

FOUNDATION FOR INTELLIGENT PHYSICAL AGENTS

FIPA Interaction Protocol Library Specification

Document title	FIPA Interaction Protocol Library Specification		
Document number	XC00025E	Document source	FIPA TC C
Document status	Experimental	Date of this status	2001/08/10
Supersedes	FIPA00003		
Contact	fab@fipa.org		
Change history			
2001/01/29	Approved for Experimental		
2001/08/10	Line numbering added		

© 2000 Foundation for Intelligent Physical Agents - <http://www.fipa.org/>

Geneva, Switzerland

Notice

Use of the technologies described in this specification may infringe patents, copyrights or other intellectual property rights of FIPA Members and non-members. Nothing in this specification should be construed as granting permission to use any of the technologies described. Anyone planning to make use of technology covered by the intellectual property rights of others should first obtain permission from the holder(s) of the rights. FIPA strongly encourages anyone implementing any part of this specification to determine first whether part(s) sought to be implemented are covered by the intellectual property of others, and, if so, to obtain appropriate licenses or other permission from the holder(s) of such intellectual property prior to implementation. This specification is subject to change without notice. Neither FIPA nor any of its Members accept any responsibility whatsoever for damages or liability, direct or consequential, which may result from the use of this specification.

19 **Foreword**

20 The Foundation for Intelligent Physical Agents (FIPA) is an international organization that is dedicated to promoting the
21 industry of intelligent agents by openly developing specifications supporting interoperability among agents and agent-
22 based applications. This occurs through open collaboration among its member organizations, which are companies and
23 universities that are active in the field of agents. FIPA makes the results of its activities available to all interested parties
24 and intends to contribute its results to the appropriate formal standards bodies.

25 The members of FIPA are individually and collectively committed to open competition in the development of agent-
26 based applications, services and equipment. Membership in FIPA is open to any corporation and individual firm,
27 partnership, governmental body or international organization without restriction. In particular, members are not bound to
28 implement or use specific agent-based standards, recommendations and FIPA specifications by virtue of their
29 participation in FIPA.

30 The FIPA specifications are developed through direct involvement of the FIPA membership. The status of a
31 specification can be either Preliminary, Experimental, Standard, Deprecated or Obsolete. More detail about the process
32 of specification may be found in the FIPA Procedures for Technical Work. A complete overview of the FIPA
33 specifications and their current status may be found in the FIPA List of Specifications. A list of terms and abbreviations
34 used in the FIPA specifications may be found in the FIPA Glossary.

35 FIPA is a non-profit association registered in Geneva, Switzerland. As of January 2000, the 56 members of FIPA
36 represented 17 countries worldwide. Further information about FIPA as an organization, membership information, FIPA
37 specifications and upcoming meetings may be found at <http://www.fipa.org/>.

38 **Contents**

39 1 Scope 1
40 2 Overview 2
41 2.1 Interaction Protocols 2
42 2.2 Status of a FIPA-Compliant Interaction Protocol 2
43 2.3 FIPA Interaction Protocol Library Maintenance 3
44 2.4 Inclusion Criteria 3
45 3 AUML Sequence Diagrams for Interaction Protocol Specification 4
46 3.1 Introduction 4
47 3.2 Extending UML by Protocol Diagrams 5
48 3.2.1 Protocol Diagrams 5
49 3.2.2 AgentRoles 7
50 3.2.3 Agent Lifeline 8
51 3.2.4 Threads of Interaction 10
52 3.2.5 Messages 11
53 3.2.6 Complex Messages 13
54 3.2.7 Nested Protocols 14
55 3.2.8 Complex Nested Protocols 15
56 3.2.9 Threads of Interaction and Messages Inside and Outside Nested Protocols 16
57 3.2.10 Parameterised Protocols 17
58 3.2.11 Bound Elements 18
59 4 References 21
60

60 **1 Scope**

61 This document contains:

62

63 Specifications for structuring the FIPA Interaction Protocol Library (IPL) including a status of a FIPA Interaction
64 Protocols (IPs), maintenance of the library and inclusion criteria for new IPs.

65

66 A description of how to understand and express IPs using AUML (Agent Unified Modeling Language).

67

68 The FIPA IP registry list.

69

70 This specification is primarily concerned with defining the structure of the FIPA IPL and the requirements for an IP to be
71 included in the library.

72

72 2 Overview

73 This specification focuses on the organization, structure and status of the FIPA IPL and discusses the main
74 requirements that an IP must satisfy in order to be FIPA-compliant. The objectives of standardising and defining a
75 library of FIPA compliant IPs are:

- 76
- 77 To provide tested patterns of agent interaction that may be of use in various aspects of agent-based systems,
- 78
- 79 To facilitate the reuse of standard agent IPs, and,
- 80
- 81 To express IPs in a standard agent unified modelling language (AUML).
- 82

83 In the following, we present the basic principles of the FIPA IPL which help to guarantee that the IPL is stable, that there
84 are public rules for the inclusion and maintenance of the IPL, and that developers seeking a public IPs can use the IPL.
85

86 2.1 Interaction Protocols

87 Ongoing conversations between agents often fall into typical patterns. In such cases, certain message sequences are
88 expected, and, at any point in the conversation, other messages are expected to follow. These typical patterns of
89 message exchange are called *interaction protocols*. A designer of agent systems has the choice to make the agents
90 sufficiently aware of the meanings of the messages and the goals, beliefs and other mental attitudes the agent
91 possesses, and that the agent's planning process causes such IPs to arise spontaneously from the agents' choices.
92 This, however, places a heavy burden of capability and complexity on the agent implementation, though it is not an
93 uncommon choice in the agent community at large. An alternative, and very pragmatic, view is to pre-specify the IPs, so
94 that a simpler agent implementation can nevertheless engage in meaningful conversation with other agents, simply by
95 carefully following the known IP.

96

97 This section of the specification details a number of such IPs, in order to facilitate the effective inter-operation of simple
98 and complex agents. No claim is made that this is an exhaustive list of useful IPs, nor that they are necessary for any
99 given application. The IPs are given pre-defined names and the requirement for adhering to the specification is:

100

101 *A FIPA ACL-compliant agent need not implement any of the standard IPs, nor is it restricted from using other IP names.*
102 *However, if one of the standard IP names is used, the agent must behave consistently with the IP specification given*
103 *here.*

104

105 *These IPs are not intended to cover every desirable interaction type. Individual IPs do not address a number of*
106 *common real-world issues in agent interaction, such as exception handling, messages arriving out of sequence,*
107 *dropped messages, timeouts, cancellation, etc. Rather, the IPs defined in this specification set should be viewed as*
108 *interaction patterns, to be elaborated according to the context of the individual application. This strategy means that*
109 *adhering to the stated IPs does not necessarily ensure interoperability; further agreement between agents about the*
110 *issues above is required to ensure interoperability in all cases.*

111

112 Note that, by their nature, agents can engage in multiple dialogues, perhaps with different agents, simultaneously. The
113 term *conversation* is used to denote any particular instance of such a dialogue. Thus, the agent may be concurrently
114 engaged in multiple conversations, with different agents, within different IPs. The remarks in this section, which refer to
115 the receipt of messages under the control of a given IP, refer only to a particular conversation.
116

117 2.2 Status of a FIPA-Compliant Interaction Protocol

118 The definition of an IP belonging to the FIPA IPL is normative, that is, if a given agent advertises that it employs an IP in
119 the FIPA Content Language Library (see [FIPA00007]), then it must implement the IP as it is defined in the FIPA IPL.
120 However, FIPA-compliant agents are not required to implement any of the FIPA IPL IPs themselves, except those
121 required for Agent Management (see [FIPA00023]).
122

123 By collecting IP definitions in a single, publicly accessible registry, the FIPA IPL facilitates the use of standardized IPs
124 by agents developed in different contexts. It also provides a greater incentive to developers to make their IPs generally
125 applicable.

126
127 FIPA is responsible for maintaining a consistent list of IP names and for making them publicly available. In addition to
128 the list of encoding schemes, each IP in the FIPA IPL may specify additional information, such as stability information,
129 versioning, contact information, different support levels, etc.
130

131 **2.3 FIPA Interaction Protocol Library Maintenance**

132 The most effective way of maintaining the FIPA IPL is through the use of the IPs themselves by different agent
133 developers. This is the most direct way of discovering possible bugs, errors, inconsistencies, weaknesses, possible
134 improvements, as well as capabilities, strengths, efficiency, etc.
135

136 In order to collect feedback on the IPs in the library and to promote further research, FIPA encourages coordination
137 among designers, agent developers and FIPA members.
138

139 **2.4 Inclusion Criteria**

140 To populate the FIPA IPL, setting fundamental guidelines for the selection of specific IPs is necessary. The minimal
141 criteria that must be satisfied for an IP to be FIPA compliant are:

142 A clear and accurate representation of the IP is provided using AUML protocol diagram,

143
144 Substantial and clear documentation must be provided, and,

145
146 The usefulness of a new IP should be made clear.
147

148
149 FIPA does not enforce the use of any particular IP.
150
151

151 3 AUML Sequence Diagrams for Interaction Protocol Specification

152 3.1 Introduction

153 During the 1970s, structured programming was the dominant approach to software development. Along with it, software
 154 engineering technologies were developed in order to ease and formalize the system development lifecycle: from
 155 planning, through analysis and design and finally to system construction, transition, and maintenance. In the 1980s,
 156 object-oriented languages experienced a rise in popularity, bringing with it new concepts such as data encapsulation,
 157 inheritance, messaging, and polymorphism. By the end of the 1980s and beginning of the 1990s, a jungle of modelling
 158 approaches grew to support the object-oriented marketplace. To make sense of and unify these various approaches, an
 159 Analysis and Design Task Force was established on 29 June 1995 within the Object Management Group (OMG). And
 160 by November 1997, a de jure standard was adopted by the OMG members called the Unified Modelling Language
 161 (UML - see [OMGuml]).

162
 163 UML unifies and formalizes the methods of many object-oriented approaches, including analysis and design [Booch94
 164 and Booch95], modelling [Rumbaugh91] and software engineering [Jacobson94]. It supports the following kinds of
 165 models:

166 **Static models**

167 Such as class and package diagrams describe the static semantics of data and messages. Within system
 168 development, class diagrams are used in two different ways, for two different purposes. First, they can model a
 169 problem domain conceptually and since they are conceptual in nature, they can be presented to the customers.
 170 Second, class diagrams can model the implementation of classes which guides developers. At a general level, the
 171 term *class* refers to the encapsulated unit and at the conceptual level, models types and their associations; the
 172 implementation level models implementation classes. While both can be more generally thought of as classes, their
 173 usage as concepts and implementation notions is important both in purpose and semantics. Package diagrams
 174 group classes in conceptual packages for presentation and consideration. (Physical aggregations of classes are
 175 called *components* that are in the implementation model family, mentioned below.)

177 **Dynamic models**

178 These include interaction diagrams (that is, sequence and collaboration diagrams), state charts and activity
 179 diagrams.

181 **Use cases**

182 The specification of actions that a system or class can perform by interacting with outside actors. They are
 183 commonly used to describe how a customer communicates with a software product.

185 **Implementation models**

186 These describe the component distribution on different platforms, such as component models and deployment
 187 diagrams

189 **Object Constraint Language (OCL)**

190 This is a simple formal language to express more semantics within an UML specification. It can be used to define
 191 constraints on the model, invariant, pre- and post-conditions of operations and navigation paths within an object net.

192
 193
 194 For modelling agents and agent-based systems, UML is insufficient. Compared to objects, agents are active because
 195 they act for reasons that emerge from themselves. The activity of agents is based on their internal states, which include
 196 goals and conditions that guide the execution of defined tasks. While objects need control from outside to execute their
 197 methods, agents know the conditions and intended effects of their actions and hence take responsibility for their needs.
 198 Furthermore, agents do not only act solely but together with other agents. Multi-agent systems can often resemble a
 199 social community of interdependent members that act individually.

200
 201 However, no sufficient specification formalism exists yet for agent-based system development. To employ agent-based
 202 programming, a specification technique must support the whole software engineering process—from planning, through
 203 analysis and design, and finally to system construction, transition, and maintenance.

204 A proposal for a full life-cycle specification of agent-based system development is beyond the scope of this
205 specification. Here, we suggest a subset of an agent-based extension to the standard UML, called AUML, for the
206 specification of agent interaction protocols (AIPs).

207
208 It has to be distinguished between generic (or parameterised) protocols (and their instantiations) and domain-specific
209 protocols.
210

211 **3.2 Extending UML by Protocol Diagrams**

212 In the following, we provide sequence diagrams for AUML [Odell2000], an extension to UML. We refer to these
213 sequence diagrams as *protocol diagrams* (PDs) which show well-defined interactions among agents. Note that we do
214 not define formal semantics for the communicative acts for AUML, but instead use the UML meta-model.
215

216 **3.2.1 Protocol Diagrams**

217 Adapted from [OMGuml], section 3.59.
218

219 3.2.1.1 Semantics

220 A PD represents an interaction, which is a set of messages exchanged among different agent roles within a
221 collaboration to effect a desired behaviour of other AgentRoles or agent instances.
222

223 3.2.1.2 Notation

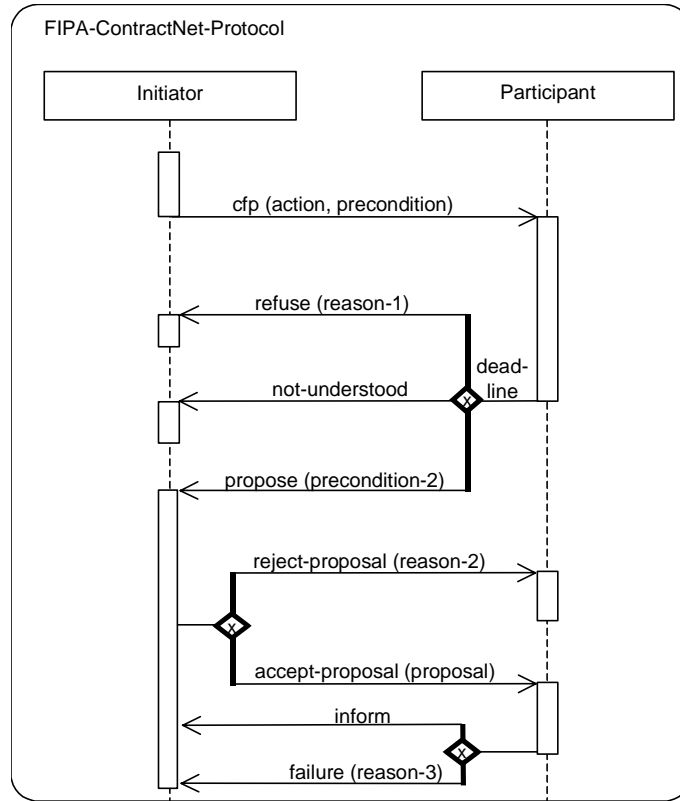
224 A PD has two dimensions: the vertical dimension represents time, the horizontal dimension represents different
225 AgentRoles. Normally the time proceeds down the page and usually only time sequences are important, but in real-time
226 applications the time axis could be an actual metric. There is no significance to the horizontal ordering of the
227 AgentRoles.
228

229 3.2.1.3 Presentation Options

230 The axes can be interchanged, so that time proceeds horizontally to the right and different AgentRoles are shown as
231 horizontal lines.
232

233 Various labels (such as timing marks, generated goals depending on the received message, etc.) can be shown either
234 in the margin or near the lifelines or messages that they label.
235
236

236 3.2.1.4 Example
237



238
239

240 3.2.1.5 Mapping

241 The mapping is analogous defined as for sequence diagrams (see [OMGuml]).

242

243 A PD maps like a sequence diagram into an Interaction and an underlying Collaboration. An Interaction specifies a
 244 sequence of communications; it contains a collection of partially ordered Messages, each specifying a communication
 245 between a sender role and a receiver role. Collections of agent roles that conform to the ClassifierRoles in the
 246 Collaboration owning the Interaction, communicate by dispatching Stimuli that conform to the Messages in the
 247 Interaction. An AgentRole maps into a ClassifierRole. A PD presents one collection of AgentRoles and arrows mapping
 248 to AgentRole and Stimuli that conform to the ClassifierRoles and Messages in the Interaction and its Collaboration.
 249

250

251 In a PD, each AgentRole box with its lifeline maps into an agent role that conforms to a ClassifierRole in the
 252 Collaboration. The name fields maps into the name of the agent, the role name into the Classifier's name and the class
 253 field maps into the names of the Classifier (in this case AgentClasses being Classes) being the base Classifiers of the
 254 ClassifierRole. The splitting of lifelines has a concurrency Association defining either AND/OR parallelism or decision
 255 Association denoting threads (<<thread>>). The associations among roles are not shown on the sequence diagram
 256 since they must be obtained in the model from a complementary collaboration diagram or other means. A message
 257 arrow maps into a Stimulus connected to two AgentRoles. the sender and receiver AgentRole. The Stimulus conforms
 258 to a Message between the ClassifierRoles corresponding to the two AgentRoles' lifelines that the arrow connects. The
 259 Link is used for the communication of the Stimulus and plays the role specified by the AssociationRole connected to the
 260 Message. Unless the correct Link can be determined from a complementary collaboration diagram or other means, the
 261 Stimulus is either not attached to a Link (not a complete model), or it is attached to an arbitrary Link or to a dummy Link
 262 between the Instances conforming to the AssociationRole implied by the two ClassifierRoles due to the lack of complete
 263 information. The name of the communicative act is mapped onto the behaviour associated by the action performing,
 264 requested information, information passing, negotiation or error handling connected to the Message. Different
 265 alternatives exist for showing the arguments of the Stimulus. If references to the actual Instances being passed as
 266 arguments are shown, these are mapped onto the arguments of the Stimulus. If the argument expressions are shown
 instead, these are mapped onto the Arguments of the action performing, requested information, information passing,

267 negotiation or error handling connected to the dispatching communicative act. Finally, if the types of the arguments are
 268 shown together with the name of the communicative act, these are mapped onto the parameter types of the
 269 communicative act. A timing label placed on the level of an arrow endpoint maps into the name of the corresponding
 270 Message. A constraint or guard placed on the diagrams maps into a Constraint on the entire Interaction. The cardinality
 271 label restricts the number of sending and receiving instances of agent roles accordingly to the numbers denoted at the
 272 beginning (sender) and end (receiver) of the message.

273

274 An arrow with the arrowhead pointing to an AgentRole symbol within the frame of the diagram maps into a Stimulus
 275 dispatched by a `CreateAction`, that is, the Stimulus conforms to a Message in the Interaction which is connected to
 276 the `CreateAction`. The interpretation is that the AgentRole instance (not an arbitrary agent role, nor a set of
 277 AgentRole instances) is created by dispatching the Stimulus, and the AgentRole instance conforms to the receiver role
 278 specified in the Message. After the creation of the AgentRole instance, it may immediately start interacting with other
 279 AgentRoles. This implies that the creation of the AgentRole dispatches these Stimuli. If an AgentRole instance
 280 termination symbol ("X") is the target of the of an arrow, the arrow maps into a Stimulus which will cause the receiving
 281 agent role instance to be removed. The Stimulus conforms to a Message in the Interaction with a `DestroyAction`
 282 attached to the Message or the agent instance terminates itself.

283

284 The order of the arrows in the diagram map onto a pair of associations between the Messages that correspond to the
 285 Stimuli the arrows maps onto. A predecessor association is established between Messages corresponding to
 286 successive arrow ends in the vertical sequence. In case of concurrent arrows preceding an arrow, the corresponding
 287 Message has a collection of predecessors. In case of exclusive-or and inclusive-or arrows preceding an arrow the
 288 corresponding message has one and at least one out of the collection of possible predecessors, respectively.
 289 Moreover, each Message has an activator (thread of interaction) association to the Message corresponding to the
 290 incoming arrow of the activation or pro-active sending of a message.

291

292 A nested protocol maps into a PD. The name compartment of a nested protocol maps into the name of the
 293 Collaboration. The guard and constraint compartment maps into a constraint on the complete Interaction.

294

295 A complex nested protocol maps into a PD. The order of the messages within the protocol is defined according to the
 296 combination of the complex nested protocol. The ordering of the messages in the nested protocol is the ordering of
 297 these protocols. Depending on the combination the messages are sent in AND/OR parallelism or decision ordering.

298

299 3.2.2 AgentRoles

300 In the framework of agent oriented programming an agent satisfying a distinguished role behaves in a special way. In
 301 contrast to this semantics *role* in UML is an instance focused term. Moreover the term *multi-object* does not fit to
 302 describe AgentRoles but it is used to show operations that address the entire set, rather than a single object in it.
 303 However, there is a communication with one instance of this multi-object. By *AgentRole* a set of agents satisfying
 304 distinguished properties, interfaces or having a distinguished behaviour are meant.

305

306 UML distinguishes between:

307

308 multiple classifications where a retailer agent can act as well as a buyer as well as a seller agent, for example, and,

309

310 dynamic classification where an agent can change its classification during its existence.

311

312 Agents can perform various roles within one IP. Using a contract-net protocol, for example, between a buyer and a
 313 seller of a product, the initiator of the protocol has the role of a buyer and the participant has the role of a seller. But the
 314 seller can as well be a retailer agent, which acts as a seller in one case and as a buyer in another case, i.e. agents
 315 satisfying a distinguished role can support multiple classification and dynamic classification. Another example can be
 316 found in [FIPA00023] which defines the functionality of the Directory Facilitator (DF) and the Agent Management
 317 System (AMS). These functionalities can be implemented by different agents, but the functionality of the DF and AMS
 318 can also be integrated into one agent.

319

320 An AgentRole can be seen as a set of agents satisfying a distinguished interface, service description or behaviour.
 321 Therefore the implementation of an agent can satisfy different roles.

322
323
324
325

Note that within FIPA the notion of role is not used, but in the framework of specifying agent-based systems this notion is appropriate.

326
327
328
329
330

3.2.2.1 Semantics

An AgentRole describes two different variations that can apply within a protocol definition. A protocol can be defined between different concrete agent instances or a set of agents satisfying a distinguished role and/or class. An agent satisfying a distinguished AgentRole and class is called agent of a given AgentRole and class, respectively.

331
332
333

3.2.2.2 Notation

An AgentRole is shown as a rectangle that is filled with some string of one of the following forms:

334
335
336

role

This denotes arbitrary agents satisfying the distinguished AgentRole.

337
338
339

instance / role-1 ... role-*n*

This denotes a distinguished agent instance that satisfies the AgentRoles 1-*n* where $n > 0$.

340
341
342

instance / role-1 ... role-*n* : class

This denotes a distinguished agent instance that satisfied the AgentRoles 1-*n* where $n > 0$ and class it belongs to.


343
344
345
346

3.2.2.3 Presentation Options

The second case can be abbreviated as instance if *n* equals zero, that is, a concrete agent is meant independent of the role(s) it satisfies and class it belongs to.

347
348

3.2.2.4 Example



349
350

351
352
353

3.2.2.5 Mapping

See *Section 3.2.1.5, Mapping*.

354
355
356
357
358

3.2.3 Agent Lifeline

The agent lifeline defines the time period when an agent exists. For example a user agent is created when a user logs on to the system and the user agent is destroyed when the user logs off. Another example is when an agent migrates from one machine to another.

359
360
361
362
363

3.2.3.1 Semantics

A PD defines the pattern of communication, that is, the steps in which the communicative acts are sent between agents of different AgentRoles. The agent lifeline describes the time period in which an agent of a given AgentRole exists. Only during this time period an agent can participate on a protocol.

364
365
366

The lifeline starts when the agent of a given AgentRole is created and ends when it is destroyed. The lifeline can be split in order to describe AND/OR parallelism and decisions and may merge together at some subsequent point.

367
368
369
370
371

3.2.3.2 Notation

An agent lifeline is shown as a vertical dashed line. The lifeline represents the existence of an agent of a given AgentRole at a particular time. If the agent is created or destroyed during the period of time shown on the PD, then its lifeline starts or stops at the appropriate point; otherwise it goes from the top of the diagram to the bottom. An AgentRole is drawn at the head of the lifeline. If an agent of a given AgentRole is created during the PD, then the

372 message that creates it is drawn with its arrowhead on the AgentRole. Note, that the AgentRole (see Section 3.2.3.4,
 373 *Example*) that receives the message is responsible for the creation of the agent instance, that is, the arrowhead ends at
 374 the dashed line of the AgentRole receiving the message and the AgentRole is fixed at the left-hand or right-hand side of
 375 the lifeline or the thread of interaction. If an agent instance is destroyed during the PD, then its destruction is marked by
 376 a large "X", either at the message that causes the destruction or (in the case of self destruction) at the final action of the
 377 AgentRole. The termination is restricted to concrete instances of an agent role.

378
 379 AgentRoles that exist when a protocol starts is shown at the top of the diagram (above the first message arrow). An
 380 AgentRole that exists when the protocol finishes has its lifeline continued beyond the final arrow of the diagram.

381
 382 The lifeline may split into two or more lifelines to show AND/OR parallelism and decisions. Each separate track
 383 corresponds to a branch in the message flow. The lifelines may merge together at some subsequent point. The splitting
 384 of the lifeline for:

- 385
 386 AND parallelism starts at a horizontal heavy bar,
- 387
 388 OR parallelism (inclusive-or) starts at a horizontal heavy bar with a non-filled diamond, and,
- 389
 390 decision (exclusive-or) starts at a horizontal heavy bar with a non-filled diamond with "x" inside the diamond and is
 391 continued with a set of parallel vertical lifelines connected to the heavy bar.

392
 393 The merging is done the analogous way, that is, the parallel vertical lifelines stop at some of the horizontal heavy bars
 394 and one lifeline is continued from at the heavy bar.

396 3.2.3.3 Presentation Options

397 None.


399 3.2.3.4 Example

400

401

402

403



also

Section

3.2.1.4,

404 *Example.*

405

406 3.2.3.5 Mapping

407 See Section 3.2.1.5, *Mapping*.

408

409 **3.2.4 Threads of Interaction**

410 The sending of messages can be done either in parallel or as a decision between different communicative acts.
411 Receiving different communicative acts usually results in different behaviour and different answers, that is, the
412 behaviour of an AgentRole depends on the received message.

413

414 Adapted from [OMGuml], section 7.4.

415

416 3.2.4.1 Semantics

417 Since the behaviour of an AgentRole depends on the incoming message and different communicative acts are allowed
418 as an answer to a communicative act, the thread of interaction, that is, the processing of incoming messages, has to be
419 split up into different threads of interaction. The lifeline of an AgentRole is split and the thread of interaction defines the
420 reaction to received messages.

421

422 The thread of interaction shows the period during which an AgentRole is performing some task as a reaction to an
423 incoming message. It represents only the duration of the action in time, but not the control relationship between the
424 sender of the message and the receiver. A thread of interaction is always associated with the lifeline of an AgentRole.
425 Note we do not mean a physical thread in this context. The specification is independent of the implementation using
426 threads or other mechanisms.

427

428 3.2.4.2 Notation

429 A thread of interaction is shown as a tall thin rectangle whose top is aligned with its initiation time and whose bottom is
430 aligned with its completion time. It is drawn over the lifeline of an AgentRole. The task being performed may be labelled
431 as text next to the thread of interaction or in the left margin, depending on the style; alternately the incoming message
432 may indicate the task, in which case it may be omitted on the thread of interaction itself.

433

434 If the distinction between the reaction to different incoming communicative acts can be neglected, the entire lifeline may
435 be shown as one thread of interaction.

436

437 3.2.4.3 Presentation Options

438 Variation

439 A thread of interaction may can take only a short period of time. To simplify diagrams, for compactification reasons
440 of the diagram the parallelism and the decisions can be abbreviated by omitting the splitting/merging and putting the
441 different threads of interaction one after another on the lifeline.

442


443 Variation

444 A break of the rectangle describes a change in the thread of interaction.


445

446 3.2.4.4 Example

447



can be abbreviated as



448
449

450 3.2.4.5 Mapping
451 See Section 3.2.1.5, Mapping.
452

453 **3.2.5 Messages**

454 The main issue of protocols is the definition of communicative patterns, especially the sending of messages from one
455 AgentRole to another. This sending can be done in different ways, for example, with different cardinalities, depending
456 on some constraints or using AND/OR parallelism and decisions.
457

458 Adapted from [OMGuml], section 7.5 and section 8.9.
459

460 3.2.5.1 Semantics

461 A message or sending of a communicative act is a communication from one AgentRole to another that conveys
462 information with the expectation that the receiving AgentRole would react according to the semantics of the
463 communicative act. The specification of the protocol says nothing about the implementation of the processing of the
464 communicative act.
465

466 3.2.5.2 Notation

467 A message sending is shown as a horizontal solid arrow from a thread of interaction of an AgentRole to another thread
468 of interaction of another AgentRole. In case of a message is sent from an AgentRole to itself (note that there might be
469 many individual agents in an AgentRole), the arrow may start and end on the same lifeline or thread of interaction. Such
470 a nested thread of interaction is denoted by a thread of interaction that is shifted a little bit to the right side in the actual
471 thread of interaction.
472

473 Nested protocols are represented by a separate thread of interaction, along with an arrow initiating the nested protocol
474 and one or more arrows terminating the nested protocol. The initiating arrow is drawn starting with a small solid filled
475 circle, and a terminating arrow ends with a circle surrounding a small solid filled circle.
476

477 Each arrow is labelled with a message label that has the following syntax:

478 *predecessor guard-condition sequence-expression communicative-act argument-list*
479

480 Where:

481 *predecessor*

482 This consists of at most one natural number followed by a slash (/) defining the sequencing of a parallel construct
483 or the number of the input and output parameter in the context of Section 3.2.9, *Threads of Interaction and*
484 *Messages Inside and Outside Nested Protocols*, xxxx. The clause is omitted if the list is empty.
485
486

487 *guard-condition*

488 This is a usual UML guard condition, with the semantics, that the message is sent iff the guard is true. The guard
489 conditions must be defined on the behavioural semantics of the agents, that is, the internal state of the agent must
490 not be used in the definition of the guard.
491

492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507

sequence-expression

This is a constraint, especially with $n..m$ which denotes that the message is sent n up to m times with n , m $\{*\}$ ¹. The keyword `broadcast` denotes the broadcast sending of a message; the keyword `deadline` denotes a string that is encoded according to [ISO8601] and represents the deadline by which a message is useful.

communicative-act

This is either the name, that is, a string representation with an underlined name, of a concrete communicative act instance, the name of a concrete communicative act instance and its associated communicative act, written as *name:communicative-act* or only the name of the communicative act, for example, `inform`.

argument-list

This is a comma-separated list of arguments enclosed in parentheses. The parentheses can be omitted if the list is empty. Each argument is an expression in pseudo-code or an appropriate programming language or an OCL expression.

508 3.2.5.3 Presentation Options

509 Variation: Cardinality

510 The cardinality of a message (for example, n senders and m receivers of a message) is shown by writing natural
511 numbers at the beginning and at the end of the arrow. This variation is only allowed if the sender and/or receiver is
512 not an instance of an agent.

513
514 Variation: Asynchronous Message Passing

515 An asynchronous message is drawn with a stick arrowhead (\longrightarrow). It shows the sending of the message without
516 yielding control.

517
518 Variation: Synchronous Message Passing

519 A synchronous message is drawn with a filled solid arrowhead (\longrightarrow). It shows the yielding of the thread of control
520 (wait semantics), that is, the AgentRole waits until an answer message is received and nothing else can be
521 processed.

522
523 Variation: Time intensive Message Passing

524 Normally message arrows are drawn horizontally. This indicates the duration required to send the message is
525 atomic, that is, it is brief compared to the granularity of the interaction and that nothing else can take place during
526 the message transmission. That is the correct assumption within many computers. If the messages requires some
527 time to arrive for mobile communication, for example, during which something else can occur then the message
528 arrow may be slanted downward so that the arrowhead is below the arrow tail (\searrow).

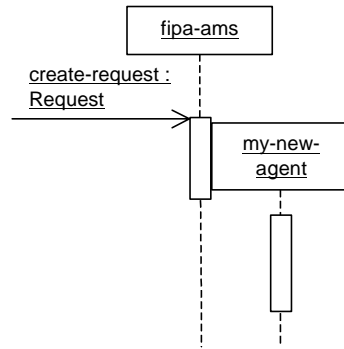
529
530 Variation: Repetition

531 The repetition of parts of a PD is represented by an arrow or one of its variations usually marked by some guards or
532 constraints ending at a thread of interaction which is according to the time axis before or after the actual time point.
533 Note, that in this case the time ordering on the PDs is violated.

534
535 3.2.5.4 Example

536

¹ The asterisk represents repetition an arbitrary number of times.



537
538

539 3.2.5.5 Mapping
540 See Section 3.2.1.5, Mapping.
541

542 **3.2.6 Complex Messages**

543 Besides the already presented kinds of messages, more complex messages can be used.
544

545 3.2.6.1 Semantics

546 A complex message may be the parallel sending of messages or exclusively sending of exactly one message out of a
547 set of different messages.
548

549 3.2.6.2 Notation

550 Three kinds of complex messages are distinguished. All three complex messages substitute an arrow from one thread
551 of interaction to another thread of interaction by an arrow starting at one thread of interaction ending either:

- 552 at a heavy bar (for AND parallelism),
- 553 at a heavy bar with a non-filled diamond (for OR parallelism; inclusive-or), or,
- 554 at a heavy bar with a non-filled diamond (for decisions; exclusive-or) with an "x" inside the diamond.

555 From these different kinds of heavy bars new arrows start in a right angle at the bar and end at possibly different
556 threads of interaction, which are possibly combined in a parallel or decisional way.


561 The merging of different messages is done in the analogous way, that is, the parallel horizontal message arrows stop at
562 one vertical bar and one message arrow is continued from the heavy bar.
563
564

565 3.2.6.3 Presentation Options

566 None.
567

568 3.2.6.4 Example

569



570
571
572

² This shows the restriction that request is sent before query.

573

574 3.2.6.5 Mapping

575 See *Section 3.2.1.5, Mapping*.

576

577 **3.2.7 Nested Protocols**

578 Nested protocols are applied to specify complex systems in a modular way. Moreover the reuse of parts of a
579 specification increases the readability of them.

580

581 A nested protocol can be defined and applied, if it is used several times within the same specification. In contrast to a
582 parameterised protocol it is only an abbreviation for a fixed (part of a) protocol. Additionally nested protocols are used
583 for the definition of repetition of a nested protocol according to guards and constraints.

584

585 Interleaved protocols show that between different agents a protocol is performed and more over in order to
586 finish/proceed the protocol an agent has to perform another protocol with other agents.

587

588 3.2.7.1 Semantics

589 If the nested protocol is marked with some guard then the semantics of the nested protocol is the semantics of the
590 protocol under the assumption that the guard evaluates to true, otherwise the semantics is the semantics of an empty
591 protocol, that is, nothing is specified.

592

593 If the nested protocol is marked with some constraints the nested protocol is repeated as long as the constraints
594 evaluate to true.

595

596 3.2.7.2 Notation

597 A nested protocol is shown as a rectangle with rounded corners. It may have one or more compartments. The
598 compartments are optional. They are as follows:

599

600 **Name compartment**

601 This holds the (optional) name of the nested protocol as a string. Nested protocols without names are anonymous.
602 It is undesirable to show the same named nested protocol twice in the same diagram except when they define the
603 same nested protocol. The compartment is written in the upper left-hand corner of the rectangle.

604

605

Guard compartment

This holds the (optional) guard of the nested protocol in the usual guard notation as [guard-condition]. Nested protocols without guards are equivalent with nested protocols with guard [true]. The guard compartment is written together with the constraint compartment in the lower left-hand corner of the rectangle.

Constraint compartment

This holds the (optional) constraint of the nested protocol in the usual constraint notation as {constraint-condition}. Nested protocols without constraints are equivalent with nested protocols with constraint {1}. The constraint compartment is written together with the guard compartment in the lower left-hand corner of the rectangle. In addition to the constraint condition used in UML the constraint $n..m$ denotes that the nested protocol is repeated n up to m times with $n = 0, m = *$.

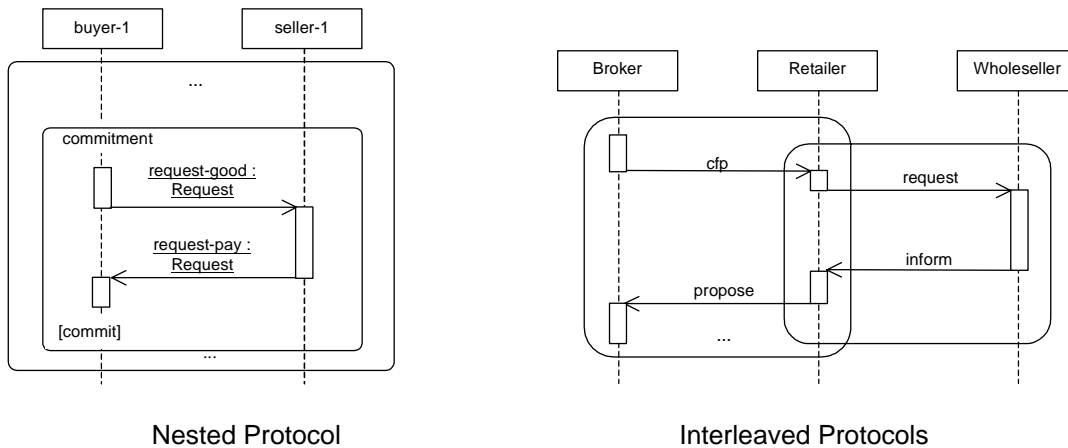
Another nested protocol can completely be drawn within the actual nested protocol denoting that the inner one is part of the outer one.

3.2.7.3 Presentation Options

The abbreviations n and $*$ can be applied to denote $n..n$ and $0..*$, respectively. Beyond the above usage of nested protocols for simple protocols, nested protocols can also be used applying parameterised protocols or instantiated parameterised protocols.

Another presentation option is the definition of interleaved protocols. For a nested protocol being part of another protocol the rectangle representing it has to be completely drawn within the other one. If interleaved protocols are defined, that is, during performing one IP another IP has to be processed, the rectangles are not drawn within each other.

3.2.7.4 Example



3.2.8 Complex Nested Protocols

Beyond the already presented nested and interleaved protocols, other kinds of complex nested protocols can also be defined.

3.2.8.1 Semantics

A complex nested protocol defines the parallel or decisional combination of nested protocols. It has to take into consideration the thread of interaction at the beginning and at the end of the complex nested protocol. Furthermore the starting and ending point within the nested protocols have to be considered.

645 3.2.8.2 Notation

646 Three kinds of complex nested protocols are distinguished. All three complex nested protocols are drawn over the
 647 lifeline and threads of interaction within a PD. Each individual nested protocol in a complex nested protocol is
 648 introduced by a line that is terminated by the rectangle of a nested protocol. These lines are connected either by:

649

650 a heavy bar defining AND parallelism,

651

652 a heavy bar with a non-filled diamond defining OR parallelism (inclusive-or), or,

653

654 a heavy bar with a non-filled diamond defining decisions (exclusive-or) with an "x" inside the diamond.

655

656 The threads of interaction which are continued in the different nested protocols are drawn as usual.

657


658 3.2.8.3 Presentation Options

659 None.

660

661 3.2.8.4 Example

662



663

664

665 3.2.8.5 Mapping

666 See *Section 3.2.1.5, Mapping*.

667

668 **3.2.9 Threads of Interaction and Messages Inside and Outside Nested Protocols**

669 Usually, nested protocols have input and output parameters, namely threads of interaction and messages.

670

671 3.2.9.1 Semantics

672 Nested protocols are defined in detail either within a PD where it is used or outside another PD. Threads of interaction
 673 and messages inside and outside nested protocols define the input and output parameter for nested protocols.

674

675 The input parameters are the threads of interaction, which are carried on in the nested protocol, and the messages
 676 which are received from other IPs.

677

678 The output parameters are on the one side the threads of interaction which are started within the nested protocol and
 679 are carried on outside the nested protocol and the messages which are sent from inside the nested protocol to
 680 AgentRoles not involved in the actual nested protocol. A message or thread of interaction ending at an input or starting
 681 at an output parameter of a nested protocol describes the connection of a whole PD with the embedded nested
 682 protocol.

683

684 3.2.9.2 Notation

685 The input and output parameters for the threads of interaction of a nested protocol are shown as a tall thin rectangle
 686 (like a thread of interaction) which is drawn beyond the bounds of over the top line and bottom line of the nested
 687 protocol rectangle, respectively.
 688


689 The input and output message parameters are shown by arrows starting with a small solid filled circle, and arrows
 690 ending at a circle surrounding a small solid filled circle (a bull's eye), respectively.

691 3.2.9.3 Presentation Options

692 The message arrows can be marked like usual messages. In this context, the predecessor denotes the number of the
 693 input/output parameter. The input/output thread of interaction can be marked with natural numbers to define the exact
 694 number of the parameter.
 695

696 3.2.9.4 Example

697

698
699

700 3.2.9.5 Mapping

701 See Section 3.2.1.5, *Mapping*.

702

703 **3.2.10 Parameterised Protocols**

704 Adapted from [OMGuml], section 5.11.

705

706 3.2.10.1 Semantics

707 A parameterised protocol is the description for an IP with one or more unbound formal parameters. It therefore defines
 708 a family of protocols, each protocol specified by binding the parameters to actual values. Typically the parameters
 709 represent agent roles, constraints, instances of communicative acts and nested protocols. The parameters used within
 710 the parameterised protocol are defined in terms of the formal parameters so they are become bound when the
 711 parameterised protocol itself is bound to the actual values.
 712

713 A parameterised protocol is not a directly-usable protocol because it has unbound parameters. Its parameters must be
 714 bound to actual values to create a bound form that is a protocol.
 715

716 3.2.10.2 Notation

717 A small dashed rectangle is superimposed on the upper right-hand corner of the rectangle with rounded corners
 718 when defining a nested protocol. The dashed rectangle contains a parameter list of formal parameters for the protocol.
 719 The list must not be empty, although it might be suppressed in the presentation. The name of the parameterised
 720 protocol is written as a string in the upper left-hand corner.
 721

722 The parameter list is a comma-separated list of arguments (formal parameters) defined by identifiers, like names for
723 AgentRoles, constraint expressions, communicative acts or nested protocol names.
724

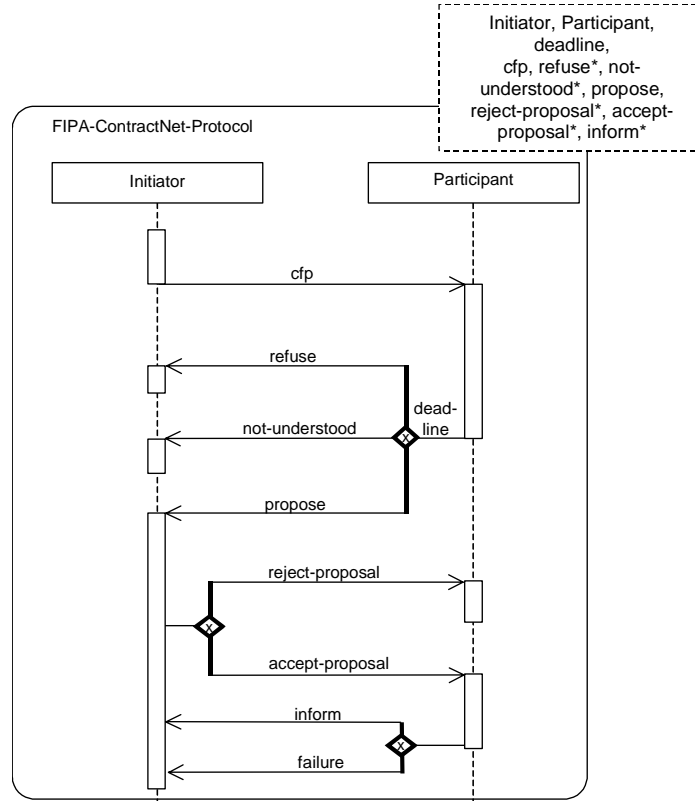
725 3.2.10.3 Presentation Options

726 The input/output parameters like messages and threads of interactions can be used and defined as for nested
727 protocols.

728 Communicative act can be marked with an asterisk in the parameter specification, denoting different kinds of messages
729 that can alternatively be sent in this context.
730

731 3.2.10.4 Example

732



733
734

735 3.2.10.5 Mapping

736 See Section 3.2.1.5, Mapping.

737

738 3.2.10.6 Comment

739 Note the difference between interleaved, nested and parameterised protocols. An interleaved protocol is used to show
740 that during the execution of one protocol another one is started/performed. Nested protocols are used to show
741 repetitions of sub-protocols, identifying fixed sub-protocols, reference to a fixed sub-protocol, like asking the DF for
742 some information, or guarding a sub-protocol. Parameterised protocols are used to prepare patterns which can be
743 instantiated in different contexts and applications, for example, the FIPA Contract Net Protocol for appointment
744 scheduling and negotiation about some good which should be sold.
745

746 3.2.11 Bound Elements

747 Adapted from [OMGuml], section 5.12.

748

749 3.2.11.1 Semantics

750 A parameterised PD cannot be used directly in an ordinary interaction description, because it has free parameters that
 751 are not meaningful outside of a scope that declares the parameter. To be used, a formal parameter of a parameterised
 752 protocol must be bound to actual values. The actual value for each parameter is an expression defined within the scope
 753 of use. If the referencing scope is itself a parameterised protocol, then the parameters of the referencing parameterised
 754 protocol can be used as actual values in binding the referenced parameterised protocol, but the parameter names in the
 755 two templates cannot be assumed to correspond, because they have no scope outside of their respective templates.
 756 We can assume without loss of generality that the parameter names of the different parameterised protocols are
 757 different.
 758

759 3.2.11.2 Notation

760 A bound element is indicated in the name string of an element, as follows:

761
 762 *parameterised-protocol-name* < *role-list*, *constraint-expression-list*, *value-list* >
 763

764 Where:

765
 766 *parameterised-protocol-name*
 767 This is identical to the name of the parameterised protocol.
 768

769 *role-list*
 770 This is a comma-delimited list of role labels. *constraint-expression-list* is a comma-delimited list of constraint terms.
 771

772 *value-list*
 773 This is a comma-delimited non-empty list of pairs, separated by a colon consisting of a value expression and a
 774 communicative act. The communicative act is optional.
 775

776 The number and types of the values must match the number and types of the parameterised protocol formal
 777 parameters for the parameterised protocol of the given name. The bound element name may be used anywhere that
 778 protocol of the parameterised kind could be used.
 779

780 3.2.11.3 Presentation Options

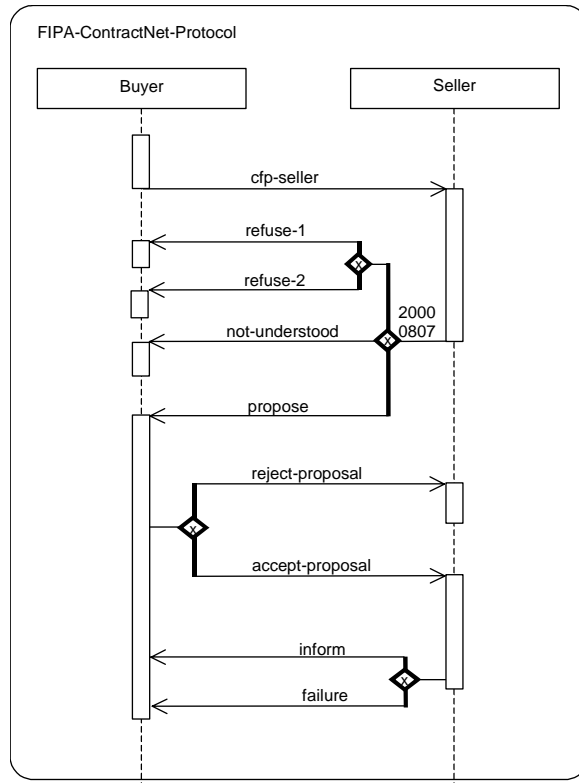
781 None.
 782
 783

783 3.2.11.4 Example

```

784
785 FIPA-ContractNet-Protocol
786 <
787   Buyer, Seller
788   20000807
789   cfp-seller : cfp,
790   refuse-1   : refuse,
791   refuse-2   : refuse, not-understood, propose, reject-proposal, accept-proposal,
792               cancel, inform, failure
793 >
794

```



795
796

797 3.2.11.5 Mapping

798 The use of the bound element syntax for the name of a symbol maps into a Binding dependency between the
799 dependent ModelElement corresponding to the bound element symbol and the provider ModelElement whose name
800 matches the name part of the bound element without the arguments. If the name does not match a parameterised
801 protocol or if the number of arguments in the bound element does not match the number of formal parameters in the
802 parameterised protocol, then the model is ill-formed. Each argument in the bound element maps into a ModelElement
803 bearing a templateArgument association to the Namespace of the bound element. The Binding relationship bears the
804 list of actual argument values.

805
806

806 4 References

- 807 [Booch94] Booch, G., Object-Oriented Analysis and Design with Applications. Benjamin/Cummings, 1994.
808 [Booch95] Booch, G., Object Solutions: Managing the Object-Oriented Project. Addison-Wesley, 1995.
809 [FIPA00007] FIPA Content Language Library Specification. Foundation for Intelligent Physical Agents, 2000.
810 <http://www.fipa.org/specs/fipa00007/>
811 [FIPA00023] FIPA Agent Management Specification. Foundation for Intelligent Physical Agents, 2000.
812 <http://www.fipa.org/specs/fipa00023/>
813 [ISO8601] Date Elements and Interchange Formats, Information Interchange – Representation of Dates and
814 Times, ISO 8601:1988(E), 1988.
815 [Odell2000] Odell, J., Parunak, H. van Dyke and Bauer, B., Extending UML for Agents. In: AOIS Workshop at AAAI,
816 2000.
817 [OMGuml] OMG Unified Modelling Language Version 1.1, Object Management Group, 1999.
818 <http://www.omg.org/uml/>
819 [Rumbaugh91] Rumbaugh, J., Blaha, M., Premerlani, W., Eddy, F. and Lorenzen, W., Object-Oriented Modeling and
820 Design. Prentice Hall, 1991.