), t_D) \text{ (in which case the negotiation fails),}$$

692 where $\Gamma'_i \subseteq \Gamma_i$ is the set of arguments already communicated to the responding agent
693 from the initiating agent. This assumption is a critical element in our negotiation
694 protocol as it facilitates a stepwise evaluation of arguments to move closer to a

695 conclusion of the negotiation. An example of arguments supporting a requested task
 696 in the sensor network application includes the approximate direction of the target
 697 detected, the current, activated sensor sector, and the helpfulness of the agent to its
 698 neighbors. The approximate direction of the detected target would help the
 699 responding agent deliberate whether it has available sensor sector for the task, for
 700 example.

701 **3.3.7. Initiating behavior** Here we outline the axioms that link an agent's commu-
 702 nication and its internal states for conducting negotiations as an initiating agent.
 703 Before moving further, we would like to note that the implications in the following
 704 axioms are basically bridge rules. The pre-conditions of these rules do give the
 705 impression that multiple contexts are present. As previously described in Section 3.3.6,
 706 we mean to have them as separate contexts which happen to be conditions required to
 707 be present to trigger the rules. The formulation of these contexts in the pre-conditions
 708 allows us to develop a time schedule for the various mental states to hold true.

709 **Initiate:** When an initiating agent (i) believes that it intends to negotiate with the
 710 responding agent (r) to perform a task ρ , it initiates a negotiation request to the
 711 responding agent.

$$I : I_i(\text{Negotiate}(r, Do(r, \rho)), t_I) \Rightarrow C_{out} : \text{initiate}(i, r, Do(r, \rho), t_{cout,initiate})$$

713 where $\text{During}(t_{cout,initiate}, t_I)$. The predicate initiate encapsulates the act of composing
 714 an INITIATE-type message and sending the message to agent r .

715 **Failure 1:** When an initiating agent (i) receives a NO_GO message from a
 716 responding agent (r), it believes that the responding agent r cannot perform the
 717 requested task ρ and stops intending r to perform the task.

$$C_{in} : \text{no_go}(r, i, Do(r, \rho), t_{cin,no_go}) \wedge I : I_i(\text{Negotiate}(r, Do(r, \rho)), t_I) \Rightarrow \\ I : I_i(\neg \text{Negotiate}(r, Do(r, \rho)), t'_I) \wedge B : B_i(\neg \text{CanDo}(r, \rho), t_B) \wedge \\ D : D_i(\neg Do(r, \rho), t_D) \wedge \text{rejected}$$

719 where $\text{During}(t_{cin,no_go}, t_I) \wedge \text{Meets}(t_I, t'_I) \wedge \text{Meets}(t_I, t_B) \wedge \text{Starts}(t_B, t_D)$. The no_go
 720 communicative act is the encapsulation of receiving and parsing a NO_GO (an
 721 outright refusal to negotiate) message. This rule allows an agent to move on to the
 722 next neighbor after the responding agent has outright refused to negotiate. This is
 723 real-time motivated:⁷ instead of trying to come up with a lesser task and trying to
 724 establish a negotiation with the responding agent, the initiating agent simply gives up
 725 and shifts its focus to other neighbors.

726 The proposition rejected indicates the failure of a negotiation.

727 Note also that in the above rule, the changes in the internal states of the agent are
 728 triggered by the incoming message. This is one of our design characteristics and
 729 goals: agents communicate only when necessary since the environment is real-time
 730 and resource constrained, and information is exchanged only during negotiation.

731 **Success 1** When an initiating agent (i) receives an AGREE message from a
 732 responding agent (r), it believes that the responding agent r intends to perform and
 733 will perform the requested task ρ .

$$\begin{aligned}
C_{in} &: \text{agree}(r, i, Do(r, \rho), t_{cin,agree}) \wedge I : I_i(\text{Negotiate}(r, Do(r, \rho)), t_I) \Rightarrow \\
B &: B_i(D_r(Do(r, \rho)), t_B) \wedge I : I_i(\neg \text{Negotiate}(r, Do(r, \rho)), t'_I) \wedge \\
D &: D_i(\neg Do(r, \rho), t_D) \wedge \text{success}
\end{aligned}$$

735 where $\text{During}(t_{cin,agree}, t_I) \wedge \text{Meets}(t_I, t'_I) \wedge \text{Meets}(t_I, t_B) \wedge \text{Starts}(t_B, t_D)$. The *agree*
736 communicative act is the encapsulation of receiving and parsing an AGREE mes-
737 sage. In this axiom, if an agent receives an AGREE message from the responding
738 agent, then (1) it believes that the responding agent desires to perform the task, (2) it
739 intends no longer to negotiate, and (3) it desires no longer that the responding agent
740 perform the task. This third desire may seem counter-intuitive at first glance. Its
741 purpose is to say “If I believe that you have the desire to do the task, then I don’t
742 have to desire you to do the task anymore,” and that does not prevent the agent to
743 desire the responding agent to perform the task in the future.

744 The proposition *success* indicates the success of a negotiation.

745 *Info 1* When an initiating agent (*i*) receives a RESPOND message from a responding
746 agent (*r*), it (1) believes that the responding agent *r* intends to negotiate and (2) intends
747 to obtain a successful negotiation. Consequently, the initiating agent *i* intends to help *r*
748 to desire to perform the task by supplying available necessary information.

$$\begin{aligned}
C_{in} &: \text{respond}(r, i, Do(r, \rho), t_{cin,respond}) \wedge I : I_i(\text{Negotiate}(r, Do(r, \rho)), t_I) \wedge \\
B &: B_i(I_r(\text{Negotiate}(i, Do(r, \rho))), t_B) \\
&\wedge I : I_i(\text{succeed}(\text{Negotiate}(r, Do(r, \rho))), t'_I) \wedge \exists p : (p \in \Gamma_i \wedge p \notin \Gamma'_i) \\
&\Rightarrow C_{out} : \text{info}(i, r, Do(r, \rho), t_{cout,info})
\end{aligned}$$

750 where

$$\begin{aligned}
&\text{During}(t_{cin,respond}, t_I) \wedge \text{Finishes}(t'_I, t_I) \wedge \text{Equals}(t_B, t'_I) \wedge \\
&\text{Starts}(t_{cout,info}, t'_I) \wedge \text{Before}(t_{cin,respond}, t_{cout,info}).
\end{aligned}$$

752 The communicative act *respond* is the encapsulation of receiving and parsing a
753 RESPOND message, and the communicative act *info* is the encapsulation of com-
754 posing and sending an INFO message, including selecting a *p* from the set of
755 arguments, γ_i , that is not a member the set of arguments already sent, Γ'_i to *r*. Clause
756 1 of the axiom explicitly derives the motivation for the initiating agent to continue
757 negotiating as it intends to have a successful negotiation since it now believes that the
758 responding agent intends to negotiate as well. Clause 2 of the axiom drives the agent
759 to send over more arguments to help bring a successful conclusion to the negotiation.

760 Note that as discussed in Section 3.3.2, the syntax of a message is
761 $\text{msg}(\langle \text{sender} \rangle, \langle \text{receiver} \rangle, \langle \text{type} \rangle, \langle \text{request} \rangle, \langle \text{contents} \rangle)$. When sending out an INFO-
762 type message, the $\langle \text{contents} \rangle$ holds the arguments.

763 *Info_null 1* When an initiating agent (*i*) receives a RESPOND message from a
764 responding agent (*r*), it (1) believes that the responding agent *r* intends to negotiate
765 and (2) intends to obtain a successful negotiation. However, if *i* has run out of
766 information or arguments, then it notifies *r* that it can no longer provide arguments.

$$\begin{aligned}
& C_{in} : \text{respond}(r, i, Do(r, \rho), t_{cin,respond}) \wedge I : I_i(\text{Negotiate}(r, Do(r, \rho)), t_I) \\
& \wedge B : B_i(I_r(\text{Negotiate}(i, Do(r, \rho))), t_B) \\
& \wedge I : I_i(\text{succeed}(\text{Negotiate}(r, Do(r, \rho))), t'_I) \wedge \exists p : (p \in \Gamma_i \wedge p \notin \Gamma'_i) \\
& \Rightarrow C_{out} : \text{info_null}(i, r, Do(r, \rho), t_{cout,info_null})
\end{aligned}$$

768 where

$$\begin{aligned}
& \text{During}(t_{cin,respond}, t_I) \wedge \text{Finishes}(t'_I, t_I) \wedge \text{Equals}(t_B, t'_I) \\
& \wedge \text{Starts}(t_{cout,info_null}, t'_I) \wedge \text{Before}(t_{cin,respond}, t_{cout,info_null}).
\end{aligned}$$

770 This rule is the counterpart to *Info 1* previously discussed. When an agent runs out of
771 arguments, it notifies the responding agent about it. Instead of giving up on the
772 negotiation right away – since obviously the initiating agent knows that it has not
773 been able to persuade the responding agent and now it has run out of arguments, the
774 initiating agent informs the responding agent of its situation and hopefully the
775 responding agent will be able to counter-offer. So, in a way, this shifts the respon-
776 sibility of achieving a successful negotiation to the responding agent from the ini-
777 tiating agent. Up until this point, the initiating agent has been responsible for
778 keeping the negotiation going by supplying arguments/information to the respond-
779 ing agent, trying to convince it. Finally when the initiating agent can argue no
780 further, the decision shifts to the responding agent.

781 *Info 2* When an initiating agent (i) receives a MORE_INFO message from a
782 responding agent (r), it simply supplies more unused arguments.

$$\begin{aligned}
& C_{in} : \text{more_info}(r, i, Do(r, \rho), t_{cin,more_info}) \\
& \wedge I : I_i(\text{succeed}(\text{Negotiate}(r, Do(r, \rho))), t'_I) \\
& \wedge \exists p : (p \in \Gamma \wedge p \notin \Gamma'_i) \Rightarrow C_{out} : \text{info}(i, r, Do(r, \rho), t_{cout,info})
\end{aligned}$$

784 where $\text{During}(t_{cin,more_info}, t'_I) \wedge \text{During}(t_{cout,info}, t'_I) \wedge \text{Before}(t_{cin,more_info}, t_{cout,info})$. This
785 axiom is similar to *Info 1*.

786 *Info_null 2* When an initiating agent (i) receives a MORE_INFO message from a
787 responding agent (r), if it does not have any more arguments, then it notifies r of its
788 status.

$$\begin{aligned}
& C_{in} : \text{more_info}(r, i, Do(r, \rho), t_{cin,more_info}) \wedge \\
& I : I_i(\text{succeed}(\text{Negotiate}(r, Do(r, \rho))), t'_I) \\
& \wedge \neg \exists p : (p \in \Gamma_i \wedge p \notin \Gamma'_i) \Rightarrow C_{out} : \text{info_null}(i, r, Do(r, \rho), t_{cout,info_null})
\end{aligned}$$

790 where $\text{During}(t_{cin,more_info}, t'_I) \wedge \text{During}(t_{cout,info_null}, t'_I) \wedge \text{Before}(t_{cin,more_info}, t_{cout,info_null})$.
791 This axiom is similar to *Info_null 1*.

792 *Info 3* When an initiating agent (i) receives a counter-offer (ρ') from a responding
793 agent (r), i believes that r desires to perform ρ' . However, if ρ' is not acceptable, then
794 the agent i continues to send unused arguments to r .

$$\begin{aligned}
& C_{in} : counter(r, i, Do(r, \rho), t_{cin,counter}) \wedge B : B_i(D_r(Do(r, \rho')), t_B) \\
& \wedge B : B_i(\neg acceptable(\rho'), t'_B) \\
& \wedge \exists p : (p \in \Gamma_i \wedge p \notin \Gamma'_i) \wedge I : I_i(succeed(Negotiate(r, Do(r, \rho))), t'_I) \\
& \Rightarrow C_{out} : info(i, r, Do(r, \rho), t_{cout,info})
\end{aligned}$$

796 where

$$\begin{aligned}
& During(t_{cin,counter}, t'_I) \wedge Meets(t_{cin,counter}, t_B) \wedge During(t'_B, t'_I) \\
& \wedge Finishes(t'_B, t_B) \wedge Meets(t_B, t_{cout,info}) \\
& \wedge During(t_{cout,info}, t'_I).
\end{aligned}$$

798 This axiom is similar to *Info 1*. The communicative act *counter* is the encapsulation of
799 receiving and parsing a COUNTER-type message, in which $\langle contents \rangle$ holds the
800 counter-offer ρ' . Clause 1 states that when an initiating agent receives a counter-offer
801 from the responding agent, it believes that the responding agent desires to perform the
802 counter-offer. Now, the initiating agent checks the acceptability of the counter-offer
803 (see Section 5). If the counter-offer is not acceptable and the agent still has unused
804 arguments, then it sends over more arguments. This is how an initiating agent counter-
805 offers a counter-offer: sending over more arguments in hope that the responding agent
806 will come back with a better counter-offer, closer to the original request.

807 Note also the temporal relationships among $t_{cout,info}$, t_B , and t'_B . As soon as the
808 initiating agent realizes that the counter-offer is not acceptable, both its beliefs that
809 the responding agent desires to perform the counter-offer and that the counter-offer
810 is unacceptable terminate and trigger the communicative act *info*. In other words,
811 when the initiating agent counters a counter-offer, all beliefs regarding the counter-
812 offer no longer hold.

813 *Info_null 3* When an initiating agent (i) receives a counter-offer (ρ') from a
814 responding agent (r), i believes that r desires to perform ρ' . However, if ρ' is not
815 acceptable and the agent does not have any more unused arguments, it notifies r that
816 it can no longer provide arguments.

$$\begin{aligned}
& C_{in} : counter(r, i, Do(r, \rho), t_{cin,counter}) \wedge B : B_i(D_r(Do(r, \rho')), t_B) \\
& \wedge B : B_i(\neg acceptable(\rho'), t'_B) \wedge \\
& \neg \exists p : (p \in \Gamma_i \wedge p \notin \Gamma'_i) \wedge I : I_i(succeed(Negotiate(r, Do(r, \rho))), t'_I) \\
& \Rightarrow C_{out} : info_null(i, r, Do(r, \rho), t_{cout,info_null})
\end{aligned}$$

818 where

$$\begin{aligned}
& During(t_{cin,counter}, t'_I) \wedge Meets(t_{cin,counter}, t_B) \wedge During(t'_B, t'_I) \\
& \wedge Finishes(t'_B, t_B) \wedge Meets(t_B, t_{cout,info_null}) \wedge During(t_{cout,info_null}, t'_I).
\end{aligned}$$

820 This axiom is the counterpart of *Info 3*.

821 *Success 2* When an initiating agent (i) receives a counter-offer (ρ') from a responding
822 agent (r), i believes that r desires to perform ρ' . If ρ' is acceptable, then i agrees.

$$\begin{aligned}
& C_{in} : counter(r, i, Do(r, \rho), t_{cin,counter}) \wedge B : B_i(D_r(Do(r, \rho')), t_B) \\
& \wedge B : B_i(acceptable(\rho'), t'_B) \wedge I : I_i(succeed(Negotiate(r, Do(r, \rho))), t'_I) \\
& \Rightarrow C_{out} : agree(i, r, Do(r, \rho), t_{cout,agree}) \wedge B : B_i(\neg CanDo(r, \rho), t''_B) \\
& \wedge D : D_i(\neg Do(r, \rho), t_D) \wedge \\
& I : I_i(\neg Negotiate(r, Do(r, \rho)), t''_I) \wedge I : I_i(\neg succeed(Negotiate(r, Do(r, \rho))), t''_I) \\
& \wedge success
\end{aligned}$$

824 where

$$\begin{aligned}
& During(t_{cin,counter}, t'_I) \wedge During(t'_B, t'_I) \wedge Finishes(t'_B, t_B) \wedge Meets(t_B, t_{cout,agree}) \\
& \wedge Meets(t'_I, t''_I) \wedge Meets(t_B, t''_B) \wedge Starts(t_D, t''_I) \wedge During(t_{cout,agree}, t''_I).
\end{aligned}$$

826 The communicative act *agree* is the encapsulation of composing and sending an
827 AGREE message, in which $\langle contents \rangle$ holds the counter-offer ρ' . With this axiom, as
828 soon as the initiating agent agrees to a counter-offer by the responding agent, it (1)
829 believes that the responding agent cannot do the originally requested task, (2) desires
830 no longer that the responding agent performs the task, (3) intends to negotiate no
831 further regarding the task, and (4) intends no longer to have a successful negotiation
832 regarding the task. However, the negotiation still ends with a *success* tag because
833 even though the initiating agent does not get what it wanted originally, it does obtain
834 a portion of its original request.

835 *Failure 2* When an initiating agent (i) receives a STOP message from a responding
836 agent (r), it believes that the responding agent r does not desire to perform the
837 requested task ρ and thus stops negotiating with r to perform the task, and the
838 negotiation fails.

$$\begin{aligned}
& C_{in} : stop(r, i, Do(r, \rho), t_{cin,stop}) \wedge I : I_i(succeed(Negotiate(r, Do(r, \rho))), t'_I) \Rightarrow \\
& B : B_i(\neg CanDo(r, \rho), t_B) \wedge D : D_i(\neg Do(r, \rho), t_D) \\
& \wedge I : I_i(\neg Negotiate(r, Do(r, \rho)), t''_I) \\
& \wedge I : I_i(\neg succeed(Negotiate(r, Do(r, \rho))), t''_I) \wedge rejected
\end{aligned}$$

840 where $During(t_{cin,stop}, t'_I) \wedge Meets(t'_I, t_B) \wedge Meets(t'_I, t_D) \wedge Meets(t'_I, t''_I)$. The commu-
841 nicative act *stop* is the encapsulation of receiving and parsing a STOP message. This
842 rule is similar to *Failure 1*. In addition, it also states that the agent intends to not
843 negotiate. In our current design we do not differentiate between an outright failure
844 (failure type 1) and an opt-out failure (failure type 2) – both end with a *rejected* tag.

845 *Failure 3I* This rule is similar to *Failure 2* except for that it deals with an ABORT
846 message and ends with an *abort* tag. See Appendix A.

847 *Failure 4I* This rule is similar to *Failure 2* except for that it deals with an OUT_-
848 OF_TIME message and ends with an *out_of_time* tag. See Appendix A.

849 *Abort I* When an initiating agent (i) no longer intends to negotiate with a
850 responding agent (r) to perform a requested task ρ , it aborts the negotiation.

$$\begin{aligned}
I &: I_i(\neg \text{Negotiate}(r, \text{Do}(r, \rho)), t'_I) \Rightarrow C_{out} : \text{abort}(i, r, \text{Do}(r, \rho), t_{cout, abort}) \wedge \\
D &: D_i(\neg \text{Do}(r, \rho), t_D) \wedge I : I_i(\neg \text{succeed}(\text{Negotiate}(r, \text{Do}(r, \rho))), t''_I) \wedge \text{abort}
\end{aligned}$$

852

853 where $\text{During}(t_{cout, abort}, t'_I) \wedge \text{During}(t_{cout, abort}, t_D) \wedge \text{During}(t_{cout, abort}, t''_I) \wedge \text{Finishes}(t''_I,$
854 $t'_I)$. The communicative act *abort* is the encapsulation of composing and sending an
855 ABORT message. This rule says that if an initiating agent aborts a negotiation, then
856 it informs the responding agent.

857 *Out_of_time I* When an initiating agent (*i*) runs out of its allocated time for the
858 negotiation with a responding agent (*r*) to perform a requested task ρ , it aborts the
859 negotiation.

$$\begin{aligned}
B &: B_i(\neg \text{time}(\text{Negotiate}(r, \text{Do}(r, \rho))), t_B) \\
&\wedge I : I_i(\text{succeed}(\text{Negotiate}(r, \text{Do}(r, \rho))), t'_I) \Rightarrow \\
C_{out} &: \text{out_of_time}(i, r, \text{Do}(r, \rho), t_{cout, out_of_time}) \wedge I : I_i(\neg \text{Negotiate}(r, \text{Do}(r, \rho)), t''_I) \wedge \\
I &: I_i(\neg \text{succeed}(\text{Negotiate}(r, \text{Do}(r, \rho))), t''_I) \wedge \text{out_of_time}
\end{aligned}$$

861 where $\text{During}(t_B, t'_I) \wedge \text{During}(t_{cout, out_of_time}, t''_I) \wedge \text{Meets}(t'_I, t''_I)$. The *time* predicate
862 encapsulates the acts of obtaining and comparing the time elapsed for the negotia-
863 tion against the time allocated for the negotiation (Section 3.4). The communicative
864 act *out_of_time* predicate is the encapsulation of composing and sending an
865 OUT_OF_TIME message to the responding agent. This rule states that when the
866 agent has run out of time allocated for the negotiation, it no longer intends to
867 negotiate. This is real-time motivated.

868 *No_response I* When an initiating agent (*i*) detects receives no response from a
869 responding agent (*r*) during a negotiation, then it unilaterally quits the negotiation
870 with a failure.

$$\begin{aligned}
B &: B_i(\text{no_response}(r), t_B) \wedge I : I_i(\text{succeed}(\text{Negotiate}(r, \text{Do}(r, \rho))), t'_I) \Rightarrow \\
I &: I_i(\neg \text{Negotiate}(r, \text{Do}(r, \rho)), t''_I) \wedge I : I_i(\neg \text{succeed}(\text{Negotiate}(r, \text{Do}(r, \rho))), t''_I) \\
&\wedge \text{channel_jammed}
\end{aligned}$$

872 where $\text{During}(t_B, t'_I) \wedge \text{Meets}(t'_I, t''_I)$. The *no_response* predicate is one of our real-
873 time enabling functional predicates to be discussed in Section 3.4. This axiom allows
874 an agent to bail out of a negotiation when the negotiation partner fails to respond.

875 Figure 3 shows the time lines of the initiating agent's behavior when faced with an
876 outright rejection or agreement. These are the simple cases of the axioms above. The
877 length of the process is based on $t_{cout, initiate}$, t_{cin, no_go} , and $t_{cin, agree}$. During such time,
878 $I : I_i(\text{Negotiate}(r, \text{Do}(r, \rho)), t_I)$ holds true. After that, the intention can be removed
879 or modified.

880 Figure 4 shows the negotiation process, from the initiating agent's point of view
881 once the responding agent agrees to negotiate. The negotiation process is a mani-
882 festation of the axioms discussed above, proving a flow of communicative acts and
883 BDI states that drives the completion of the negotiation. It is with the temporal BDI




Figure 3. Time lines of the initiating agent's behavior when faced with an outright rejection or agreement.

884 axioms that we are able to produce Figure 4, an explicit outline of the interactions of
 885 the communicative acts with various BDI states – specifying when and how long
 886 certain states must hold true, cannot be modified, can change, can be accessed, or are
 887 of no concern. It is also through the axioms that we are able to guarantee the
 888 completion of a negotiation process within a certain time. For example, if the initi-
 889 ating agent goes through the following steps: *initiate*, *parse respond*, *info*, *parse*
 890 *more_info*, *info*, *parse more_info*, *info_null*, *parse counter*, check to see whether the
 891 counter-offer is acceptable, and agree, then we know how much time it takes to do so
 892 by summing up the temporal intervals associated with each step. The acceptability of
 893 the counter-offer has to be held constant throughout the agreement step. This
 894 explicit declaration of time constraints is important since each of our agents is

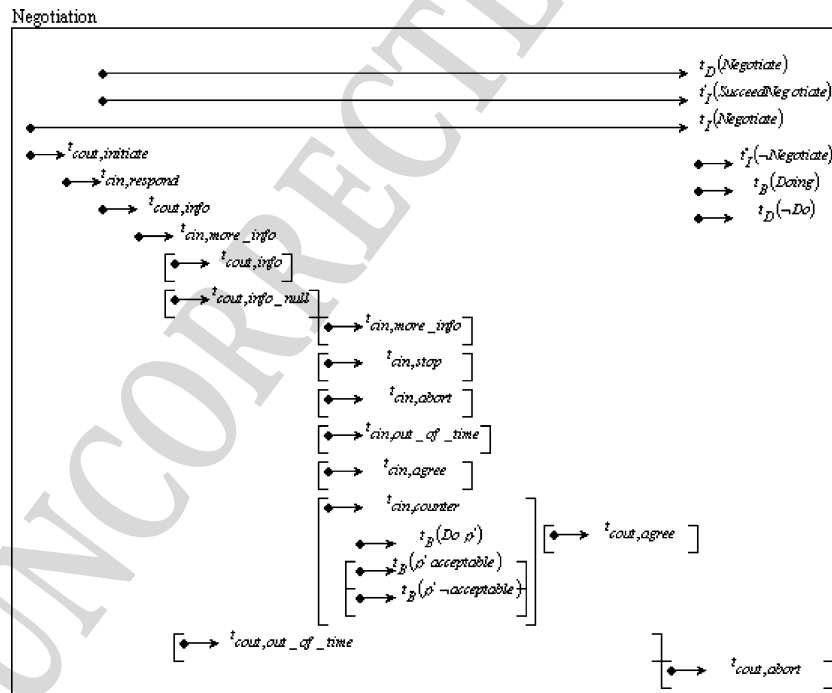


Figure 4. Time lines of the negotiation from the initiating agent's point of view (a simplified version). Temporal intervals under brackets are options. Note that once the initiating agent receives a RESPOND message from the responding agent, it is committed to negotiate successfully since it believes that the responding agent desires to negotiate.

895 multi-threaded, where several threads may access and need to modify the same
 896 variable at the same time. Without the axioms, a variable such as the acceptability of
 897 the counter-offer might *accidentally* be modified by another negotiation thread,
 898 rendering the current negotiation process ambiguous. Moreover, with the temporal
 899 intervals, we are able to fine-tune the system by observing the time usage of each
 900 communicative act for speedup. For example, with the BDI states, we know the
 901 minimal time we need to hold the value of a variable constant. The shorter the time
 902 needed for a variable to be held constant, the more frequent the variable can be
 903 updated and accessed, allowing other threads to proceed.

904 **3.3.8. Responding behavior** Here we outline axioms that link an agent's commu-
 905 nication and its internal states for conducting negotiations as a responding agent.

906 *No_go* When a responding agent (r) believes that it cannot perform a requested
 907 task ρ from an initiating agent i , it outright refuses to negotiate.

$$C_{in} : initiate(i, r, Do(r, \rho), t_{cin,initiate}) \wedge B : B_r(\neg CanDo(r, \rho), t_B) \Rightarrow \\ C_{out} : no_go(r, i, Do(r, \rho), t_{cout,no_go}) \wedge exit$$

909 where $Meets(t_{cin,initiate}, t_B) \wedge During(t_{cout,no_go}, t_B)$. The communicative act *no_go*
 910 encapsulates the act of composing a NO_GO message and sending the message to
 911 agent i . Agents are responsible in that if a responding agent refuses to negotiate, it
 912 informs the initiating agent.

913 *Agree 1* When a responding agent (r) (1) believes that it is already performing a
 914 requested task ρ or (2) desires to perform ρ , it agrees to perform the task.

$$C_{in} : initiate(i, r, Do(r, \rho), t_{cin,initiate}) \wedge (B : B_r(Doing(r, \rho), t_B) \vee D : D_r(Do(r, \rho), t_D)) \\ \Rightarrow C_{out} : agree(r, i, Do(r, \rho), t_{cout,agree}) \wedge D : D_r(Do(r, \rho), t'_D) \wedge success$$

916 where

$$(Overlaps(t_{cin,initiate}, t_B) \wedge Before(t_{cin,initiate}, t_{cout,agree}) \wedge During(t_{cout,agree}, t_B) \\ \wedge Meets(t_{cout,agree}, t'_D)) \\ \vee (Overlaps(t_{cin,initiate}, t_D) \wedge Before(t_{cin,initiate}, t_{cout,agree}) \wedge Finishes(t_{cout,agree}, t_D) \\ \wedge Meets(t_D, t'_D)).$$

918 The communicative act *agree* encapsulates the acts of composing an AGREE
 919 message and sending the message to the initiating agent. The rule *extends temporally*
 920 the desire of the agent to continue performing the task. As previously discussed in
 921 Section 3.3.4, an agent may be still performing a task while it no longer desires to do
 922 so because a task may be atomic and non-interruptible. So, when the responding
 923 agent realizes the initiating agent requests for the same task to be performed, then it
 924 re-asserts its desire to ensure the continuation of the task. We also use the relation
 925 $Meets(t_D, t'_D)$ to transition the desire.

926 Note also that with this rule, the responding agent agrees to help only because it is
 927 already performing the task, but the *agree* predicate does not reveal that to the

928 initiating agent. This simplifies our agent design in two ways: (1) the responding
 929 agent does not have to explain to the initiating agent why it agrees to perform a
 930 requested task, and (2) the initiating agent does not have to remember why the
 931 responding agent agreed to a requested task.

932 *Respond* When a responding agent (r) believes that it can perform a requested task
 933 ρ , and there is no desire to perform ρ nor belief that it is performing ρ , it responds to
 934 the negotiation request, i.e., it agrees to negotiate.

$$\begin{aligned} C_{in} &: \text{initiate}(i, r, Do(r, \rho), t_{cin,initiate}) \wedge B : B_r(\text{CanDo}(r, \rho), t_B) \wedge \\ &\quad \neg \exists D : D_r(Do(r, \rho), t_D) \wedge \neg \exists B : B_r(\text{Doing}(r, \rho), t'_B) \\ \Rightarrow C_{out} &: \text{respond}(r, i, Do(r, \rho), t_{cout,respond}) \wedge \\ I &: I_r(\text{Negotiate}(i, Do(r, \rho)), t_I) \wedge I : I_r(\text{succeed}(\text{Negotiate}(i, Do(r, \rho))), t'_I) \end{aligned}$$

936 where

$$\begin{aligned} & \text{Meets}(t_{cin,initiate}, t_B) \wedge \text{During}(t_B, t_D) \wedge \text{During}(t_B, t'_B) \\ & \wedge \text{During}(t_{cout,respond}, t_B) \wedge \\ & \text{Meets}(t_{cout,respond}, t_I) \wedge \text{Meets}(t_{cout,respond}, t'_I) \wedge \text{Finishes}(t'_I, t_I) \\ & \wedge \text{Before}(t_{cin,initiate}, t_{cout,respond}). \end{aligned}$$

938 The communicative act *respond* encapsulates the acts of composing a RESPOND
 939 message and sending the message to the initiating agent. If the responding agent be-
 940 lieves it can perform the task, and yet it currently does not have a desire to do so, and
 941 does not believe it is performing the task, then it decides to negotiate. Note that the
 942 implicit assumptions discussed in Section 3.3.6 are at play here. Because of the coop-
 943 erativeness of the agent, it intends to negotiate and intends to negotiate successfully.
 944 These two intentions motivate the responding agent to continue negotiating.

945 *More_info* When a responding agent (r) receives arguments from an initiating
 946 agent (i) for a requested task ρ , it processes the arguments to update evidence
 947 support for the task. If the support is still lacking, and the negotiation is on pace or
 948 the task is discrete, then it asks for more arguments.

$$\begin{aligned} C_{in} &: \text{info}(i, r, Do(r, \rho), t_{cin,info}) \wedge I : I_r(\text{succeed}(\text{Negotiate}(i, Do(r, \rho))), t'_I) \\ & \wedge I : I_r(\text{update}(\Gamma'_i, t''_I)) \wedge \\ & \neg \exists D : D_r(Do(r, \rho), t_D) \wedge (B : B_r(\neg \text{slow}(\text{Negotiate}(i, Do(r, \rho))), t_B) \\ & \vee B : B_r(\text{discrete}(\rho), t'_B)) \\ \Rightarrow C_{out} &: \text{more_info}(r, i, Do(r, \rho), t_{cout,more_info}) \end{aligned}$$

950 where

$$\begin{aligned} & \text{Meets}(t_{cin,info}, t''_I) \wedge \text{During}(t_{cin,info}, t'_I) \wedge \text{During}(t''_I, t'_I) \\ & \wedge \text{Meets}(t''_I, t_D) \wedge \text{During}(t_D, t'_I) \\ & \wedge \text{During}(t_D, t_B) \wedge \text{During}(t_D, t'_B) \wedge \text{During}(t_{cout,more_info}, t_D) \\ & \wedge \text{Before}(t_{cin,info}, t_{cout,more_info}). \end{aligned}$$

952 *Auxiliary to More_info* The action *update* examines the arguments collected, Γ'_i
 953 during $t_{evidence}$ (Section 3.3.6) where $Finishes(t_{evidence}, t''_I)$, to see if the proposition
 954 $\Gamma'_i \cup \Gamma_r \models D : D_r(Do(r, \rho), t_D)$. If so, then $D : D_r(Do(r, \rho), t_D)$ where $Meets(t''_I, t_D)$.
 955 First, the communicative act *info* encapsulates the actions of receiving and
 956 parsing an INFO-type message from the initiating agent, in which the $\langle contents \rangle$
 957 holds the arguments $p \in \Gamma_i$ during t_{comm} and $p \notin \Gamma'_i$ during t_{comm} where
 958 $Before(t_{comm}, t_{cin,info})$. The responding agent then examines the arguments by
 959 invoking the predicate *update*. Since *update* is an action, it is self-terminating and
 960 $Equal(t'_I, [update(\Gamma'_i)])$. If the arguments are sufficient, then *update* results in
 961 $D : D_r(Do(r, \rho), t_D)$; otherwise, the negotiation continues. If at the meantime the
 962 agent believes that the pace of the negotiation is not slow or that the task is
 963 discrete, then it continues to ask for more information from the initiating agent.
 964 Note that in our protocol, a responding agent can only make a counter-offer when
 965 the requested task is non-discrete. There are two conditions that prompt a
 966 responding agent to counter-offer: (1) when the initiating agent does not have any
 967 more arguments (as discussed later), or (2) when the pace of the negotiation is
 968 slow. The predicate *slow* measures the pace of a negotiation. The predicates *update*,
 969 *slow*, and *discrete* are part of our real-time enabling functional predicates and will
 970 be discussed further in Section 3.4.

972 *Agree 2* When a responding agent (r) receives arguments from an initiating agent
 973 (i) for a requested task ρ , it processes the arguments to update evidence support for
 974 the task. If the support is enough, then it agrees to perform ρ .

$$\begin{aligned}
 & C_{in} : info(i, r, Do(r, \rho), t_{cin,info}) \wedge I : I_r(succeed(Negotiate(i, Do(r, \rho))), t'_I) \wedge \\
 & I : I_r(update(\Gamma'_i), t''_I) \wedge D : D_r(Do(r, \rho), t_D) \Rightarrow C_{out} : agree(r, i, Do(r, \rho), t_{cout,agree}) \wedge \\
 & I : I_i(\neg Negotiate(i, Do(r, \rho)), t'''_I) \wedge I : I_i(\neg succeed(Negotiate(i, Do(r, \rho))), t'''_I) \\
 & \wedge success
 \end{aligned}$$

976 where

$$\begin{aligned}
 & Meets(t_{cin,info}, t'_I) \wedge During(t_{cin,info}, t'_I) \wedge During(t''_I, t'_I) \wedge Meets(t''_I, t_D) \wedge \\
 & During(t_D, t'_I) \wedge During(t_{cout,agree}, t_D) \wedge During(t_{cout,agree}, t'''_I) \wedge Meets(t'_I, t'''_I).
 \end{aligned}$$

978 This axiom is a counterpart of *More_info*. If it turns out that the arguments are
 980 sufficient, then the responding agent agrees to the request. It uses the communicative
 981 act *agree* to compose and send an AGREE-type message to the initiating agent. The
 982 negotiation ends with a *success* tag. Also, the agent also stops intending to negotiate
 983 and to negotiate successfully.

984 *Counter 1* When a responding agent (r) receives arguments from an initiating agent
 985 (i) for a requested task ρ , it processes the arguments to update evidence support for
 986 the task. If the support is not enough, the pace of the negotiation is slow, and ρ
 987 involves non-discrete resource/task, then r makes a counter-offer (ρ').

$$\begin{aligned}
& C_{in} : info(i, r, Do(r, \rho), t_{cin,info}) \wedge I : I_r(succeed(Negotiate(i, Do(r, \rho))), t'_I) \wedge \\
& I : I_r(update(\Gamma'_i, t'_I) \wedge \neg \exists D : D_r(Do(r, \rho), t_D) \wedge B : \\
& B_r(slow(Negotiate(i, Do(r, \rho))), t_B) \wedge \\
& B : B_r(\neg discrete(\rho), t'_B) \Rightarrow C_{out} : counter(r, i, Do(r, \rho), t_{cout,counter})
\end{aligned}$$

989 where

$$\begin{aligned}
& Meets(t_{cin,info}, t'_I) \wedge During(t_{cin,info}, t'_I) \wedge During(t''_I, t'_I) \wedge Meets(t'_I, t_D) \wedge \\
& During(t_D, t'_I) \wedge Finishes(t_D, t_B) \wedge Finishes(t_D, t'_B) \wedge During(t_{cout,counter}, t_D).
\end{aligned}$$

991 The communicative act *counter* encapsulates the acts of finding a counter-offer (ρ'),
992 composing a COUNTER-type message, and sending the message to the initiating
993 agent. The counter-offer ρ' is stored in the $\langle contents \rangle$ of the message. This is a
994 companion rule to *More_info* as discussed above. If the negotiation is off pace, then
995 instead of asking for more information/arguments, the responding agent counter-
996 offers. This is motivated by the intention of the agent to achieve a successful outcome
997 to the negotiation. Note that, as mentioned earlier, when a task involves a discrete
998 resource (see also Section 3.4.3), our protocol does not allow for a counter-offer.
999 Thus, the *More_info* rule overwrites the *Counter 1* rule. Further, even though the
1000 responding agent makes a counter-offer, it does not have the desire to perform the
1001 task counter-offered. It only has the desire to do so after the initiating agent agrees to
1002 it. This is represented later in axiom *Success 3*.

1003 Note that the motivation behind a counter-offer is to speed up the pace of the
1004 negotiation or as a last-ditch effort to salvage a failing negotiation. We do not
1005 perform counter-offer as part of the normal interaction – to evaluate and re-plan
1006 proposals at each negotiation step would have slowed down our negotiations and
1007 that is not applicable to a real-time problem.

1008 *Stop* When a responding agent (r) is notified by an initiating agent (i) that it has no
1009 more arguments for a requested task ρ , and the task is discrete, the agent r stops the
1010 negotiation, and the negotiation fails.

$$\begin{aligned}
& C_{in} : info_null(i, r, Do(r, \rho), t_{cin,info_null}) \wedge I : I_r(succeed(Negotiate(i, Do(r, \rho))), t'_I) \wedge \\
& I : I_r(update(\Gamma'_i(Do(r, \rho)), t''_I) \wedge \neg \exists D : D_r(Do(r, \rho), t_D) \wedge B : B_r(discrete(\rho), t'_B) \\
& \Rightarrow C_{out} : stop(r, i, Do(r, \rho), t_{cout,stop}) \wedge I : I_r(\neg Negotiate(i, Do(r, \rho)), t'''_I) \wedge \\
& I : I_r(\neg succeed(Negotiate(i, Do(r, \rho))), t''_I) \wedge stop
\end{aligned}$$

1012 where

$$\begin{aligned}
& Meets(t_{cin,info_null}, t'_I) \wedge During(t_{cin,info_null}, t'_I) \wedge During(t''_I, t'_I) \\
& \wedge Meets(t'_I, t_D) \wedge During(t_D, t'_I) \wedge \\
& Finishes(t_D, t_B) \wedge Finishes(t_D, t'_B) \wedge During(t_{cout,stop}, t_D) \wedge Meets(t'_I, t'''_I) \\
& \wedge Starts(t_{cout,stop}, t''_I).
\end{aligned}$$

1014 The communicative act *info_null* is the encapsulation of receiving and parsing an
 1015 INFO_NULL-type message from the initiating agent, while the communicative act
 1016 *stop* encapsulates the actions of composing a STOP-type message and sending the
 1017 message to the initiating agent. This is when the responding agent gives up on the
 1018 negotiation, after being informed that no more arguments are on the way. As a
 1019 result, it no longer intends to negotiate, and it opts out by informing the initiating
 1020 agent as a responsible gesture.
 1021 The proposition *stop* indicates the failure of a negotiation because the responding
 1022 agent is not convinced to perform the requested task.

1023 *Counter 2* When a responding agent (*r*) is notified by an initiating agent (*i*) that it
 1024 has no more arguments for a requested task ρ , and the task is discrete, the agent *r*
 1025 makes a counter-offer (ρ').

$$\begin{aligned} C_{in} &: info_null(i, r, Do(r, \rho), t_{cin,info_null}) \wedge I : I_r(succeed(Negotiate(i, Do(r, \rho))), t'_I) \wedge \\ & I : I_r(update(\Gamma'_i, t''_I) \wedge \neg \exists D : D_r(Do(r, \rho), t_D) \wedge B : B_r(\neg discrete(\rho), t'_B) \\ & \Rightarrow C_{out} : counter(r, i, Do(r, \rho), t_{cout,counter}) \end{aligned}$$

1027 where

$$\begin{aligned} & Meets(t_{cin,info}, t'_I) \wedge During(t_{cin,info}, t'_I) \wedge During(t'_I, t'_I) \wedge Meets(t'_I, t_D) \wedge \\ & During(t_D, t'_I) \wedge Finishes(t_D, t_B) \wedge Finishes(t_D, t'_B) \wedge During(t_{cout,counter}, t_D). \end{aligned}$$

1029 This axiom is a counterpart of *Stop*, and closely resembles *Counter 1*. Driven by the
 1030 intention to succeed in the negotiation, the responding agent, after being notified of
 1031 no more arguments coming in from the initiating agent, voluntarily makes a counter-
 1032 offer if the task is non-discrete.

1033 *Failure 3R* This axiom is similar to *Failure 3I*. See Appendix A.

1034 *Failure 4R* This axiom is similar to *Failure 4I*. See Appendix A.

1035 *Abort R* This axiom is similar to *Abort I*. See Appendix A.

1036 *Out_of_time R* This axiom is similar to *Out_of_time I*. See Appendix A.

1037 *Success 3* When a responding agent (*r*) receives an AGREE message from an
 1038 initiating agent (*i*) to its counter-offer (ρ'), the responding agent desires to perform
 1039 the counter-offer, and the negotiation ends with success.

$$\begin{aligned} C_{in} &: agree(i, r, Do(r, \rho), t_{cin,agree}) \wedge I : I_r(succeed(Negotiate(i, Do(r, \rho))), t'_I) \\ & \Rightarrow D : D_r(Do(r, \rho'), t_D) \wedge I : I_i(\neg Negotiate(i, Do(r, \rho)), t'_I) \wedge \\ & I : I_i(\neg succeed(Negotiate(i, Do(r, \rho))), t'_I) \wedge success \end{aligned}$$

1041 where $Finishes(t_{cin,agree}, t'_I) \wedge Meets(t'_I, t_D) \wedge Meets(t'_I, t'_I)$. The communicative act
 1042 *agree* predicate encapsulates the acts of receiving and parsing an AGREE-type
 1043 message, in which $\langle contents \rangle$ holds the information regarding ρ' . This rule says that if

1044 the initiating agent agrees to the counter-offer, then the responding agent (1) desires
 1045 to perform the counter-offered task and (2) intends no longer to continue with the
 1046 negotiation regarding the originally requested task.

1047 *No_response R* This is similar to *No_response I*. See Appendix A.

1048 Figure 5 shows the time lines of the initiating agent's behavior when faced with an
 1049 outright rejection or agreement. These are the simple cases of the axioms above. See
 1050 for example that when the responding agent agrees to a request, it extends the desire
 1051 to do the requested task.

1052 Figure 6 shows the negotiation process, from the responding agent's point of view
 1053 once it agrees to negotiate. Similar to Figure 4, Figure 6 allows us to explicitly
 1054 describe the temporal relationships among the BDI states and the communicative
 1055 acts. That description, in turn, allows us to guarantee the behavior of the negotiation
 1056 and to fine-tune its efficiency. See Section 3.3.7. Note that when a responding agent
 1057 has a desire to perform the requested task, it will monitor its negotiation progress, as
 1058 triggered by the responding agent's intention to update its belief of whether the task
 1059 requested should be agreed to. Also, when the task is non-discrete and the progress is
 1060 slow, the responding agent will make a counter-offer.

1061 3.4. Functional predicates

1062 In this section, we describe the functional predicates mentioned in the previous
 1063 section that present the logical framework for the rules of encounter between two
 1064 agents. These predicates are the infrastructure to our real-time negotiation protocol.
 1065 To simplify the discussion, we have touched upon 11 communicative acts such as
 1066 *initiate*, *respond*, *no_go*, *agree*, etc. (discussed in Section 3.3) and six negotiation-
 1067 related functional predicates: *slow*, *time*, *discrete*, *no_response*, *acceptable*, and *up-*
 1068 *date*. Here we will elaborate further on the six predicates as they are an integral part
 1069 that enables the real-time negotiation between the agents.

1070 **3.4.1. Slow** This predicate takes the form of $slow(action)$, and given an action (or
 1071 a task), it measures the pace of the action and returns true or false. An action has
 1072 two temporal intervals: the actual real-time interval, $[action]$, and the planned/pre-
 1073 dicted interval, $\|action\|$. Suppose that $[action]$ has a duration between $t_{s,action}$ and
 1074 $t_{f,action}$, and the set of states as a result of the action is $S_{action} = \{S_{0,action}, \dots, S_{N,action}\}$.




Figure 5. Time lines of the responding agent's behavior when faced with an outright rejection or agreement.

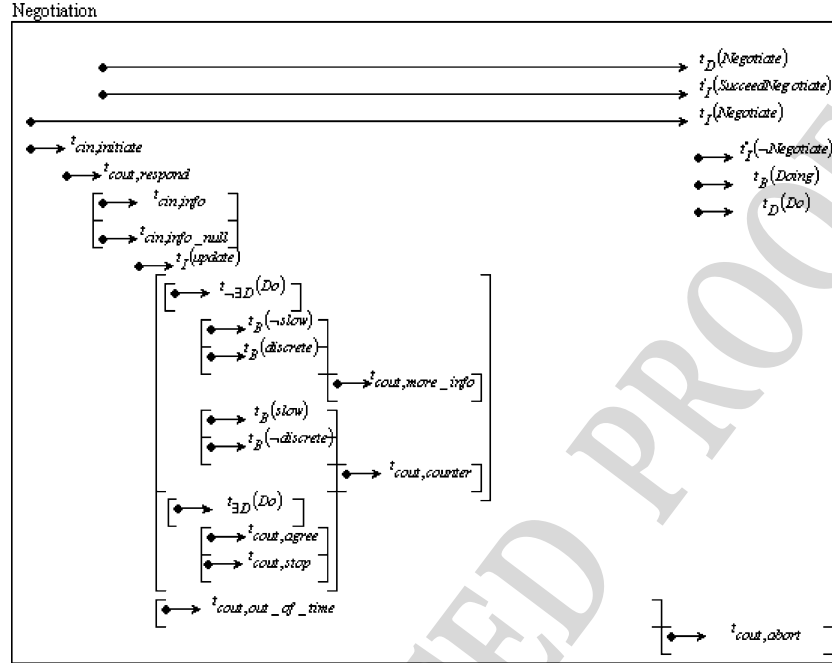


Figure 6. Time lines of the negotiation from the responding agent's point of view (a simplified version). Temporal intervals under brackets are options. Note that once the responding agent sends out a RESPOND message to the initiating agent, it is committed to negotiate successfully.

1075 Each state, $s_{n,action}$, holds during an actual interval $[s_{n,action}]$ such that (1)
 1076 *During*($[s_{n,action}], [action]$), (2) *Overlaps*($[s_{n,action}], [s_{n+1,action}]$), and (3) the temporal
 1077 interval of two such states has a duration between the start of $s_{n,action}$ and the latest
 1078 finish time of the two states, $[s_{n,action} s_{n+1,action}]$. Similarly, we can obtain
 1079 $\|s_{n,action} s_{n+1,action}\|$

1080 When *slow*(*action*) is invoked, it retrieves the current state of the action,
 1081 $s_{current,action} \in S_{action}$. If

$$\frac{[s_{0,action} s_{current,action}]}{[action]} > \frac{\|s_{0,action} s_{current,action}\|}{\|action\|}$$

1083 is true, then *slow*(*action*) returns true; otherwise, it returns false. Note that this
 1084 predicate is binary since we use it to trigger a *counter_offer* act. A more general
 1085 approach is to use a *degree* of slowness that would not only trigger a *counter_offer*
 1086 but also dictate how conceding the responding agent should be.

1087 **3.4.2. Time** This predicate takes the form of *time*(*action*) and indicates whether
 1088 an agent has run out of the allocated time for the *action*. Borrowing the notations
 1089 from Section 3.4.1, suppose we have $[action]$ and $\|action\|$. When an agent invokes
 1090 *time*(*action*), the predicate measures the time elapsed so far, $t_{elapsed} = t_{current} - t_{s,action}$.
 1091 If $t_{elapsed} \geq \|action\|$, then *time*(*action*) returns true; otherwise it returns false.

1092 **3.4.3. Discrete** A discrete task is when the task cannot be broken up or attenu-
 1093 ated. For example, if the initiating agent asks the responding agent to turn on a
 1094 sensor, then the responding agent may only respond with yes or no. However, if the
 1095 initiating agent asks the responding agent to turn on a sensor in five seconds, then
 1096 the responding agent, in addition to yes or no, may also counter with “yes, but in 10
 1097 seconds”. The introduction of the time factor makes the task a non-discrete one,
 1098 allowing the responding agent to counter-offer. In our negotiation protocol, a task ρ
 1099 may be qualified by time and resource amount. The predicate *discrete* returns true if
 1100 both qualities are absent, otherwise, it returns false. Note that we intentionally leave
 1101 out the qualification of tasks in Section 3.3 to simplify our discussions.

1102 **3.4.4. No_response** This predicate takes one argument – the agent from which the
 1103 current agent is waiting for a message. In general, when an agent i invokes
 1104 *no_response*(r) on another agent r , it is after agent i has performed a communicative
 1105 act, C . Hence, *Before*($[C], [no_response(r)]$). If agent i does not receive a response in
 1106 the interval t_{window} (and *During*(t_{window}, t_I) given that the intention I is to negotiate),
 1107 then *no_response*(r) returns true; otherwise, it returns false.

1108 **3.4.5. Acceptable** This predicate is only invoked by an initiating agent, I , and
 1109 takes as argument a task. Suppose the agent i desires to achieve goal G . To achieve
 1110 G , there is a set of subtasks $\vec{\rho} = \{\rho_1, \rho_2, \dots, \rho_{|C_i(\vec{\rho})|}\}$, and as mentioned in Section
 1111 3.3.5, a coalition, $C_i(\vec{\rho})$, exists to distribute the subtasks among the coalition
 1112 members. As a result,

$$D_i(G, t_D) \Rightarrow D_i(Do(C_i(\vec{\rho}), \vec{\rho}), t_D),$$

1114 where $D_i(Do(C_i(\vec{\rho}), \vec{\rho}), t_D) = \left\{ D_i(Do(r_1, \rho_1), t_{D1}), \dots, D_i(Do(r_{|C_i(\vec{\rho})|}, \rho_{|C_i(\vec{\rho})|}), t_{D|C_i(\vec{\rho})|}) \right\}$.
 1115 As the agent progresses in real time, G may become G' such that
 1116 $D_i(G', t'_D) \Rightarrow D_i(Do(C_i(\vec{\rho}'), \vec{\rho}'), t'_D)$ where *During*(t'_D, t_D). For example, an agent has
 1117 received commitments for some negotiated resources, so it no longer desires the
 1118 original set of resources if asked for. So, when the initiating agent receives a counter
 1119 offer ρ' from the responding agent, it invokes *acceptable*(ρ') to compare the counter-
 1120 offered task ρ' with the corresponding ρ'_r where $D_i(Do(r, \rho'_r), t'_D)$. If $\rho' \in \rho'_r$ then
 1121 *acceptable*(ρ') returns true; otherwise, it returns false.

1122 Whether a resource request is agreeable by a responding agent and whether a
 1123 counter-offer is acceptable by an initiating agent depend on how the request is
 1124 determined in the first place by the initiating agent. And this determination of which
 1125 and the amount of resources to request from which responding agents is task allo-
 1126 cation in our coalition. Note that in [50], we detail how task allocation is carried out
 1127 in our coalition. Briefly, the initiating agent makes use of the potential utilities to
 1128 carry out task allocations and assignments. Based on the overall potential utility of
 1129 the initial coalition, the initiator may want to lower its demands to improve the
 1130 chance of forming a coalition. By the same token, if the potential utility of a candi-
 1131 date is high, then the initiator may want to increase its demand with that candi-
 1132 date. We are currently investigating various algorithms such as greedy, lazy, worried,
 1133 and weary. An initiating agent becomes greedy during a negotiation when (1) it
 1134 tries to minimize its own rationalization and computing process, (2) it selects the

1135 candidate with the higher overall utility values to approach hoping for a successful
1136 negotiation, (3) it cares mostly about high-priority tasks, (4) it tries to maximize its
1137 chance of getting a particular task done – by including sub-utilities in the focused
1138 utility evaluation, and (5) it hopes to shift its responsibility (partially) to the candi-
1139 dates via successful negotiations – expecting the candidates to spawn their own
1140 coalitions to help respond to the problem at hand. In a lazy algorithm, the initiator
1141 prefers to concentrate its effort on a few candidates, as it does not want to spend
1142 resources or time on too many. In a worried algorithm, the agent asks for more than
1143 it needs to ensure, that if some of the candidates refuse to help, it will get what it
1144 needs. This translates to insurance policies. Finally, in a weary algorithm, the ini-
1145 tiator prefers not to upset a candidate – especially one that has a high uniqueness –
1146 by being over-demanding. This leads to demand caps that make sure that an addi-
1147 tional demand does not hurt the negotiation.

1148 **3.4.6. Update** This functional predicate is invoked only by a responding agent *r*,
1149 as previously mentioned in an auxiliary to the *More_info* axiom. See the discussion in
1150 Section 3.3.7.

4. Implementation

1152 The driving application for our system is multi-sensor target tracking, a distributed
1153 resource allocation and constraint satisfaction problem. The objective is to track as
1154 many targets as possible and as accurately as possible using a network of sensors.
1155 Each sensor has a set of consumable resources, such as beam-seconds (the amount of
1156 time a sensor is active), battery power, and communication channels, that each
1157 sensor desires to utilize efficiently. Each sensor is at a fixed physical location and, as a
1158 target passes through its coverage area, it has to collaborate with neighboring sen-
1159 sors to triangulate their measurements to obtain an accurate estimate of the position
1160 and velocity of the target. As more targets appear in the environment, the sensors
1161 need to decide which ones to track, when to track them, and when not to track them,
1162 always being aware of the status and usage of sensor resources. Each sensor can at
1163 any time scan one of three sectors, each covering a 120-degree swath. Sensors are
1164 connected to a network of CPU platforms on which the agents controlling each
1165 sensor reside. The physical sensors are 9.35 GHz Doppler MTI radars that com-
1166 municate using a 900 MHz wireless, RF transmitter. The agents (and sensors) must
1167 communicate over an eight-channel RF link, leading to potential channel jamming
1168 and lost messages. Our agents may reside on the same CPU platform or different
1169 platforms.

1170 We are using a tracker module that receives radar measurements and returns to
1171 the agent the estimated position and velocity of a target. These estimates are im-
1172 proved with more measurements from different sensors within a short time interval,
1173 since this allows better target triangulation. The tracker software is not an agent, and
1174 simply implements a target tracking algorithm. The communication between the
1175 tracker module and an agents is via socket connections.

1176 We have implemented the real-time, case-based negotiating agents described in the
1177 previous sections and tested them using real sensors and targets moving in a physical
1178 environment. The agents exhibit all of the behavior described: they use CBR to select
1179 and adapt a negotiation strategy, use a RTSS to request CPU resources and to have
1180 time and system awareness, negotiate for radar use, and learn the new negotiation
1181 strategies they have developed. See [50] for more details.

1182 4.1. Agents and environments

1183 We have implemented our multi-agent system in the C++ programming language.
1184 Each agent has $3 + n$ threads: (1) a core thread that performs the reasoning, message
1185 checking, task handling chores of the agent and thus is always active, (2) a com-
1186 munication thread that is responsible for polling for incoming messages and sending
1187 out messages and thus is always active, (3) an execution thread that performs sensor-
1188 related tasks such as calibration, target searching, and tracking, and thus is some-
1189 times active and sometimes dormant, and (4) n negotiation threads that each waits to
1190 be awoken to perform a negotiation and goes back to a dormant state after the
1191 negotiation is over. This setup allows an agent to carry out various lines of tasks
1192 concurrently. It also allows an agent to conduct multiple negotiations in parallel.
1193 Each negotiation thread can be an initiating thread or a responding thread,
1194 depending on the dynamic, real-time instructions given by the core thread.

1195 The current implementation of our agents is able to (1) detect a target, (2) form a
1196 coalition, perform CBR to determine its negotiation strategy, initiate or respond to
1197 negotiations, (3) argue to persuade its partner to perform a task or reasons to
1198 whether to agree to perform a task, (4) monitor its own status such as its sensor,
1199 noise, tasks, and CPU resource usage, (5) interact with either a software simulation
1200 or the actual physical hardware setup, and (6) obtain real-time data from the
1201 operating system supporting its execution. More importantly, each agent is auton-
1202 omous and can sense and react to real-time events in the environment. Moreover,
1203 there is no hierarchy within the multi-agent system – all agents are peers.

1204 We have also implemented the complete real-time argumentative negotiation
1205 protocol in the negotiator module of an agent. With the formalisms encoded, the
1206 negotiator conducts a negotiation with high efficiency and autonomy. We have
1207 implemented all communicative acts: parsing messages and converting them to belief
1208 states and converting belief states and composing messages out of them. We have
1209 also implemented the belief, desire, and intention states as functions, procedures, and
1210 clauses. As for the temporal definitions and constraints of those states, we have
1211 implemented recursive mutexes (semaphore-like designs) to manage read/write
1212 accesses: when to acquire the value of a state, when to release a lock on the value of a
1213 state, when must a state be ready in order for a certain task to start, and so on. For
1214 example, an agent may be tracking a target and negotiating at the same time. While
1215 tracking, the agent may realize that the target is no longer visible. This directly
1216 affects the on-going negotiation since now the agent's sensor has become available,
1217 leading to a lower threshold, for example, of a counter-offer. But, the negotiator
1218 module of a negotiation thread cannot afford to constantly check the states of the

1219 tracking. It does so occasionally and only when it is necessary, and can only be
1220 interrupted at certain points over the course of the negotiation. This is dictated by
1221 the temporal definitions of the states found in Section 3.

1222 4.2. *Real-time scheduling service*

1223 We have implemented a RTSS in the C programming language, on top of the KU
1224 Real-Time system (KURT) [53] that adds real-time functionality to Linux. First, the
1225 RTSS provides an interface between the agents and the system timers, allowing
1226 agents to (1) query the operating system about the current time; (2) ask the RTSS to
1227 notify them after the passage of certain length of time; and (3) ask the RTSS to ping
1228 them at fixed time intervals. This allows agents to know when to, for example,
1229 conclude a negotiation process or turn on a radar sector. Second, the agents may ask
1230 the RTSS to notify them when certain system-level events occur, such as process
1231 threads being activated, or communication messages going out or coming into the
1232 system. Third, the agents can ask the RTSS to allocate them a percentage of the CPU
1233 for each one of their threads (such as the ones controlling the radar and tracking or
1234 the ones used in negotiations) and to schedule this allocation within an interval of
1235 time. This RTSS allows an agent to monitor the progress of its own negotiations and
1236 the usage status of its allocated CPU resource.

1237 4.3. *Case-based argumentative negotiation*

1238 We have implemented the CBR Manager to maintain the case base of an agent. The
1239 implementation includes similarity-based retrieval, both difference- and outcome-
1240 driven adaptations, and the incremental and refinement learning. The CBR Manager
1241 retrieves a case for each negotiation. A case provides the appropriate negotiation
1242 strategy for the agent. For a responding agent, the negotiation strategy also includes
1243 the persuasion threshold. When the responding agent receives an argument from the
1244 initiating agent supporting its request, it measures the support of this argument using
1245 CLIPS⁸ rules and compares the support collected so far against the persuasion
1246 threshold. We have implemented the entire negotiation protocol, as depicted in
1247 Figure 1, into each negotiation thread (the number of threads in an agent is static for
1248 a specific agent implementation. In our experiments we used two negotiation threads
1249 per agent). Each thread is capable of monitoring the pace of its own negotiation,
1250 retrieving messages via the communication thread, parsing incoming messages,
1251 making decisions and reasoning, composing an outgoing message and sending
1252 messages via the communication thread. Each thread is autonomous in a way that
1253 the core thread of the agent does not have to tell the thread how to conduct a
1254 negotiation once it has gotten underway.

1255 Each thread forks off a child process that automatically invokes CLIPS. The com-
1256 munication between a negotiation thread and its child CLIPS process is through pipes.
1257 After receiving an acknowledge signal from the CLIPS child process, the negotiation
1258 thread informs the core thread that it is ready to accept a negotiation task and waits.
1259 When it finally receives an activation signal, the negotiation thread downloads the

1260 relevant information regarding the negotiation task. From the information, the thread
 1261 decides its identity – either an initiating thread or a responding thread. Then the
 1262 negotiation thread negotiates following the negotiation protocol described in Section
 1263 3. Once the negotiation is done, the thread updates its status and waits for a signal from
 1264 the core thread before resetting itself for the next negotiation task. Meanwhile, the core
 1265 thread of the agent periodically checks the status of the active negotiation threads. If a
 1266 negotiation is completed, the core thread downloads the updated data and signals the
 1267 negotiation thread that it is okay to reset.

1268 In our case-based strategy selection approach, a negotiation strategy dictates a set
 1269 of tactics for a negotiation. For example, an initiating agent needs to know which
 1270 arguments are more important to send over first to the responding agent. We use
 1271 CBR to help us determine that. When an agent encounters a negotiation problem, it
 1272 searches its casebase for the most similar case, in which the problem description in
 1273 that case resembles the current negotiation problem. Then, based on the differences
 1274 between the two problem descriptions, the CBR module of the agent performs an
 1275 adaptation on the solution. The modified solution becomes the negotiation strategy
 1276 (details are provided in [51]).

1277 4.4. Real-time enabling functional predicates

1278 In Section 3.4, we presented the logical model of our real-time enabling functional
 1279 predicates. Here, we describe the implementation that, even though is domain- and
 1280 application-specific, may serve as a useful example to other designs of the predicates.
 1281 In this section, we also describe how our agents make a counter-offer in real-time.
 1282 Note that in our implementation, we employ CBR to derive a negotiation strategy
 1283 for an agent for each negotiation task. A case has a set of belief states (situated input
 1284 parameters), a set of desires (a parametric negotiation strategy), and the outcome of
 1285 the negotiation. In addition, in our agent design, a negotiation is handled by one of
 1286 the negotiation threads that an agent dispatches. So, in the following, we will use the
 1287 term “negotiation thread” quite often and make use of the belief and desire states.

1288 **4.4.1. Slow** In a case, the desires include the number of negotiation steps allowed
 1289 and the time allowed. That is, a negotiation thread desires to complete a particular
 1290 negotiation in n iterations and s seconds. A small n means that fewer messages are
 1291 exchanged and the negotiation may avoid incurring too much overhead cost per
 1292 transmission. A small s means that the negotiation is to be completed in a short time.

1293 Suppose we denote the number of negotiation steps allowed as $step_{allowed}$, the time
 1294 allowed as $time_{allowed}$, the number of steps performed so far as $step_{sofar}$, and the time
 1295 elapsed so far as $time_{sofar}$. We define the *slow* predicate for a responding agent α 's
 1296 negotiation (intending to achieve the desire for performing a requested task ρ) as

$$slow(I_{\alpha}(D_{\alpha}(Do(\alpha, \rho))), t) = \left(\frac{time_{sofar}}{step_{sofar}} > \frac{time_{allowed}}{step_{allowed}} \right)$$

1298 This definition is an example of the logical model discussed in Section 3.4.1.

1299 **4.4.2. Time** To implement this predicate, a negotiation thread registers its process
 1300 ID with a real-time system-level service and makes use of a time-based notification
 1301 mechanism. A negotiation asks the notification mechanism to signal the thread after
 1302 s seconds. One unique characteristic of the mechanism is that the negotiation thread,
 1303 after registration, may find out how much of the s seconds has elapsed after the
 1304 notification was first registered (for example, 25, 50, 75, 100%) by consulting the
 1305 notification flag: t_{flag} . The value of s is determined by the $time_{allowed}$ of the desire
 1306 states of a case.

1307 In our current design, we define the *time* predicate for an agent α 's negotiation as

$$time(negotiation) = (t_{flag} < 1)$$

1309 When $t_{flag} = 1$, that means the time elapsed has reached 100% of $time_{allowed}$.

1310 **4.4.3. Discrete** First, we denote the set of discrete tasks Θ_{dis} and the set of non-
 1311 discrete tasks Θ_{con} . Then we define the *discrete* predicate of a requested task ρ as

$$discrete(\rho) = (\rho_{request} \in \Theta_{dis})$$

1313 where $\rho_{request} \in \rho$ is part of the requested task.

1314 **4.4.4. No_response** Our implementation is the following: After an agent sends out
 1315 a message, it polls its message queue for a response before moving on. After $t_{polling}$
 1316 seconds, if the negotiation thread receives no messages from a particular negotiation
 1317 partner, then *no_response* returns true. If there is a consistently typed message, then
 1318 the negotiation thread reacts to it based on the negotiation protocol outlined in
 1319 Section 3.3.

1320 **4.4.5. Acceptable** This predicate is used only by an initiating agent when it re-
 1321 ceives a counter-offer from a responding agent and refers to non-discrete (continu-
 1322 ous) tasks, θ_{con} . Let us denote a continuous task as $\vec{r} \in R_{con}$. There are three key
 1323 parameters in ρ_{con} : $\rho_{con} = \{\rho_{con,name}, \rho_{con,res}, \rho_{con,amount}\}$ where $\rho_{con,name}$ is the name of
 1324 the task, $\rho_{con,res}$ designates the resource involved in the task, and $\rho_{con,amount}$ indicates
 1325 the amount of the resource involved.

1326 In our design, when an agent α realizes $B : B_{\alpha}(\rho_{con,amount}^{needed})$ where $\rho_{con,amount}^{needed}$ is
 1327 non-zero, it initiates negotiations – each with a different $\rho_{con,amount}^{requested}$, to the coalition
 1328 members – attempting to obtain enough $\rho_{con,amount}$ from the members to meet
 1329 $\rho_{con,amount}^{needed}$, i.e., $\sum_{C_x(\vec{p})} \rho_{con,amount}^{requested} \geq \rho_{con,amount}^{needed}$. At each agent cycle, the agent α updates its
 1330 $B : B_{\alpha}(\rho_{con,amount}^{needed})$ and $B : B_{\alpha}(\sum_{C_x(\vec{p})} \rho_{con,amount}^{requested})$. If $B : B_{\alpha}(\rho_{con,amount}^{needed} = 0)$ then it has
 1331 achieved its target, and it can abort all current negotiations associated with that
 1332 particular resource. The process checks (1) the current usage, (2) the anticipated
 1333 usage, (3) the current allocation, and (4) the agreed additional allocation. As
 1334 $\rho_{con,amount}^{needed}$ gets smaller, the remaining negotiations become less demanding in their
 1335 requests, and vice versa. As each negotiation completes gradually, $\sum_{C_x(\vec{p})} \rho_{con,amount}^{requested}$
 1336 changes.

1337 Suppose that the k th negotiation thread of the agent α is negotiating to obtain
 1338 $\rho_{con,amount,k}^{requested}$ from a responding agent and the responding agent has just counter-offered
 1339 ρ'_k with an offered amount of $\rho'_{con,amount,k}$. Then the acceptability of the counter-offer ρ'_k
 1340 is defined as

$$acceptable(\rho'_k) = \left(\sum_{C_2(\bar{\rho})} \rho_{con,amount}^{requested} - \rho_{con,amount,k}^{requested} + \rho'_{con,amount,k} \geq \rho_{con,amount}^{needed} \right)$$

1342

1343 **4.4.6. Update** The *update* predicate is used when a responding agent receives
 1344 information or arguments from an initiating agent. The objective of this predicate is to
 1345 find out whether the evidence support for a requested task is convincing enough for the
 1346 responding agent to perform it. One key parameter of a negotiation strategy of a
 1347 responding agent, r , is the persuasion threshold, $T_{persuasion,Do(r,\rho)}$, for a requested task ρ .
 1348 This is a value created by the responding agent r for ρ ; to agree to $Do(r,\rho)$, the argu-
 1349 ments sent by the initiating agent must provide for ρ that is greater than $T_{persuasion,\rho}$.

1350 Suppose we denote the evidence support for $Do(r,\rho)$, with a persuasion threshold,
 1351 $T_{persuasion,Do(r,\rho)}$, at temporal interval t as $Support(\rho, T_{persuasion,Do(r,\rho)}, t)$, and
 1352 $Support(\rho, T_{persuasion,Do(r,\rho)}, time0) = 0$. The objective of the initiating agent is to
 1353 obtain $Support(\rho, T_{persuasion,Do(r,\rho)}, t) \geq T_{persuasion,Do(r,\rho)}$ in order to convince the
 1354 responding agent to perform the requested task. Thus we have the following axiom:

1355 *Axiom desire to do* If a responding agent r is negotiating with an initiating agent i
 1356 regarding a requested task ρ , and at temporal interval t , it has
 1357 $Support(\rho, T_{persuasion,Do(r,\rho)}, t) \geq T_{persuasion,Do(r,\rho)}$ then $D : D_r(Do(r,\rho), t)$. (This sup-
 1358 plements the *More_info*, *Agree 2*, *Counter 1*, *Stop*, and *Counter 2* axioms discussed in
 1359 Section 3.3.)

1360 When arguments are received by the responding agent, the support value changes
 1361 based on the following definitions of an agent behavioral model:

- 1362 (1) Agents that have cooperated before will tend to cooperate again.
- 1363 (2) A responding agent is willing to trust an initiating agent's perception.
- 1364 (3) An agent is more inclined to help another agent if that another agent has relied
 1365 on the agent for help before.
- 1366 (4) An agent is more inclined to help another agent if it knows that it is one of the
 1367 few possible solutions to the requested task.
- 1368 (5) An agent is more inclined to help another agent if it knows that that another
 1369 agent is busy.

1370 Note that the persuasion threshold and the evidence support work together to *imi-*
 1371 *tate a joint intention*⁹ between the two negotiating agents. On one hand, the
 1372 responding agent is helpful and desires to help the initiating agent, but it also intends
 1373 to help when it is worthwhile. This implies cooperativeness with a touch of selfish-
 1374 ness in a local sense that translates into global optimization of resource allocation.
 1375 To make sure that the requested task is worthwhile to do, the responding agent uses
 1376 a persuasion threshold, derived from its past experience and its current status. On the
 1377 other hand, the initiating agent intends that the responding agent help with its task.

1378 It collects its belief states and sends over whatever it thinks are useful as arguments
 1379 for its intention. These arguments modify an evidence support value. Therefore,
 1380 Definition 4 (“An agent is more inclined to help another agent if that another agent
 1381 has relied on the agent for help before.”) merges the two intentions, seen from two
 1382 different perspectives, and the joint intention is to achieve a successful negotiation.
 1383 This deviates from the model proposed by Cohen and Levesque [9, 10] but if we view
 1384 *achieving a successful negotiation as a team action*, then both members of the team
 1385 (the two negotiating agents) are jointly committed to completing the so-called team
 1386 action, with the agent that performs the task believing that it performs the task while
 1387 the other agent believes that the same performs the task as well.

1388 We have implemented the *update* function as a CLIPS-based operation. Evaluation
 1389 heuristics are coded as CLIPS rules and arguments received from an initiating
 1390 agent are fed into the CLIPS process (associated with each negotiation thread) to re-
 1391 compute the evidence support for a requested task. The CLIPS process then sends
 1392 back the updated evidence support to its negotiation thread and waits for another set
 1393 of arguments. The negotiation thread and the CLIPS process communicate via a
 1394 synchronized pipe connection. The negotiation thread compares the updated evi-
 1395 dence support with its persuasion threshold. If the former is greater than or equal to
 1396 the latter, then the negotiation thread agrees to the requested task, composes a
 1397 message to notify the initiating agent of the deal, and completes its own negotiation.
 1398 The core thread of the agent then schedules the requested task in its activity.

1399 Here we list one example of the CLIPS rules used. The first rule indicates that if
 1400 the initiating agent has been helping the responding agent in the past with a rate of
 1401 66.7%, then the evidence support is incremented by 0.05 of that rate in favor of the
 1402 initiating agent with its current request.

```
(defruleworld-help-rate
(world(helpRate?y&:(>?y0.6667)))
=>
(bind?*evidenceSupport*(+(*0.05?y)?*evidenceSupport*))
(bind?*world-help-rate-count*(+?*world-help-rate-count*1)))
```

1404 In all, we have 15–20 rules in our CLIPS rulebase.

1405 **4.4.7. Counter offer** When dealing with continuous resources, the responding
 1406 agent has a persuasion function, modified by two parameters: (1) *kappa* – a will-
 1407 ingness factor, and (2) *beta* – a conceding factor, and bounded by the maximum
 1408 resource that it is willing to give up, $\rho_{con,amount}^{max}$. An agent can use any function¹⁰ to
 1409 express the continuous persuasion value; in our implementation, we examined two: a
 1410 linear and an exponential persuasion function.

1411 In our model, the linear persuasion function is:

$$P_{linear}(\tau) = \beta \cdot \tau + \kappa$$

1413 and, the exponential persuasion function is:

$$P_{exp}(\tau) = \kappa \cdot e^{\tau/\beta}$$

1415 The variable τ is the evidence support collected so far, i.e., $\tau = \text{Support}$
 1416 $(\rho, T_{\text{persuasion}, \text{Do}(r, \rho)}, t) \geq T_{\text{persuasion}, \text{Do}(r, \rho)}$. So, in the beginning, at t_s when the support
 1417 is zero, a responding agent is willing to give up $P_{\text{linear}}(0) = \text{kappa}$ or $P_{\text{exp}}(0) = \text{kappa}$.
 1418 That is why *kappa* is called the willingness factor. Then for $\tau > 0$, the conceding rate
 1419 depends on *beta*: the larger this value is, the more conceding the persuasion function
 1420 is. The key differences between our persuasion functions and others are (1) the
 1421 conceding and willingness factors are determined by past experiences, and adapted
 1422 to fit the current situation, (2) each function applies to an evidence support based on
 1423 arguments, (3) each function is implicitly bounded by *time_{allowed}* for a negotiation,
 1424 and (4) each function is bounded by $\rho_{\text{con}, \text{amount}}^{\text{max}}$, the maximum amount of a resource
 1425 that the agent is willing to give up.

1426 Thus, when a responding agent is about to make a counter-offer, it checks τ , and
 1427 the counter-offer, $\rho'_{\text{con}, \text{amount}}$ is given as

$$\rho'_{\text{con}, \text{amount}} = \begin{cases} P(\tau) \\ \rho_{\text{con}, \text{amount}}^{\text{max}} & \text{if } P(\tau) > \rho_{\text{con}, \text{amount}}^{\text{max}} \end{cases}$$

1429

1430 **5. Experiments and results** We conducted a set of experiments with a varying
 1431 number of real sensors and vehicles (as described in Section 4). The most complex
 1432 experiment involved eight 9.35 GHz Doppler MTI radars and two targets moving in
 1433 oval trajectories (Figure 7). During the experiments we used three different types of

COLOUR FIG.




Figure 7. Experimental setup. The two ovals represent the trajectory of the two targets. The sensors are labeled as S1 to S8. The X and Y -axes are the dimensions of the area in which the targets travel and the sensors are located.

1434 agents so as to compare the behavior and performance of other agent formalisms
1435 with that of the performance of negotiating, real-time agents described in this paper.
1436 Our hypothesis was that negotiating agents can track targets better since they can
1437 coordinate radar measurements and achieve better triangulation. The MTI sensors
1438 and the vehicles were simulated by the Radsim simulator version 2.08 [32]. Agents
1439 connect to Radsim and use a standard API to send control signals to their respective
1440 sensor platforms. Radsim, in turn, uses a model of the platforms' behavior to return,
1441 in the case of taking a measurement, a hypothetical value of amplitude and fre-
1442 quency based on the current locations of the targets. For each setup, we had three
1443 runs. Each run was run for 250,000 simulation ticks, representing 250 real-time
1444 seconds.

1445 Our negotiating agents were compared to (1) simple, reactive ones, and (2)
1446 negotiating agents that use a single, static negotiation strategy instead of an adap-
1447 tive, case-based one. The reactive agents operated as follows: When an agent de-
1448 tected a target, it started tracking it without attempting to form a coalition.
1449 Comparing our negotiating agents to reactive ones allowed us to establish a baseline
1450 performance in which the agents simply track without coordinating with other
1451 agents (as in the simple, reactive scenario) and study how coordinated tracking in a
1452 coalition (as in our negotiating agents) helped improve the tracking. When the agent
1453 no longer sensed the target, it returned to a "wait" state. The single-strategy nego-
1454 tiating agents used our negotiation protocol, but always used the same negotiation
1455 strategy. Comparing our negotiating agents between using a single, static strategy
1456 and using an adaptive, case-based one allowed us to investigate how agents equipped
1457 with learning capability are able to improve their negotiation and ultimately target-
1458 tracking efforts.

1459 We used two metrics of quality for the performance of the three techniques. First,
1460 we used the number of times that an agent selected the correct sensor sector to turn
1461 on, indicating that it was correctly tracking the target. A higher percentage of correct
1462 sector selections would also lead to better triangulation and tracking, implying better
1463 collaboration among the agents. Second, we used the *quality* of the measurements
1464 made. A target is visible by a sensor sector for a period of time. During this time, its
1465 visibility to the sensor changes as the target moves closer or further away from the
1466 sensor. The *exact time* at which a sensor makes a measurement establishes the *quality*
1467 *of the measurement as a function of the target visibility*. In other words, the first metric
1468 indicates whether the agent selects the correct sector to turn on at the write time, and
1469 the second metric indicates whether the appropriate sensor is tracking the target.

1470 To establish the optimal (or correct) sector that an agent should be looking at, we
1471 used the omniscient agent experiment performed by researchers at the University of
1472 South Carolina (<http://www.cse.sc.edu/~vargasje/targetshare/>). In their experiment
1473 the South Carolina researchers built a controller that knew the real location of the
1474 targets most of the time. The controller used that information to command a cen-
1475 tralized agent to change the sectors of the sensors so that they pointed most of the
1476 time in the direction of the best sector for tracking the target. Their experiment
1477 determined the best baseline performance that a centralized, omniscient agent could
1478 achieve and provided us with the best "ground truth" available sector. Further, the
1479 negotiating agents with CBR did significantly better in tracking Target 1 than in

1480 tracking Target 2. This is possibly due to resource contention for the sensors S4 and
 1481 S5, where Target 1 was first engaged and acquired by the sensors, which caused
 1482 Target 2 to be “off” in terms of the sensors’ timing.

1483 The results of this evaluation are summarized in the Table 1 below.

1484 Our real-time negotiation protocol with CBR does much better in creating
 1485 coalitions of the appropriate sectors to track a vehicle, especially for the left
 1486 target.

1487 Next, we used the *quality* of the measurements made as a proxy of the quality of
 1488 the coalition formation and the task execution. The *exact time* at which a sensor
 1489 makes a measurement establishes the *quality of the measurement as a function of the*
 1490 *target visibility*. An example is given in Figure 8: The red arc represents the visibility
 1491 of a target by a sector of a sensor (sensor named “node 1” and sector 1). The *x-axis* is
 1492 time in milliseconds. The *y-axis* is the quality represented by the sensitivity of the
 1493 target position to measurement errors. In practice average measurement errors are
 1494 on the order of between 5 and 10 feet, so the vertical range of the chart roughly
 1495 corresponds to the range of sensitivity over which a measurement has some theo-
 1496 retical usefulness for tracking purposes. Notice that for any given sensor most of the
 1497 time measurements are of negligible value (for tracking); however, at any given time
 1498 there exists a sensor which would produce valuable measurements. Also, the smaller
 1499 the value of *quality*, the better, since it represents the expected average error. The
 1500 vertical green lines are the times when the sector was measuring. Clearly, we want the
 1501 measurements to coincide with the existence of a target (although some measure-
 1502 ments will not, since the radar will occasionally be on sentry duty, searching for
 1503 targets). In Figure 8, sector 1 of node 1 has made multiple measurements when the
 1504 target was first visible starting around time 180, but then did not turn itself off and
 1505 failed to start measuring the target that appeared around time 305. For *quality* of
 1506 measurements we reward an agent for taking high-quality measurements of a target,
 1507 and penalize it for measuring when there is no target.

1508 In our experiments, different sensors had different quality totals based on their
 1509 location relative to the target trajectories (for example, sensor 2 had the best per-
 1510 formance because it could dedicate all its resources to one target only; sensors that
 1511 had to track both targets had worse performance). Our results are presented both per
 1512 sensor and as an average. The results of the quality of measurement are summarized
 1513 in Table 2.

1514 Our real-time negotiation protocol with CBR does much better in creating
 1515 coalitions of the appropriate sectors to track a target, especially for the left target.

Table 1. Percentage of correctly selecting a radar sector for tracking a target for the three agent designs.

Agent designs	Correctly selecting radar sector for Target 1 (left target)	Correctly selecting radar sector for Target 2 (right target)
Negotiating agents with CBR	88.60%	75.69%
Negotiating agents with single strategy	68.20%	74.50%
Reactive agents	69.95%	70.72%




Figure 8. A graph explaining the concept of “quality of measurement” as a function of the visibility of a target and the time a measurement is taken. Target trajectories are represented by the arcs and the sensor measurements are represented by the vertical lines. The x-axis is time in milliseconds. The y-axis is the quality represented by the sensitivity of the target position to measurement errors.

1517

1519

1521

1523

Table 2. Average quality of each sensor in tracking targets for the three agent designs.

Agent designs	Average quality								
	Sensor1	Sensor2	Sensor3	Sensor4	Sensor5	Sensor6	Sensor7	Sensor8	All
Negotiating agents with CBR	190.3 ¹¹	1.9	680.4	462.5	1800.1	6696.8	34.4	468.9	1322.5
Negotiating agents with single strategy	1159.0	2.3	932.6	1787.7	1390.3	4849.3	37.6	2463.9	1577.8
Reactive agents	497.2	2.6	949.0	2040.1	1767.1	4596.3	37.4	2129.3	1502.4

Note that our measure of “quality” is actually the expected average error. Thus, the smaller the value of quality, the better it is.

1524 Overall, our technique scheduled and executed higher quality measurements. In
 1525 two sensors, number 5 and number 6, we did worse. The number 5 sensor’s per-
 1526 formance is comparable to that of the reactive system, while sensor 6 does sub-
 1527 stantially worse than both other techniques. We suspect that these poor
 1528 performances might be due to the initial casebases of the agents controlling these two
 1529 sensors, and how the targets moved. Conceptually, sensor 6 corresponds to sensors
 1530 1, 2, and 8. These sensors with our CBR-powered negotiating agents were able to
 1531 perform much better than with the other agent designs. Similarly, sensor 5 corre-
 1532 sponds to sensor 4, which also had a better performance using the CBR-powered
 1533 negotiating agent design. However, practically, it is possible that, due to how the two
 1534 targets moved along the tracks, sensors 5 and 6 encounter situations not covered in
 1535 the initial casebases and the agents do not have powerful-enough adaptations to
 1536 learn to address those situations. We will investigate this further. In all other cases,
 1537 though, our technique performed from exceptionally better (e.g. sensors 1, 4 and 8)
 1538 to a little better (e.g. sensor 7), and did better on average.

1539 One added measurement was the number of times a radar was turned on and
 1540 measuring. Our technique averaged 63.1 target illuminations per radar per each 250-
 1541 second experiment; the negotiating agents with a single strategy averaged 86.8; and
 1542 the reactive agents averaged 89.1. In other words, our agents achieved higher quality
 1543 measurements taken from the right radar sector at the right time, while *conserving*
 1544 radar resources by not turning them on too often.

1545 In summary, our results showed that in a dynamic, real-time environment of
 1546 multi-sensor target tracking our proposed technique proved to be superior to a
 1547 reactive agent architecture and a BDI architecture of negotiating agents that use a
 1548 single negotiation strategy.

6. Related work

1550 Agent-based negotiations involve information exchanges between two agents where
 1551 the initiating agent desires to persuade the responding agent to accept a task. Two
 1552 agents need to exchange information as each has only a partial set of the information
 1553 or knowledge contained in the entire multi-agent environment, especially where

1554 information cannot be updated and distributed quickly and accurately to all agents.
1555 Thus, when an initiating agent encounters a task that it believes that another agent
1556 can perform, it negotiates by informing that agent of the description of the task and
1557 its own constraints. In this way, each agent maintains a local information base and
1558 information is exchanged only when necessary.

1559 Faratin et al. [15] presented a formal model of negotiation between autonomous
1560 agents. The authors defined *negotiation issues* such as price, volume, duration,
1561 quality, etc. For example, if an agent is trying to sell an item to another agent and
1562 both are negotiating about the price, then the price is the negotiation issue. The
1563 authors defined *negotiation tactics* as the set of functions that determine how to
1564 compute the value of an issue by considering a single criterion, including time,
1565 resources, previous offers and counter offers. There are time-dependent, resource-
1566 dependent, and behavior-dependent tactics such as polynomial, exponential, Boul-
1567 ware and Conceder functions. The authors ultimately defined a *negotiation strategy*
1568 as follows: First, for an issue, there is a weighted counter proposal, made up of a
1569 linear combination of the tactics that generates the value at the current time step of a
1570 negotiation, and the set of tactics is finite. Given a set of tactics, different types of
1571 negotiation behavior can be obtained by weighting the tactics in a different way and
1572 the weights are specified in a matrix. A negotiation strategy is then any function that
1573 is based on the matrix of the weights and the mental state of the agent. This defi-
1574 nition provides a strong formal foundation for generating multi-tactic strategies and
1575 for modifying the negotiation strategy in real-time. However, the negotiation tactics
1576 are functional and thus prescribed. In our approach, we treat negotiation tactics as
1577 situational, temporal and even historical as our agents are able to learn from their
1578 past experiences and adapt old tactics to new situations.

1579 Parsons et al. [13, 39, 40] proposed a comprehensive argumentative negotiation
1580 model that (1) used bridge rules for inference in different units [36] and (2) incor-
1581 porated explicitly multi-context BDI agents [22, 43, 44] and the formalism described
1582 in [25] to construct arguments to evaluate proposals and counterproposals in
1583 negotiation. In particular, the authors made a strong case for sending over rules of
1584 an agent to another agent. This allows the other agent to know the actual reasoning
1585 process behind its counterpart's claims. In an agency of heterogeneous and com-
1586 petitive agents, this makes a lot of sense since it is logical to have agents that reason
1587 differently. Thus, naturally, to convince another agent, an agent lets its counterpart
1588 know about its reasoning steps. This approach promotes diversification and efficient
1589 *knowledge* (in addition to merely information) distribution within an agency. Of
1590 course, in a system of homogeneous agents in which all agents share the exactly same
1591 reasoning behavior, then such an argumentation-based negotiation model of Parsons
1592 and Jennings might not be necessary. The model is truly argumentative since agents
1593 not only argue about what evidence there is, but they also argue about how one
1594 reasons evidentially. Matos et al. [35] also investigated the use of genetic algorithms
1595 to evolve the above negotiation strategies. Strategies with high fitness values were
1596 chosen to survive to the next population generation. The fitness measure was a
1597 combination of gain (the seller's and buyer's scores) and cost (message exchanges).
1598 By the same token, there are different levels of acceptance, based on notions of
1599 argument defeat [28].

1600 Amgoud et al. [4–6] further dealt with preferences between two arguments. First,
1601 an argument is undercut if and only if there exists an argument for the negation of an
1602 element of its support. Second, an argument arg_1 is preferred to another argument
1603 arg_2 according to $Pref$ – a (partial or complete) preordering on the propositions – if
1604 and only if the level of the set of propositions in arg_1 is more referred than the
1605 corresponding set in arg_2 . Using these two definitions, one is able to distinguish
1606 different types of relation between arguments. Suppose arg_1 and arg_2 are two
1607 arguments. Then, arg_2 strongly undercuts arg_1 if and only if arg_2 undercuts arg_1 and
1608 it is not the case that $arg_1 \gg^{Pref} arg_2$, where \gg^{Pref} stands for the strict pre-order
1609 associated with $Pref$. If arg_2 undercuts arg_1 then arg_1 defends itself against arg_2 if
1610 and only if $arg_1 \gg^{Pref} arg_2$. A set of arguments S defends arg_1 if there is some
1611 argument in S which strongly undercuts every argument arg_2 where arg_2 defends arg_1
1612 and arg_1 cannot defend itself against arg_2 . The set S of acceptable arguments of the
1613 argumentation system is the least fix-point of a set of propositions in the system that
1614 is defended by S . These definitions in turn guide the search for the set of arguments
1615 and counter-arguments (or defeaters) and provide the motivation for a dialogue. In
1616 our model, the preference of arguments is prescribed by the CBR system, and the
1617 dynamic part of this preference schemes lies with the prescribed *ordering* of this
1618 arguments. Our CBR system, based on its experience, ranks the order of the argu-
1619 ments to be communicated. Here, the ordering preference is based on the likelihood
1620 of an argument in persuading the other agent in the shortest time. In a way, this
1621 resembles the *degree of strength* of the argument in defending its associated request.

1622 Kraus and Wilkenfeld [26] proposed a negotiation model on hostage crisis. The
1623 model is a modification of Rubinstein’s model [38] of alternative offers which focuses
1624 on the passage of time and the preferences of the players for different agreements as
1625 well as for opting out of the negotiations. The authors assumed that agents cared
1626 only about the nature of the agreement or opting out, and the time at which the
1627 outcome was reached, and not about the sequence of offers and counter-offers that
1628 led to the agreement. In particular, no agent regrets either making an offer that was
1629 rejected or rejecting an offer. They also assumed disagreement is the worst outcome
1630 and the initiator gained over time and the receiver lost over time, since this was
1631 applied to hostage crisis. According to this model, if there is an agreement zone
1632 between the two players in the negotiation, an agreement will be reached in the first
1633 or the second iteration of the negotiation as the players either have complete
1634 knowledge of each other to begin with or have complete knowledge of each other
1635 after the first offer as time constraints are incorporated [24, 27]. The model is par-
1636 ticularly useful when inter-agent communication has to be kept absolutely minimal
1637 (one or two interactions), individual agent knowledge is small, and less intelligent as
1638 a result. However, it does not generally apply to multi-agent negotiations where
1639 uncertain and incomplete knowledge is usually the case among autonomous agents.

1640 Zlotkin and Rosenschein [46, 47, 59, 60] outlined a theoretical negotiation model
1641 for rational agents. Basically, the utility of an agent from a deal is the difference
1642 between the cost of achieving its goal alone and its expected part of the deal. A deal
1643 is *individual rational* if the utility of that deal to each partner of the deal is not
1644 negative, and *pareto-optimal* if there does not exist another deal that dominates it –
1645 there does not exist another deal that is better for one of the agents and not worse of

1646 the other. A *negotiation set* is the set of all the deals that are both *individual rational*
1647 and *pareto-optimal*. In order for the negotiation set to be non-empty, a necessary
1648 condition is that there is no contradiction between the goals of the two agents.
1649 However, this condition is not sufficient since there may still be a conflict between the
1650 agents. A *conflict* is where any joint plan that satisfies the union of goals will cost one
1651 agent (or both) more than it would have spent achieving its own goal in isolation –
1652 i.e., no deal is individual rational. They presented a unified negotiation protocol
1653 (UNP) for conflict resolutions and agents taking partial cooperative steps. The
1654 protocol can be applied to two general agent scenarios. First, agents have been
1655 centrally designed to coexist in a single system and are predisposed toward coop-
1656 erative activity – or in the least, there is some notion of global utility that the system
1657 is trying to maximize. Second, agents are self-serving, have their own utility func-
1658 tions and no global notion of utility. These agents have disparate goals and are
1659 individually motivated. Negotiation can be used to share the work associated with
1660 carrying out a joint plan (for the agents' mutual benefits), or to resolve outright
1661 conflicts arising from limited resources. When a conflict arises, the authors proposed
1662 a conflict resolution via coin toss based on weights. The authors further extended
1663 their model to task-oriented, state-oriented, and worth-oriented functions [61, 62,
1664 63].

1665 Lesser et al. [11, 12, 14] conducted research in which agent-based negotiations
1666 were essentially used for the collection and combination of various information
1667 pieces to derive meaningful and correct observations of the world. The approach
1668 stems from a partial global planning model [14], where an agent has the ability to (1)
1669 represent its own expected interpretation activities (of sensory data), (2) communi-
1670 cate about these expectations with others, (3) model the collective activities of
1671 multiple systems, (4) propose changes to one or more systems' interpretation
1672 activities to improve group performance, and (5) modify its planned local activities
1673 in accordance with the more coordinated proposal. Thus, the object of negotiation
1674 here is to combine information from several agents in order to reach an outcome for
1675 the entire group of agents. During negotiations, the goals, the long-term strategy and
1676 the rating of a plan are exchanged. Decker and Lesser [12] extended the work further
1677 to generalized partial global planning.

1678 Based on the same partial global planning approach, Lander and Lesser [30, 31]
1679 introduced a framework called TAEMS that implements the cooperative (or nego-
1680 tiated) search and conflict resolution among heterogeneous, reusable, expert agents.
1681 The authors identified three negotiated-search strategies: (1) local search, (2) inte-
1682 gration of local search, and (3) general negotiated search. A local search is performed
1683 individually by an agent within its current view of the shared solution space. This
1684 search includes an initiation, a critique – that describes the features of the proposal
1685 the agent agrees with and those it disagrees with [29], and a solution relaxation. In an
1686 integration of local search, solutions are critiqued by other agents, and the critique is
1687 used to determine local conflicts. The agent then performs relaxation (which lowers
1688 the utility of the solution to the agent) and computes the acceptability of the solu-
1689 tion, based on the number of agents considering the solution to be acceptable.
1690 Finally, a general negotiated search is an opportunistic search augmented by the
1691 communication and assimilation of conflict information. Note that there is no

1692 explicit *face-to-face* negotiation involved in the above approach. Agents exchange
1693 information through blackboards, and re-plan based on the critiques and conflicts.
1694 Then they re-post their solutions to the blackboard and the framework controller
1695 sends the information to every agent in the group. The solutions are refined until
1696 acceptable by all agents. However, the negotiated search in addressing inconsisten-
1697 cies such as relaxation is very important in agent-based negotiations, and the task
1698 group and commitment sharing is critical in negotiations as well.

1699 Zeng and Sycara [58] proposed a negotiation model called Bazaar, powered by
1700 Bayesian belief networks. Given each response from the counterpart, an agent updates
1701 its beliefs (through the belief network) regarding the counterpart's reservation price. A
1702 supplier's reservation price, for example, is the price below which the supplier agent will
1703 not accept an offer. Then, the agent makes another offer based on the newly updated
1704 reservation price. The unique characteristic of this work is its use of Bayesian belief
1705 network, which allows modeling of constraints and hypotheses formally. However, one
1706 has to come up with the initial conditional probabilities to build the belief networks.

1707 Kraus et al. [25] described a logical model for multi-agent argumentative negoti-
1708 ations and an implementation called the Automated Negotiation Agent (ANA).
1709 Given the mental states of the agents in their beliefs, desires, intentions, and goals,
1710 the authors defined argumentation as an interactive process emerging from ex-
1711 changes between agents to persuade each other and bring about a change in inten-
1712 tions. There are six different argument types such as threats, promises of a future
1713 reward, appeals to past promises, etc. To decide which arguments to use, an agent
1714 uses negotiation meta-rules. For example, if the counterpart is a memoryless agent,
1715 then do not choose the argument type "appeal to past promise". Next, the agent
1716 ranked all generated arguments in terms of their strengths. For example, a threat is
1717 stronger than a promise of a future reward. Finally, the agent evaluates a request to
1718 decide whether to accept or reject it based on whether there is a contradiction, its
1719 convincing factor and its acceptance value. This logical model provides a compre-
1720 hensive and sound foundation for competitive, selfish (to a certain degree), and
1721 argumentative agents. In our research, however, our agents are cooperative and, due
1722 to time constraints, cannot afford to be *over*-argumentative. In addition, our agents
1723 can neither pose a threat (since agents are cooperative) nor offer a promise of a
1724 future reward (since no agents in our application can guarantee its future state in a
1725 dynamic environment). Moreover, our protocol explicitly defines when certain states
1726 have to be true and for how long. This makes the real-time implementations feasible.

1727 Our use of mental states and temporal logic shared significant similarity to the
1728 SharedPlan model of Grosz and Kraus [18, 19, 20] with the key difference being in
1729 the resolution of the temporal constraints that we refine in our negotiation model. As
1730 described in [20], the SharedPlan formalization provides mental-state specifications
1731 of both SharedPlans and individual plans. Shared plans are constructed by groups of
1732 collaborating agents and include subsidiary SharedPlans formed by subgroups as
1733 well as subsidiary individual plans formed by individual participants in the group
1734 activity. The formalization distinguishes between complete plans – those in which all
1735 the requisite beliefs and intentions have been established – and partial plans. In
1736 addition to the propositional attitude of intending to do an action, it introduces the

1737 attitude of intending that a proposition hold. Stemming from this formalization,
1738 Grosz and Kraus [19] identified two axioms.

1739 First, an agent can only have an intention-toward propositions it believes are
1740 possible. The axiomization does not allow an agent to intend-that an impos-
1741 sible proposition hold. That is, $(\forall G, T_i) Int.Th(G, prop, T_i, T_{prop}, IC_{prop}) \Rightarrow$
1742 $\neg Bel(G, \neg prop, T_i)$ where if G is an agent that intends-that (for a duration of T_i)
1743 proposition $prop$ to be true for duration T_{prop} , within the intention context of IC_{prop} ,
1744 then G does not believe that the proposition is not possible during the duration of the
1745 agent's intention. Similar to our model, this SharedPlans model uses a variant of
1746 intentions, bounded by time to generate other mental states. In SharedPlans, the
1747 intention-that stems from the mutual belief of agents to make sure that the con-
1748 straints for doing a task will hold. Similarly, we treat intentions as derived from
1749 beliefs and desires in our model. In addition, our agents also believe that a propo-
1750 sition is possible if they attempt to achieve that proposition during the negotiation
1751 process. The key difference is that in our model, due to the consideration of real-time
1752 delays, we allow for a mental state (e.g., Bel) – that has initially caused another
1753 mental state (e.g., Int) to hold – to become false before the intention is completely
1754 acted out.

1755 Second, if an agent believes that the intended proposition $prop$ is possible but does
1756 not currently believe it holds and knows some actions it can take to cause propo-
1757 sition $prop$ to hold, then the agent must consider doing at least one of these actions.
1758 So, the second axiom in the SharedPlans model says the following. If an agent is
1759 uncertain about whether proposition $prop$ holds and believes there are some actions
1760 any of which might lead to $prop$ holding, then it must either (a) intend to do, or (b)
1761 be actively considering doing, one of these actions, or (c) it must have considered all
1762 of the actions and determined that it could not do any of them. Our model adopts a
1763 similar practice. However, the mental state of our agent in its intention is subtly
1764 different. In our model, when an agent intends to achieve a $prop$, it carries out a
1765 series of actions towards achieving that $prop$. These actions may lead to a modifi-
1766 cation of the original intention, and may lead to the termination of the original
1767 intention. In the face of conflicts that cannot be resolved, our agents *negotiate* to
1768 change the original intention so that an alternative $prop$ can be achieved.

1769 Gmytrasiewicz and Durfee [17] proposed a Recursive Modeling Method (RMM)
1770 approach to multi-agent coordination under time constraints, and one of their
1771 applications was anti-air defense. Even though agents do not negotiate explicitly,
1772 they make decisions based on their individual modeling of each other, and sub-
1773 sequent modeling of how each other models others.

1774 Our domain of application has certain similarities with the distributed vehicle
1775 monitoring testbed (DVMT) (among many publications on the topic: [33, 37]) which
1776 was used as a tool to investigate distributed problem solving networks. The appli-
1777 cation is vehicle monitoring where each agent receives data only from its own
1778 individual sensor. The sensor coverage areas of the agents overlap. So, the task
1779 becomes one of modeling agent interpretations to ensure local and global consis-
1780 tencies. Hence, in a way, an agent is able to produce an interpretation by itself and
1781 needs to synthesize its own with other agents'. However, in our problem, an agent
1782 needs to cooperate with other agents just to obtain an interpretation of where the

1783 target is. Hence, the DVMT problem focuses more on well-timed sensor fusion while
1784 ours more on coordination.

1785 Distributed sensor networks have also recently received attention where para-
1786 digms and strategies from the multi-agent systems are used to solve the problem of
1787 allocating resources within such networks [34]. In such networks, the sensors may be
1788 heterogeneous in terms of what and when they could sense, the quality of their
1789 measurements, and the coverage area in which they operate. There are approaches
1790 based on negotiation such as dynamic mediation, argumentation, and satisficing
1791 search. There are also approaches based on hierarchical team organization for large-
1792 scale network arbitration and approaches employing distributed constraint satis-
1793 faction techniques dealing with both soft and hard constraints and peer-based
1794 optimization. A version of our work is reported in [34], focusing on a coalition
1795 formation architecture that learns [50].

7. Conclusions

1797 This paper describes a real-time case-based negotiation model and its logical
1798 negotiation protocol. We first defined the characteristics of our agents and the
1799 cooperative environment that they work in. Next, we outlined a set of real-time
1800 constrained design guidelines and introduced our argumentative negotiation pro-
1801 tocol. We described the protocol in a state transition diagram and as a logical
1802 framework. To better represent the axioms used to guide the rules of encounter
1803 between agents, we used the multi-context BDI framework and temporal logic to
1804 define *CanDo*, *Do*, and *Doing* predicates that link agent behaviors to negotiation
1805 motivations. We also defined logically motivations for a negotiation from the
1806 different viewpoints of the initiating and responding agents. Moreover, we utilized
1807 a set of communicative acts to handle a received message and to send out a
1808 message. We further described in detail two sets of axioms that govern the
1809 negotiation behavior between two agents. The axioms are real-time constrained,
1810 facilitated by a suite of functional predicates. The functional predicates are in-
1811 frastructural tools that enable a negotiation process to monitor the pace of its
1812 negotiation, to determine whether it is on time, to determine whether and how to
1813 counter-offer, to analyze the acceptability of a counter-offer, and to evaluate evi-
1814 dence support for a requested task. These predicates are real-time and reflective.
1815 Then, we turned our attention to our case-based negotiation model. We also
1816 outlined a set of guidelines for the design of the model, particularly for enabling
1817 real-time negotiations. In addition, we designed and implemented a real-time
1818 logical framework and a set of real-time functional predicates.

1819 Subsequently, we applied the model to the distributed sensor network domain in
1820 which multiple sensors controlled by different agents are required to collaborate to
1821 track targets. Our experiments showed that our model was able to select the
1822 appropriate sensor sectors and perform good-quality measurements better than
1823 purely reactive agents and negotiating agents with only a single negotiation strategy.

1824 As for our future work, we aim to extend our model to consider multi-agent
1825 systems with heterogeneous agents, different types of resources, and multiple

1826 concurrent negotiation issues. In Sections 3.3.7 and 3.3.8, we have hinted that our
 1827 model assumes a multi-agent with agents of overlapping capabilities, such that when
 1828 an initiating agent encounters a negotiation failure, it is able to approach another
 1829 neighbor for help. In some heterogeneous multi-agent systems, this may not be
 1830 possible. Thus, if an agent is determined to solve a task and if there is only one
 1831 neighbor that has that capability, then the agent has no choice but to continue to
 1832 negotiate with the neighbor. In a way, our proposed model is general enough to
 1833 support the above scenario. However, the specific mechanism is not entirely clear.
 1834 For example, our model will require some extensions, including a *notion of persis-*
 1835 *tence* and a “re-planning” mechanism to overhaul the task description and alloca-
 1836 tion. The persistence notion will still have to trade-off time constraints and task
 1837 accomplishments. We will also look into different types of resources: highly con-
 1838 strained, sharable, consumable, and otherwise. We will investigate the role of these
 1839 negotiations in coalition formation – for example, if a coalition is deemed to fail due
 1840 to a failed negotiation, should the initiating agent still continue with the remaining,
 1841 ongoing negotiations? Finally, we will also study how real-time negotiations based
 1842 on our model perform in noisy and uncertain environments. For example, if there is
 1843 uncertainty in the requested task or there is message loss in the communication, the
 1844 negotiation protocol has to be flexible enough such that the negotiating agents will
 1845 factor uncertainty into their reasoning, and be more cautious in its communication.

1846 8. Acknowledgments

1847 The work described in this paper is sponsored by the Defense Advanced Research
 1848 Projects Agency (DARPA) and the Air Force Research Laboratory, Air Force
 1849 Materiel Command, USAF, under agreement number F30602-99-2-0502. The U.S.
 1850 Government is authorized to reproduce and distribute reprints for Governmental
 1851 purposes notwithstanding any copyright annotation thereon. The views and con-
 1852 clusions contained herein are those of the authors and should not be interpreted as
 1853 necessarily representing the official policies or endorsements, either expressed or
 1854 implied, of the Defense Advanced Research Projects Agency (DARPA), the Air
 1855 Force Research Laboratory, or the U.S. Government. The authors also would like to
 1856 thank the anonymous reviewers for their comments that improved the paper.

Appendix A. Similar Axioms

1861 *Failure 2* When an initiating agent (i) receives a STOP message from a responding
 1862 agent (r), it believes that the responding agent r does not desire to perform the
 1863 requested task ρ and thus stops negotiating with r to perform the task, and the
 negotiation fails.

$$\begin{aligned}
 & C_{in} : stop(r, i, Do(r, \rho), t_{cin, stop}) \wedge I_i(succeed(Negotiate(r, Do(r, \rho))), t'_I) \Rightarrow \\
 & B : B_i(\neg CanDo(r, \rho), t_B) \wedge D : D_i(\neg Do(r, \rho), t_D) \wedge I : I_i(\neg Negotiate(r, Do(r, \rho))), t''_I) \wedge \\
 & I : I_i(\neg succeed(Negotiate(r, Do(r, \rho))), t'''_I) \wedge rejected
 \end{aligned}$$

1865 where $\text{During}(t_{cin,stop}, t'_I) \wedge \text{Meets}(t'_I, t_B) \wedge \text{Meets}(t'_I, t_D) \wedge \text{Meets}(t'_I, t''_I)$. The com-
 1866 municative act *stop* is the encapsulation of receiving and parsing a STOP message.
 1867 This rule is similar to *Failure 1*. In addition, it also states that the agent intends to not
 1868 negotiate. In our current design we do not differentiate between an outright failure
 1869 (failure type 1) and an opt-out failure (failure type 2) – both end with a *rejected* tag.

1870 *Failure 3I* When an initiating agent (i) receives an ABORT message from a
 1871 responding agent (r), it believes that the responding agent r no longer desires to
 1872 perform the requested task ρ and thus stops negotiating with r to perform the task,
 1873 and the negotiation fails.

$$\begin{aligned} & C_{in}: abort(r, i, Do(r, \rho), t_{cin,abort}) \wedge I: I_i(\text{succeed}(\text{Negotiate}(r, Do(r, \rho))), t'_I) \Rightarrow \\ & B: B_i(\neg \text{CanDo}(r, \rho), t_B) \wedge D: D_i(\neg Do(r, \rho), t_D) \wedge I: I_i(\neg \text{Negotiate}(r, Do(r, \rho)), t''_I) \wedge \\ & I: I_i(\neg \text{succeed}(\text{Negotiate}(r, Do(r, \rho))), t'_I) \wedge abort \end{aligned}$$

1875 where $\text{During}(t_{cin,abort}, t'_I) \wedge \text{Meets}(t'_I, t_B) \wedge \text{Meets}(t'_I, t_D) \wedge \text{Meets}(t'_I, t''_I)$. The com-
 1876 municative act *abort* is the encapsulation of receiving and parsing an ABORT
 1877 message. This rule is similar to *Failure 2*.

1878 *Failure 4I* When an initiating agent (i) receives an OUT_OF_TIME message from a
 1879 responding agent (r), it believes that the responding agent r no longer desires to
 1880 perform the requested task ρ and thus stops negotiating with r to perform the task,
 1881 and the negotiation fails.

$$\begin{aligned} & C_{in}: out_of_time(r, i, Do(r, \rho), t_{cin,out_of_time}) \wedge I: I_i(\text{succeed}(\text{Negotiate}(r, Do(r, \rho))), t'_I) \Rightarrow \\ & B: B_i(\neg \text{CanDo}(r, \rho), t_B) \wedge D: D_i(\neg Do(r, \rho), t_D) \wedge I: I_i(\neg \text{Negotiate}(r, Do(r, \rho)), t''_I) \wedge \\ & I: I_i(\neg \text{succeed}(\text{Negotiate}(r, Do(r, \rho))), t'_I) \wedge out_of_time \end{aligned}$$

1883 where $\text{During}(t_{cin,out_of_time}, t'_I) \wedge \text{Meets}(t'_I, t_B) \wedge \text{Meets}(t'_I, t_D) \wedge \text{Meets}(t'_I, t''_I)$. The
 1884 communicative act *out_of_time* is the encapsulation of receiving and parsing an
 1885 ABORT message. This rule is similar to *Failure 2*.

1886 *Failure 3R* When a responding agent (r) receives an ABORT message from an
 1887 initiating agent (i), it believes that the initiating agent i no longer desires to negotiate
 1888 and thus stop negotiating with i to perform the task, and the negotiation fails.

$$\begin{aligned} & C_{in}: abort(i, r, Do(r, \rho), t_{cin,abort}) \wedge I: I_r(\text{succeed}(\text{Negotiate}(i, Do(r, \rho))), t'_I) \Rightarrow \\ & I: I_r(\neg \text{Negotiate}(i, Do(r, \rho)), t''_I) \wedge I: I_r(\neg \text{succeed}(\text{Negotiate}(i, Do(r, \rho))), t'_I) \wedge abort \end{aligned}$$

1890 where $\text{During}(t_{cin,abort}, t'_I) \wedge \text{Meets}(t'_I, t''_I)$. The communicative act *abort* is the
 1891 encapsulation of receiving and parsing an ABORT message. This axiom is similar to
 1892 *Failure 3I*.

1893 *Failure 4R* When a responding agent (r) receives an OUT_OF_TIME message
 1894 from an initiating agent (i), it believes that the initiating agent i no longer desires to
 1895 negotiate and thus stop negotiating with i to perform the task, and the negotiation
 1896 fails.

$C_{in} : out_of_time(i, r, Do(r, \rho), t_{cin, out_of_time}) \wedge I : I_r(succeed(Negotiate(i, Do(r, \rho))), t'_I) \Rightarrow$
 $I : I_i(\neg Negotiate(i, Do(r, \rho)), t''_I) \wedge I : I_i(\neg succeed(Negotiate(i, Do(r, \rho))), t''_I) \wedge out_of_time$

1898 where $During(t_{cin, out_of_time}, t'_I) \wedge Meets(t'_I, t''_I)$. This axiom is similar to *Failure 4I*.

1899 *Abort R* When a responding agent (r) no longer intends to negotiate with an ini-
 1900 tiating agent (i) to perform a requested task ρ , it aborts the negotiation.

$I : I_r(\neg Negotiate(i, Do(r, \rho)), t'_I) \Rightarrow C_{out} : abort(r, i, Do(r, \rho), t_{cout, abort}) \wedge$
 $I : I_r(\neg succeed(Negotiate(i, Do(r, \rho))), t''_I) \wedge abort$

1902 where $During(t_{cout, abort}, t'_I) \wedge Meets(t_{cout, abort}, t''_I) \wedge Finishes(t''_I, t'_I)$. The communica-
 1903 tive act *abort* is the encapsulation of composing and sending an ABORT message. It
 1904 is similar to *Abort I*.

1905 *Out_of_time R* When a responding agent (r) runs out of its allocated time for the
 1906 negotiation with an initiating agent (i) to perform a requested task ρ , it aborts the
 1907 negotiation.

$B : B_r(\neg time(Negotiate(i, Do(r, \rho))), t_B) \wedge I : I_r(succeed(Negotiate(i, Do(r, \rho))), t'_I) \Rightarrow$
 $C_{out} : out_of_time(r, i, Do(r, \rho), t_{cout, out_of_time}) \wedge I : I_r(\neg Negotiate(i, Do(r, \rho)), t'_I) \wedge$
 $I : I_r(\neg succeed(Negotiate(i, Do(r, \rho))), t''_I) \wedge out_of_time$

1909 where

1910 $During(t_B, t'_I) \wedge Finishes(t_{cout, out_of_time}, t_B) \wedge Meets(t_{cout, out_of_time}, t''_I) \wedge Meets(t'_I, t''_I)$.

1911 This axiom is similar to *Out_of_time I*.

1912 *No_response R* When a responding agent (r) detects receives no response from an
 1913 initiating agent (i) during a negotiation, then it unilaterally quits the negotiation with
 1914 a failure.

$B : B_r(no_response(i), t_B) \wedge I : I_r(succeed(Negotiate(i, Do(r, \rho))), t'_I) \Rightarrow$
 $I : I_r(\neg Negotiate(i, Do(r, \rho)), t'_I) \wedge I : I_r(\neg succeed(Negotiate(i, Do(r, \rho))), t''_I) \wedge$
channel_jammed

1916 where $During(t_B, t'_I) \wedge Meets(t'_I, t''_I)$. The *no_response* predicate is one of our real-
 1917 time enabling functional predicates to be discussed in Section 3.4. This axiom allows
 1918 an agent to bail out of a negotiation when the negotiation partner fails to respond.
 1919

Notes

1. A “good-enough, soon-enough” solution is sometimes known as a “satisficing” one [34] – a solution that satisfies the minimum problem requirements and the time constraints.
2. Since our agents must be able to perform multiple, concurrent negotiations, it is important for them to be multithreaded.

3. Take a simplified example, if a tracking coalition must be formed before a target enters a coverage zone that includes the projected path of the target and involves at least three different sensors, and if the negotiation runs longer than the allowable time, then the negotiation is considered to have lost its utility – even if a deal is reached, the coalition is no longer useful for tracking the target. This motivates our design to consider real-time issues significantly.
4. A “time step” can be any interval of time that is reasonable for an application. For target tracking, for example, a reasonable time step is in the order of tens of milliseconds.
5. This could be accomplished through a domain knowledge base that each agent has at creation, or a centralized broker or a mediating agent.
6. The utility-based ranking is detailed in [51, 52] and will not be repeated here.
7. This assumes that the agent knows of neighbors of similar capabilities such that it is able to approach a second neighbor for help when the first approached neighbor refuses to help.
8. See, for example, <http://www.ghg.net/clips/CLIPS.html> for a discussion on CLIPS.
9. Note that the workings of the persuasion threshold and evidence do not define to a joint intention in a formal sense. However, we see these two parameters as the keys to make possible for two negotiating agents to come to a deal with which both are satisfied, leading to an implicit joint intention in the end (if a deal is indeed reached).
10. There have been a variety of persuasion functions used in previous work in negotiation. Lander and Lesser [30, 31] used linear functions of local utility values over contract prices. Zlotkin and Rosenschein [61–63] suggested non-linear and exponential worth (or utility) functions and task-based, pre-defined worth functions. In [15], polynomial, exponential, Boulware tactics and Conceder are used.
11. An explanation of these numbers may be in order: Some sensors do not see a target often. During their idle time they measure the environment to see if a target is there. Such sentry duty is penalized since it results in no quality measurements. So, the actual values in the table fluctuate greatly across sensors. The important aspect is the comparison of the behavior of the same sensor using different techniques.

References

1. J. F. Allen, “Maintaining knowledge about temporal intervals,” *Commun. ACM*, vol. 26, no. 11, pp. 832–843, 1983.
2. J. F. Allen, “Time and time again: the many ways to represent time,” *Int. J. Intell. Syst.*, vol. 6, no. 4, pp. 341–355, 1991.
3. J. F. Allen and G. Ferguson, “Actions and events in interval temporal logic,” *J. Logic Comput. Spl. Issue Actions Proces.*, vol. 4, no. 5, pp. 531–579, 1994.
4. L. Amgoud, S. Parsons, and N. Maudet, “Arguments, dialogue, and negotiation,” in *Proceedings of the 14th European Conference on Artificial Intelligence*, Berlin, Germany, 2000.
5. L. Amgoud, N. Maudet, and S. Parsons, “Modelling dialogues using argumentation,” in *Proceedings of the 4th Conference on Multi-Agent Systems (ICMAS'00)*, Boston, MA, USA, 2000.
6. L. Amgoud and S. Parsons, “Agent dialogues with conflicting preferences,” in *Proceedings of the Workshop on Agent Theories, Architectures and Languages (ATAL'01)*, Seattle, WA, USA, 2001.
7. A. Bolotov and M. Fisher, “A clausal resolution method for CTL branching time temporal logic,” *J. Exp. Theor. Artif. Intell.*, vol. 11, (1999), pp. 77–93, 1999.
8. F. Brazier and J. Treur, “Compositional modelling of reflective agents,” in *Proceedings of the 10th Knowledge Acquisition for Knowledge-Based Systems Workshop (KAW'96)*, Banff, Alberta, Canada, 1996.
9. P. R. Cohen and H. J. Levesque, “Intention is choice with commitment,” *Artif. Intell.*, vol. 42, no. 2–3, pp. 213–261, 1990.
10. P. R. Cohen and H. J. Levesque, “Confirmation and joint action,” in *Proceedings of IJCAI-91*, Sydney, New South Wales, Australia, 1991, pp. 951–959.
11. K. Decker and V. Lesser, “A one-shot dynamic coordination algorithm for distributed sensor networks,” in *Proceedings of AAAI-93*, 1993, pp. 210–216.

12. K. Decker and V. Lesser, "Designing a family of coordination algorithms," in *Proceedings of IC-MAS-95*, San Francisco, CA, 1995, pp. 73–80.
13. J. E. Doran, S. Franklin, S. N. R. Jennings, and T. J. Norman, "On cooperation in multi-agent systems," *Knowl. Eng. Rev.*, vol. 12, no. 3, pp. 309–314, 1997.
14. E. H. Durfee and V. Lesser, "Partial global planning: a coordination framework for distributed hypothesis formation," *IEEE Trans. Syst. Man Cybernet.*, vol. 21, no. 5, pp. 1167–1183, 1991.
15. P. Faratin, C. Sierra, and J. R. Jennings, "Negotiation decision functions for autonomous agents," *Int. J. Robot. Autonomous Syst.*, vol. 24, no. 3–4, pp. 159–182, 1998.
16. J. R. Galliers, "A strategic framework for multi-agent cooperative dialogue," in *Proceedings of ECAI'88*, Munich, Germany, 1998, pp. 415–420.
17. P. J. Gmytrasiewicz and E. H. Durfee, "Rational coordination in multi-agent environments," *Autonomous Agents and Multiagent Systems*, vol. 3, no. 4, pp. 319–350, 2000.
18. B. Grosz and S. Kraus, "Collaborative plans for group activities," in *Proceedings of the 1993 International Joint Conference on Artificial Intelligence (IJCAI-93)*, San Mateo, CA, 1993, pp. 367–373.
19. B. Grosz and S. Kraus, "Collaborative plans for complex group action," *Artif. Intell.*, vol. 86, no. 2, pp. 269–357, 1996.
20. B. J. Grosz and S. Kraus, "The evolution of SharedPlans," in Rao, A. and M. Wooldridge (eds.), *Foundations and Theories of Rational Agency*, Kluwer Academic Publishing, 1998.
21. H. Jakobovits and D. Vermeir, "Dialectic semantics for argumentation frameworks," in *Proceedings of the Seventh International Conference on Artificial Intelligence and Law*, 1999, pp. 53–62.
22. N. R. Jennings, "Controlling cooperative problem solving in industrial multi-agent systems using joint intentions," *Artif. Intell.*, vol. 75, (1993), pp. 195–240, 1993.
23. J. Kolodner *Case-Based Reasoning*. Morgan Kaufmann, San Mateo, CA, 1993.
24. S. Kraus, "Beliefs, time, and incomplete information in multiple encounter negotiations among autonomous agents," *Ann. Math. Artif. Intell.*, vol. 20, no. 1–4, pp. 111–159, 1997.
25. S. Kraus, K. Sycara, and A. Evenchik, "Reaching agreements through argumentation: a logical model and implementation," *Artif. Intell. J.*, vol. 104, no. 1–2, pp. 1–69, 1998.
26. S. Kraus and J. Wilkenfeld, "A strategic negotiations model with applications to an international crisis," *IEEE Trans. Syst. Man Cybernet.*, vol. 23, no. 1, pp. 313–323, 1993.
27. S. Kraus, J. Wilkenfeld, and G. Zlotkin, "Multiagent negotiation under time constraints," *Artif. Intell.*, vol. 75, (1995), pp. 297–345, 1995.
28. P. Krause, S. Ambler, M. Elvang-Gøransson, and J. Fox, "A logic of argumentation for reasoning under uncertainty," *Comput. Intell.*, vol. 11, (1995), pp. 113–131, 1995.
29. B. Låasri, H. Låasri, S. Lander, and V. Lesser, "A generic model for intelligent negotiating agents," *Int. J. Intell. Coop. Inform. Syst.*, vol. 1, no. 2, pp. 291–317, 1992.
30. S. Lander and V. Lesser, "Customizing distributed search among agents with heterogeneous knowledge," in *Proceedings of CIKM-92*, Baltimore, MD, 1992, pp. 335–344.
31. S. Lander and V. Lesser, "Understanding the role of negotiation in distributed search among heterogeneous agents," in *Proceedings of IJCAI-93*, Chambéry, France, 1992, pp. 438–444.
32. J. Lawton "The Radsim simulator". V. Lesser, C. L. Ortiz Jr., M. Tambe (eds.), *Distributed Sensor Networks: A Multiagent Perspective*. Kluwer Academic Publishers, The Netherlands, 2003, pp. 11–20.
33. V. Lesser and D. Corkill, "The distributed vehicle monitoring testbed: a tool for investigating distributed problem solving networks," *AI Magazine*, vol. 4, no. 3, pp. 15–33, 1983.
34. V. Lesser, C. L. Ortiz Jr., M. Tambe (eds.) *Distributed Sensor Networks: A Multiagent Perspective*. Kluwer, Norwell MA.
35. N. Matos, C. Sierra, and J. R. Jennings, "Negotiation strategies: an evolutionary approach," in *Proceedings of International Conference on Multiagent Systems (ICMAS)*, Paris, France, 1998, pp. 182–189.
36. P. Noriega and C. Sierra, "Towards layered dialogical agents," in *Proceedings of ECAI Workshop on ATAL'96*, Budapest, Hungary, 1996, pp. 157–171.
37. C. Norman, V. Lesser, and Q. Long, "Distributed Sensor Interpretation: Modeling Agent Interpretations in DRESUN," in *University of Massachusetts Technical Report*, UMCS 93-75, 1993.
38. M. J. Osborne and A. Rubinstein *A Course in Game Theory*. MIT Press, Cambridge, MA, 1994.

39. S. Parsons and N. R. Jennings, "Negotiation through argumentation—a preliminary report," in *Proceedings of ICMAS-96*, Kyoto, Japan, 1996, pp. 267–274.
40. S. Parsons, C. Sierra, and N. R. Jennings, "Agents that reason and negotiate by arguing," *J. Logic Comput.*, vol. 8, no. 3, pp. 261–292, 1998.
41. H. V. D. Parunak, A. D. Baker, and S. J. Clark, "The AARIA agent architecture: an example of requirements-driven agent-based system design," in *Proceedings of the 1st International Conference on Autonomous Agents (ICMAS'97)*, 1997, pp. 482–483.
42. H. Prakken, "From logic to dialectics in legal argument," in *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, 1995, pp. 165–174.
43. A. Rao and M. Georgeff, "Modeling rational agents within a BDI-architecture," in *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, 1991, pp. 473–484.
44. A. Rao and M. Georgeff, "BDI agents: from theory to practice," in *Proceedings of ICMAS-95*, San Francisco, CA, 1995, pp. 312–319.
45. A. Rao and M. Georgeff, "Decision procedures for BDI logics," *J. Logic Comput.*, vol. 8, no. 3, pp. 293–342, 1998.
46. J. S. Rosenschein and G. Zlotkin *Rules of Encounter*. MIT Press, Cambridge, MA, 1994.
47. J. S. Rosenschein and G. Zlotkin, "Designing conventions for automated negotiation.," *AI Magazine*, vol. 15, no. 3, pp. 29–46, 1994.
48. T. W. Sandholm and V. R. Lesser, "Coalition formation among bounded rational agents," in *Proceedings of IJCAI-95*, Montreal, Canada, 1995, pp. 662–669.
49. M. Singh, "Developing formal specifications to coordinate heterogeneous autonomous agents," in *Proceedings of the 3rd International Conference on Multiagent Systems (ICMAS'98)*, 1998, pp. 261–268.
50. L.-K. Soh, H. Sevay, and C. Tsatsoulis, "A satisficing, learning, and negotiated coalition formation architecture," in V. Lesser, M. Tambe and C. Ortiz (eds.), *Distributed Sensor Networks: A Multiagent Perspective*, Kluwer Publishing, Chapter 7, 2003, pp.109–138.
51. L.-K. Soh and C. Tsatsoulis, "Reflective negotiating agents for real-time multisensor target tracking," in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'01)*, August 6–11, Seattle, WA, 2001, pp. 1121–1127.
52. L.-K. Soh and C. Tsatsoulis, "Agent-based argumentative negotiations with case-based reasoning," in *Working Notes of the AAAI Fall Symposium Series on Negotiation Methods for Autonomous Cooperative Systems*, November 1–4, North Falmouth, MA, 2001, pp. 16–25.
53. B. Srinivasan, S. Pather, R. Hill, F. Ansari, and D. Niehaus, "A firm real-time system implementation using commercial off-the-shelf hardware and free software," in *Proceedings of RTAS-98*, June Denver, CO, 1998, pp. 112–119.
54. K. Sycara, M. Paolucci, M. van Velsen, and J. Giampapa, "The RETSINA MAS infrastructure," Technical Report CMU-RI-TR-01-05, Robotics Institute Technical Report, Carnegie Mellon, 2001.
55. M. Tambe, "Towards flexible teamwork," *J. Artif. Intell. Res.*, vol. 7, (1997), pp. 83–124, 1997.
56. S. A. Vere, "Planning in time: windows and durations for activities and goals," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 5, no. 3, pp. 246–267, 1983.
57. M. Wooldridge and N. Jennings, "Intelligent agents: theory and practice," *Knowl. Eng. Rev.*, vol. 10, no. 2, pp. 114–152, 1995.
58. D. Zeng and K. Sycara, "Bayesian learning in negotiation," *Int. J. Human-Comput. Stud.*, vol. 48, (1998), pp. 125–141, 1998.
59. G. Zlotkin and J. S. Rosenschein, "Negotiation and task sharing among autonomous agents in cooperative domains," in *Proceedings of IJCAI-89*, Detroit, MI, August, 1989, pp. 912–917.
60. G. Zlotkin and J. S. Rosenschein, "Cooperation and conflict resolution via negotiation among autonomous agents in noncooperative domains," *IEEE Trans. Syst. Man Cybernet. Spl. Issue Distributed Artif. Intell.*, vol. 21, no. 6, pp. 1317–1324, 1991.
61. G. Zlotkin and J. S. Rosenschein, "Mechanism design for automated negotiation, and its application to task oriented domains," *Artif. Intell.*, vol. 86, no. 2, pp. 195–244, 1996.
62. G. Zlotkin and J. S. Rosenschein, "Mechanisms for automated negotiation in state oriented domains," *J. Artif. Intell. Res.*, vol. 5, (1996), pp. 163–238, 1996.
63. G. Zlotkin and J. S. Rosenschein, "Compromise in negotiation: exploiting worth functions over states," *Artif. Intell.*, vol. 84, no. 1–2, pp. 151–176, 1996.